LITERATURE REVIEW

# PROJECT - CLOUDLET

JARVIS MUTAKHA (MTKJAR001)

COMPUTER SCIENCE HONOURS
MAY 2014

**Abstract**

This document is a goes through the idea of cloudlet computing technology as a viable alternative to internet cloud computing technology by providing the advantages of cloud computing while minimizing its disadvantages (namely high network latency; the privacy, security and ownership of the information shared using these technologies) and reviews the relevant literature on the topics. The first half of the literatures reviewed mainly propose cloudlet technology as a way to offload computation from a computing device onto an available cloudlet using wireless local area networking (e.g. Wi-Fi) with the possibility of temporarily saving some data as well. The rest of the literatures talk about various psychosocial attitudes and behavior patterns of users who share digital information wirelessly while they are in close proximity to each other. The analysis of the psychosocial aspects could help in the design of the project.

**Cloud Computing**

Cloud computing is a global phenomenon. Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction (Mell & Grance, 2011). We make use of cloud services because they provide us with unique ways to collaborate, stay in touch, and engage with media socially. These services have been used to connect and make sharing easier for people across the globe. Cloud computing has enabled WhatsApp (WhatsApp Inc., 2014), a small company of 55 employees that operates a proprietary, cross-platform instant messaging subscription service for smartphones that uses the internet for communication, to scale to 18M users and handle spikes in demand. This is because instead of messages being transferred directly from mobile device to mobile device, messages are stored in the company's servers using the internet and passed along thus allowing messages to be accessed at anytime and anywhere (provided there is an internet connection).

Cloud computing technology also affects the design and implementation of wireless access points, whose purpose was to interconnect devices but now also connects them to the internet and the cloud. Currently it is not possible to create a wireless access point to share files locally on current mobile platforms without also sharing your 3G internet connection. A computer scientist could argue that using cloud services may not be the most efficient way of sharing files with the person physically next to you. This is mainly due to the fact that network latency could slow the whole process down considerably.

**Cloudlet Computing**

The same technologies that power WhatsApp, namely webservers, data stores, and databases, can also run at a smaller scale on embedded systems, such as the Raspberry Pi (Raspberry Pi Foundation, 2014). Computer hobbyists have revealed that the Raspberry Pi could make as an attractive solution for a small lightweight server because of its solid state storage, no noise, small form factor and low power consumption (Clay, 2014). So it seems even a Raspberry Pi extended with a battery, WiFi radio and local storage can provide a group of co-located friends with opportunities to share and engage with media, not with the public at large, but amongst themselves. Hence, the cloud is transformed from a global panopticon to a hyper-localized, ad-hoc instantiations

of the cloud: a cloudlet in short. To clarify, a cloudlet is a cloud that is not running on the internet but on a server(s) that are close in terms of number of network nodes a user has to cross (ideally the server would also be in close physical proximity, i.e. less than 100 metres, to the users of the cloudlet).

There are a handful of advantages cloudlets have over cloud technology which include: lower network latency and users having full ownership of the data shared. These formed as motivation for this project. This project is trying to implement cloudlet technology on computers with low processing power and storage capacity (e.g. Raspberry Pi) and enable file sharing between mobile devices.

This review will go through an evaluation of some of that literature and how the ides discussed in those papers apply to this idea. Some possible research questions include: are cloudlets a useful implementation of ICT to be used in conjunction with mobile devices? Are cloudlets practical in an urban/rural setting? Can cloudlets be implemented on cheaper computers and still be practical? Can the embedded system cloudlets be used for more than simply file sharing (e.g. computation)?

