

Fully Charged: Characterizing the Data Requirements of Services and Apps to Optimize the Use of Energy

Jorge Andrés Larracochea^{1,3}, Philippe Roose¹, Sergio Ilarri^{2,3}

¹ LIUPPA, E2S / University of Pau, 64600, Anglet, France, France

² Instituto de Investigación en Ingeniería de Aragón (I3A)

³ Universidad de Zaragoza, Mariano Esquillor s/n, 50018, Zaragoza, España

Tel. +34-976762707, e-mail: silarri@unizar.es

Introduction

Currently, the rate of adoption of ICT (Information and Communication Technologies) devices has increased at an impressive rate. This is not only due to the sales of consumer devices such as smartphones, but also due to the widespread adoption of mobile and IoT (Internet of Things) devices with the purpose of gathering more knowledge about our environment. Energy is consumed in these devices in a different manner according to their available limited resources (such as battery) in an attempt to fulfill tasks imposed by the services' processing and data demands. This has led ICT researchers to find interest in strategies for mitigating high energy consumptions when processing and transferring data. For example, efforts on delivering efficient software code that preserves energy by being architecturally efficient and conscious about resource usage have been done in the past, but we consider that other approaches could be implemented in order to further increase energy savings.

Services Operations as an Energy Saving Strategy

In contrast to traditional architectures, performing operations with services, such as whole/partial service duplication, migration etc. in current services architectures (such as a microservices approach) could be used to more efficiently manage data and services' processing. This is of a particular interest to us due to previous research proving that shortening the distance between the client and the data storage location results in less energy consumption (Schien et al. 2012) and, other

scenarios where migrating/duplicating demanding services/data to a part of the network with better resource availability could save energy. Furthermore, data/services localization and treatment strategies are still to be developed according to awareness of software/applications contextual and architectural nuances.

For example, heavy computational processes like AI computation and graphics rendering for real-time strategy games could be done on a server in the cloud, in such a way that the client is only responsible for user interaction to receive input. An inconvenience with this approach is that, under the context of a real-time strategy game, the user input relies heavily on a reactive response to what is shown on the screen. This means that the graphical rendering service cannot be separated at all from the client if a QoS constraint is found in the architecture of the application. This example is also useful to demonstrate the importance of energy consumption in data transfer during the deployment of services. Video game streaming services need to offer, on average, higher FPS (frames per second) than films (which are mostly shot at 24 FPS). This means that graphic rendering in the cloud would put an extra strain on the network and consume more bandwidth, and finally more energy in the transfer. Bringing the rendering service closer to the client could be of our interest to save energy within data management strategies.

The Relevance of Application Taxonomy in Data and Processing Management

Even though some proposals have been made in the past in order to classify applications with a

rating according to their power consumption and functions (Hindle 2016), there are still several hurdles that hinder the way to classify them, such as application updates, user behaviour during apps' usage, data consumption, and data generation (Peltonen et al. 2015).

Previous works have proposed a taxonomy for applications based on the behaviour of their data and service characteristics (Nickerson et al. 2009), but further research needs to be invested in order to define a taxonomy for the behaviour of data and processing in relation to the intrinsic characteristic of applications' architectures. This would allow us to generalize behaviours in order to apply pertinent energy saving strategies depending on the existing requirements and context.

Future Goals and Approaches

We consider that our research will lead us towards several goals:

- 1- The creation of a taxonomy for applications from the perspective of their software architecture.
- 2- Identification of strategies to apply on services and/or data to save energy according to the proposed taxonomy.
- 3- Implementation and validation of algorithms that use the aforementioned identified strategies.
- 4- A Domain Specific Language (DSL) for eco-responsible apps with services orchestration according to our algorithms.

Therefore, we aim at analyzing and categorizing a wide range of situations and example applications that may have an impact on the most suitable decisions to be taken regarding the needs of the service/application and the expected energy consumption. We will then analyze all these situations from an architectural perspective, considering the most important metrics to optimize in each case (e.g., latency, throughput) and the existing requirements (e.g., degree of user interactivity

and data consistency required), in order to propose general guidelines of possible strategies to apply in each case, such as the displacement or duplication of services and/or data when appropriate. Furthermore, we will implement and evaluate the results exhibited by testing the algorithms generated with the strategies previously defined. Finally, we will develop a DSL for eco-responsible apps in order to provide orchestration of services according to the algorithms validated, in such a way that energy savings can be considered from the start of the design in an easier way.

References

Hindle, A., 2016. Green Software Engineering: The Curse of Methodology. *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. S.l.: s.n., pp. 46-55. DOI 10.1109/SANER.2016.60.

Nickerson, R., Muntermann, J., Varshney, U. y Isaac, H., 2009. Taxonomy Development In Information Systems: Developing A Taxonomy Of Mobile Applications. *HAL, Working Papers*

Peltonen, E., Lagerspetz, E., Nurmi, P. y Tarkoma, S., 2015. Energy modeling of system settings: A crowdsourced approach. *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. S.l.: s.n., pp. 37-45. DOI 10.1109/PERCOM.2015.7146507.

Schien, D., Preist, C., Yearworth, M. y Shabajee, P., 2012. Impact of location on the energy footprint of digital media. *2012 IEEE International Symposium on Sustainable Systems and Technology (ISSST)*. S.l.: s.n., pp. 1-6. DOI 10.1109/ISSST.2012.6228017.