

Interconnecting Smart Grids and Clouds to save Energy

Anne-Cécile Orgerie
CNRS, IRISA, Rennes, France

Keywords: Energy Efficiency, Cloud Computing, Smart Grids.

Abstract: Cloud computing is becoming an essential component for Internet services. However, its energy consumption has become a key environmental and economic concern. The distributed nature of Cloud infrastructures involves that their components are spread across wide areas, interconnected through different networks, and powered by diverse energy sources and providers, making overall energy monitoring and optimizations challenging. In this paper, we present the opportunity brought by the Smart Grids to exploit renewable energy availability and to optimize energy management in distributed Clouds. The presence of smart sensors which are both integrated into the electricity Grid and connected to the Internet, indeed offers for the first time the possibility of exploiting the availability of various energy sources, and of making complete energy measurements of all the Cloud resources – computing, storage and especially networking resources – problems which have previously been intractable.

1 INTRODUCTION

Global warming, oil reserve limits and ever-growing electricity consumption are all powerful motivations for designing more energy-efficient systems. The energy consumed to power Information and Communication Technologies (ICT) has become a major concern. Data centers and network infrastructures represent an ever-growing part of this consumption (Mills, 2013). As an example, for 2010, Google used 900,000 servers and consumed 1.9 billion kWh of electricity (Kooimey, 2011). Electricity becomes the new limiting factor for deploying data center equipment. Indeed, the energy consumption of state-of-the-art servers grows inexorably as they embed more and more powerful cores and advanced features and technologies. So, electricity consumption is hampering technological advances. It also raises major environmental, economic and societal concerns. Curbing this consumption is thus crucial and urgent.

In this paper, we mainly aim at investigating how to exploit the capabilities of Smart Grids technologies in the context of Cloud computing in order to save energy. This has the potential for reducing global energy consumption and thus, inducing real economic, environmental, and societal benefits. On the Cloud's side, we also aim at increasing the user's energy-awareness and the Cloud's energy-efficiency which are the main levers needed to trigger substantial involvement from

Cloud providers and designers in order to reduce the electricity demand, and contribute to a more sustainable ICT industry.

On the generation side, it is estimated that 10% of electric energy produced by power plants is currently lost during transmission and distribution to consumers, with 40% of these losses occurring on the distribution network (Feng et al., 2009). As an example, in 2006 in the United-States, the total energy losses and distribution losses were about 1,638 billion and 655 billion kWh, respectively (Feng et al., 2009). Smart Grids provide the means to control the energy supply more efficiently and to dynamically manage peak load. We argue that their utilization will greatly improve understanding of the energy consumption of Clouds and will help to reduce it significantly.

The remainder of the paper is organized as follows. Section 2 presents the context and motivation, which is further discussed over a canonical example on Section 3. Section 4 illustrates the problem complexity. The related work is presented in Section 5. Section 6 concludes and presents our future work.

2 CONTEXT AND MOTIVATION

Recently, the maturity of virtualization techniques has led to the emergence of Cloud Computing, which is a new paradigm becoming increasingly essential

to Information Technology (IT) services. It aims to provide dynamic, reliable, customized and QoS-guaranteed environments to end-users. By putting together individual requirements and by gathering the needed equipment, Cloud providers can efficiently manage and offer virtually unlimited resources, minimizing the costs incurred by organizations when providing Internet services.

To offer better QoS and support more users, Clouds are distributed: they consist in several data centers geographically distributed. Computer networks are the crucial elements that interconnect the data centers providing the computing and storage capacities of these distributed Clouds. As the size of Clouds increases and their traffic demands diversify, computer network resources, inside and in-between the data centers, are often stretched to their limits and, in many cases, become a performance bottleneck. Besides, they represent a non negligible part of the energy consumption of these distributed systems (Orgerie et al., 2014).