**Using Cloudlets for External Information Processing**

Mobile computing devices generally have less processing power and storage compared to their desktop and laptop counterparts. This is a major disadvantage for mobile devices especially if users want to run software that requires relatively high processing power and/or storage capacity. In 2009 a group of researchers proposed that mobile computing devices could be used to carry out "compute-intensive capabilities" for example, natural language processing; speech recognition; computer vision; machine learning; and augmented reality, but not at their state at the time of writing (Satyanarayanan, Bahl, Cáceres, & Davies, 2009). The reason being, in my opinion, that mobile computing devices don't have the computing capacity to carry out such tasks at the speed that would make them practical. According to the researchers they were "resource-poor" compared to less portable computing devices, namely desktop computers. They went on to suggest a new architecture where mobile devices use virtual machine technology to run on cloudlets that are within physical proximity of the device and are access via wireless local area network. Each cloudlet would contain at least one computer with a relatively high processing and memory that is connected to the internet (with reasonable speed) and is available for use by mobile devices within the area of the LAN.

Computationally intensive tasks currently are comparable to human quality if the tasks are done on powerful computers these include language translation (Carbonell et al., 2006) and facial recognition (Adler & Schuckers, 2007). Both can be practical applications if they could be used on mobile devices as well instead of non-portable computers. Both applications can be run on the cloud but they are only as effective as the speed of the connection the user has to the servers running them. Hence cloudlet technology could be a contender as a solution. Just like the motivation of this project they also saw the limitations of traditional web-based cloud solutions (the main limitation being the latency). The latency in wide area networks could upset the usability of systems that require cloud computation on the fly. Interactive response was shown to be inversely proportional to latency (Lagar-Cavilla et al., 2007). The main advantage of a cloudlet over the cloud is minimized latency due to a one hop jump between the mobile device and the network that does the computation.

In my opinion the researchers came to similar conclusions about using cloudlets to reduce these problems, specifically the speed of cloud-based applications could be increased if they were run on cloudlets that are closer to the user. The major differences being that they proposed using cloudlet technology to carry out computations and the cloudlet technology is implemented on relatively more powerful computers. This project is not attempting to carry out computations on cloudlets. Firstly, the cloudlets in this project are meant to be used as enabling technology for file transfer between mobile devices. Secondly, one of the research questions is to check if the cloudlets can be implemented on cheaper and computationally weaker embedded systems (namely the Raspberry Pi).

The pattern in reasoning seems to follow on with other researchers proposing the implementation of the computationally expensive MapReduce framework on Virtual Machines via cloudlets (Ibrahim, Jin, Cheng, Cao, Wu, & Qi, 2009). The main reason here was MapReduce had large overheads in I/O virtualization and managing storage and computation proved difficult. The virtual machine is built on multiple physical machines that carry out complex computations that are shared amongst them.

**Using Cloudlets for Information Exchange and External Information Storage**

There are other questions project would also like to answer for example how cloudlets will actually work and which structure & features are more likely to be preferred by users. Obviously, this question cannot be answered until the system is implemented and deployed for practical testing. But there is some literature that asks similar questions pertaining file and media sharing amongst users on mobile networks locally.

A group of researchers looked into various methods in which co-located people share information with each other. The conclusion that was found was an intuitive albeit a subtle one: people's sharing patterns changed according to the situation they were in even if they were sharing with the same people (Reitmaier, Benz, & Marsden, 2013). This point was supported with the Mobiphos study they conducted whereby photos are automatically shared to people in a predefined group. The people's photo sharing behaviors changed according to who was in the group (Clawson, Voida, Patel, & et al., 2008). I think psychosocial behavior like this has to be taken into account especially for this project because I suspect the success of the implementations may hinge heavily on this.

Another important social issue that I believe has to be considered is privacy. Reitmaier and Benz's study on young people's sharing habits on a Bluetooth based quasi-cloudlet pointed out the importance of privacy and identity management in mobile co-located media sharing (Reitmaier, Benz, & Marsden, 2013). Participants in the study mentioned some issues which included: who sees what; who has control over who sees what; how they limit or specify what they are sharing and with or to whom; who owns the content. Although this particular study was based on Bluetooth technology, some of the privacy, security and ownership issues can easily be applied to wireless local area network cloudlets. It seems that cloudlets could provide an alternative to cloud-based technology to users who want to maintain information ownership.