Advances in distributed systems have historically been related to improving their performance, scalability and quality of service. Yet, it is now urgent to drastically increase the efficiency of large-scale distributed systems in order to curb the rising energy consumption of ICT (Information and Communications Technologies). Indeed, a recent study focused on Clouds (Mills, 2013) estimates that the ICT ecosystem now approaches 10% of world electricity generation. The authors promise a faster growth in ICT energy use, pointing out the information appetite of Big Data, big networks and big infrastructure which unavoidably leads to big power.

Smart Grid sensors are the monitoring devices of Smart power Grids. They provide high fidelity control through high-speed, two-way communication, sensing, and real-time coordination of all assets down to the customer meter and the end-user devices (Pratt et al., 2010). As they are literally located at the intersection between the electrical and communication networks, smart sensors enable improvement in the stability, efficiency, flexibility, and reliability of electricity services. Through their communicating sensors, they offer better support of renewable energy and more comprehensive views of the consumption of infrastructures. As for the Smart Grid actuators, the operating counterpart of smart sensors, they can act directly on the devices to enforce energy-efficient policies (Hledik, 2009). For instance in smart homes, they can adjust air conditioning level according to the parameters detected by the smart sensors: the temperature, the operation of appliances producing heat and human presence.

While the design of smart actuators is under study, smart sensors starts to be deployed at the scale of entire cities (like in Lyon, France, with the smart city project for instance). This deployment comes along with a Big Data challenge since the collected data represent huge amounts, coming from distributed places, and requiring cross-computations. While Clouds come naturally to the rescue of Smart Grids for dealing with this big data issue, little attention has been paid to the benefits that Smart Grids could bring to Clouds.

We propose exploiting Smart Grid technologies to come to the rescue of energy-hungry Clouds. Unlike in traditional electrical distribution networks, where power can only be moved and scheduled in very limited ways, Smart Grids dynamically and effectively adapt supply to demand. This recent technology offers the unique chance to monitor the energy consumption in real-time of entire and distributed infrastructures through their smart sensors, to reduce energy use through their smart actuators, and to act as the bridge enabling hitherto impossible joint energy production/consumption synergies.

3 A CANONICAL EXAMPLE

When a user submits a Google request (such as a search for instance), it is routed to a data center that processes it, computes the answer and sends it back to the user. Google owns several data centers spread across the world. For performance reasons, the center answering the user's request is more likely to be the one closest to the user. However, this data center may be supplied by a far away energy source. A more distant data center may be using a renewable energy source which happens to be close by and available at the time of the request. This request may have consumed less energy, or a different kind of energy (renewable or not), if it had been sent to this further data center. In this case, the response time would have been increased but maybe not noticeably: a different trade-off between performance and energy-efficiency could have been adopted.

However, two issues make this trade-off impossible to determine: 1) the cloud (hosting Google's services here) is not able to estimate the energy consumption of a user's request because it implies various and distributed resources (several servers to process the request and communication networks to carry the request and the answer), and 2) the cloud has no means of knowing where the energy for its data centers comes from, and whether it is sufficiently available at a given time. Such data could be provided by




Figure 1: Representation of the involved resources and layers on the example of a worldwide Cloud.

the smart sensors included in Smart Grid infrastructures. Smart sensors and Cloud's knowledge about its utilization could be combined to take near real-time allocation decisions for the Cloud's resources, and energy management optimizations through the smart actuators for the electrical Grid.

4 ILLUSTRATION OF THE PROBLEM COMPLEXITY

Figure 1 shows an example of a large-scale Cloud infrastructure. Its hardware layer is constituted of two physical components: data centers to process the user requests and communication network resources (i.e. routers) to transport the requests and to link the data centers together. Electricity, which is also a physical resource, is distributed through the electrical network from power stations to data centers and network resources.

But, as shown in Figure 1, the electrical network and the communications network of a Cloud infrastructure are not directly interconnected. So, the Cloud manager has no access to the electrical information. This implies two inherent problems for the Cloud manager which aims to optimize the energy consumption of the entire Cloud infrastructure: 1) it has no way to know the consumption of the networking resources as they belong to other (numerous) companies; 2) it has no access to the electricity availability in the power stations used by its resources.

As illustrated in Figure 1, Smart Grids constitute the autonomic management layer that enables cooperation among the electrical network and the communication network. Our goal is to make the Cloud layer and the Smart Grid layer cooperate in order to save energy in the Cloud infrastructures (with its underly-

ing physical resources) and in the electrical network. As pointed out in the figure, the communication network will play a key role in this cooperation as it physically stands in between.

Smart Grids have been designed to be deployed at large scale in order to improve the efficiency of the production and distribution of electricity in a comprehensive context. Although some smart sensors (home smart meters for instance) are already available, smart actuators are still under development. This currently leaves Smart Grids with a measurement role only (exploited for a smarter accounting for instance), and considerably limits the power of *Smart Grids*, which are here simply able to collect more detailed monitoring information than regular power grids. Thus, the current status of deployed Smart Grid technologies is not allowing them yet to dynamically adapt supply to demand. Nonetheless, we believe that they will strengthen in the forthcoming years in order to answer to the urgent need of smarter infrastructures for saving energy.

In this paper, we limit the investigation of the potential impact of Smart Grids to distributed Cloud infrastructures, impact which has not been estimated yet to the best of our knowledge. Interconnecting Cloud computing infrastructures and Smart Grid technologies is not straightforward due to the complexity of these two systems. It requires to face multi-parameter modeling and multi-objective optimization challenges with distributed algorithmic and real-time constraints. A fine-grained modelization of the considered Cloud applications is also needed in order to be able to accommodate users' expectations in terms of QoS and in order to anticipate the impact of infrastructure reconfigurations and optimization algorithms.

5 RELATED WORK

Large-scale distributed systems – such as computing data centers, Grids, Clouds and dedicated networks – consist of vast collections of computing and storage resources interconnected through a network. These systems, with different levels of scalability, interoperability and respect to user constraints, have become the building blocks for numerous applications and services ranging from weather forecast to web search.

Cloud infrastructures consist in geographically distributed data centers which are linked through communication networks (Wang et al., 2008a). With the emergence of the Future Internet and the dawning of new IT models such as cloud computing, the usage of data centers (DC), and consequently their power consumption, increase dramatically. As an example, for 2010, Google used 900,000 servers and consumed 1.9 billion kWh of electricity (Kooimey, 2011). Other big Cloud companies present similar figures and similar issues (Katz, 2009).

Electricity becomes the new limiting factor for deploying data center equipment. Indeed, the energy consumption of state-of-the-art servers grows inexorably as they embed more and more powerful cores and advanced features and technologies. So, electricity consumption is hampering technological advances. It also raises major environmental, economic and social concerns. Curbing this consumption is thus crucial and urgent.

The first way to save energy at a data center level consist in locating it close to where the electricity is generated, hence minimizing transmission losses. For example, Western North Carolina, USA, attracts data centers with its low electricity prices due to abundant capacity of coal and nuclear power following the departure of the region's textile and furniture manufacturing (Greenpeace, 2011). As of writing, this region has three super-size data centers from Google, Apple and Facebook with respective power demands of 60 to 100 MW, 100 MW and 40 MW (Greenpeace, 2011).

Other companies opt for greener sources of energy. For example, Quincy (Washington, USA) supplies electricity to data facilities from Yahoo, Microsoft, Dell and Amazon with its low-cost hydro-electrics left behind following the shutting down of the region's aluminum industry (Greenpeace, 2011). Several renewable energy sources like wind power, solar energy, hydro-power, bio-energy, geothermal power and marine power can be considered to power up super-sized facilities.

In spite of these approaches, numerous data facilities have already been built and cannot be moved. Cloud infrastructures, on the other hand, can still take

advantage of multiple locations to use green sources of energy with approaches such as follow-the-sun and follow-the-wind (Figuerola et al., 2009). As sun and wind provide renewable sources of energy whose capacity fluctuates over time, the rationale is to place computing jobs on resources using renewable energy, and migrate jobs as renewable energy becomes available on resources in other locations.