**Conclusion**

Unfortunately a comprehensive comparison of these previous implementations of cloudlets and studies of co-located information sharing cannot be done because they are essentially different. While one set of studies focused on implementation of cloudlets to intensive computation, the others focused on file transfer between users on mobile phones on other technology other than

cloudlet technology. The use of VM-based cloudlets in carrying out heavy computations for mobile devices and the use of single cloudlets to host virtual machines to do map and reduce calculations are very similar to each other. However, their goals are different to this project. The only real similarity is the MapReduce proposal used the cloudlet for temporary storage of variables.

The psychosocial studies on the other hand have more similarities to this project because even though each of the studies used different technology the social aspects seemed to remain the same.

There seems to be more literature and projects on the use of cloudlets as an alternative to web-based clouds for computationally expensive applications compared to other applications, with the key motive being cloudlets offer lower latency compared to web-based clouds. Granted, cloudlet technology is fairly new so the more obvious applications may be more likely to be implemented and studied first. Less obvious applications, I believe, include implementing cloudlet technology using smaller and weaker computers. However, using cloudlets for file sharing I think is a more obvious application.

Furthermore using cloudlets could prove as a tangible solution to the growing concerns of privacy and ownership of information stored in web-based clouds. Cloudlets may provide the advantages of the web-based clouds (i.e. high power computation at the user's disposal; and access to information over a wide area) while simultaneously curbing the major disadvantage of security.

**References (APA 5th Edition)**

Adler, A., & Schuckers, M. (2007). Comparing Human and Automatic Face Recognition Performance. *Institute of Electrical and Electronics Engineers Trans. Systems, Man, and Cybernetics, 37*(5), 1248-1255.

Carbonell et al., J. (2006). Context-Based Machine Translation. *7th International Conference of the Association for Machine Translation in the Americas* (pp. 19-28). Boston: Proceedings of the Association for Machine Translation of the Americas (AMTA-2006).

Clawson, J., Voida, A., Patel, N., & et al. (2008). Mobiphos: A Collocated-Synchronous Mobile Photo Sharing Application. *MobileHCI*, 187-195.

Clay, C. (2014, January 16). *Raspberry Pi: 11 reasons why it's the perfect small server | ZDNet*. Retrieved May 14, 2014, from ZDNet | Technology News, Analysis, Comments and Product Reviews for IT Professionals: http://www.zdnet.com/raspberry-pi-11-reasons-why-its-the-perfect-small-server-7000025206/

Ibrahim, S., Jin, H., Cheng, B., Cao, H., Wu, S., & Qi, L. (2009). CLOUDLET: Towards MapReduce Implementation on Virtual Machines. *HPDC* (pp. 1-2). Munich, Germany: Accociation for Computing Machinery.

Lagar-Cavilla et al., H. (2007). Interactive Resource-Intensive Applications Made Easy. *Proc. Middleware 2007: ACM/IFIP/Usenix 8th Int'l Middlewae Conf.* (pp. 143-163). Springer.

Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing.* Gaithersburg: U.S. Department of Commerce National Institute of Standards and Technology.

Raspberry Pi Foundation. (2014, May 06). *Raspberry Pi*. Retrieved May 06, 2014, from Raspberry Pi: http://www.raspberrypi.org/

Reitmaier, T., Benz, P., & Marsden, G. (2013). Designing and Theorizing Co-Located Interactions. *CHI 2013* (pp. 1-10). Paris, France: Association of Computing Machinery.

Satyanarayanan, M., Bahl, V., Cáceres, R., & Davies, N. (2009). The Case for VM-Based Cloudlets in Mobile Computing. *Pervasive Computing, IEEE CS, 8*(4), 2-11.

WhatsApp Inc. (2014, May 06). *WhatsApp :: Home*. Retrieved May 06, 2014, from WhatsApp :: Home: http://www.whatsapp.com/