Within the data center itself, a range of technologies can be utilized to make cloud computing infrastructures more energy efficient, including better cooling technologies, temperature-aware scheduling (Fan et al., 2007), Dynamic Voltage and Frequency Scaling (DVFS) (Snowdon et al., 2005), and resource virtualization (Talaber et al., 2009). The use of VMs (Barham et al., 2003) brings several benefits including environment and performance isolation; improved resource utilization by enabling workload consolidation; and resource provisioning on demand. Nevertheless, such technologies should be analyzed and used carefully for really improving the energy-efficiency of computing infrastructures (Miyoshi et al., 2002).

Concerning wired networking resources, which constitute an important part of Cloud infrastructures, the energy-efficient management techniques can be divided into two primary categories (Orgerie et al., 2014; Bolla et al., 2011): sleeping techniques making use of future foreseen on/off capabilities of network devices, and rate adaptation techniques. The most emblematic solution of the sleeping techniques is named Low-Power Idle (Christensen et al., 2010), standardized under the name 802.03az, and starting to be implemented on Network Interface Cards and switches available on the market. The precursor of rate adaptation techniques is named ALR (adaptive Link Rate) (Gunaratne et al., 2006), and proposes to adapt the negotiated link rates (between the two link end devices) to the actual load.

As detailed in a survey (Orgerie et al., 2014), numerous works target resource allocation and task scheduling with energy constraints for both computing (Lefèvre and Orgerie, 2010; Giori et al., 2011), and networking (Baliga et al., 2011; Bolla et al., 2011; Betti et al., 2013) resources separately. However, for both of these disjointed propositions, they either focus on energy-efficient goals (ie. reducing the amount of consumed energy), or on energy-aware objectives (ie. maximizing the amount of consumed renewable energy). Only few works consider both aims at the same time, like it is the case in this position paper; and none considered both aims for both kind of resources: computing and networking.

The Cloud workload varies widely on a short time

scale (Elson and Howell, 2008). To quickly and decentrally adapt Cloud resource allocation to renewable energy availability, efficient real-time autonomic reconfiguration mechanisms are required. Centralized solutions have been proposed to this issue (Orgerie et al., 2014). Yet, their scalability is limited, and thus, their applicability to Cloud infrastructures is arguable. In order to solve this technical barrier, we argue for the use of Software-Defined Networking techniques. These techniques allow for virtual networks decoupling the control (routing messages for instance) and data planes (users' messages) (Jain and Paul, 2013). OpenFlow is currently the mostly used SDN technology (Lara et al., 2014).

Network virtualization is also a promising solution to enable energy savings (Orgerie et al., 2014). Virtual Routers On the Move (VROOM), a network management primitive, pioneering research work, is presented in (Wang et al., 2008b). It simplifies the management of routers by virtualizing them and allowing them to freely move from one physical node to another, thus avoiding logical topology changes. The migration of virtual routers can also be used to save energy (Chen and Phillips, 2013).

On the electricity side, it is estimated that 10% of electric energy produced by power plants is currently lost during transmission and distribution to consumers, with 40% of these losses occurring on the distribution network (Feng et al., 2009). As an example, in 2006 in the United-States, the total energy losses and distribution losses were about 1,638 billion and 655 billion kWh, respectively (Feng et al., 2009). Smart Grids provide the means to control the energy supply more efficiently and to dynamically manage peak load. Their utilization has the potential to greatly improve understanding of the energy consumption of Clouds and to help to reduce it significantly.

To take advantage from this new technology, we need to deal with three major issues: 1) complexity incurred by the proliferation of energy sensors, the fact that different systems are interconnected (communication networks and distribution networks), these systems are highly multi-scale by nature, different technologies are employed (different manufacturers), and a lot of parameters are involved (energy sources and availability, location, etc.); 2) interdependability because our solution relies on the collaboration between electricity networks and communication networks; 3) performance of the system: it is mandatory for the monitoring infrastructure to provide real-time analysis in spite of the large data volumes to process, and for the Cloud management system to integrate distributed management and failure recovery mechanisms.

Current works on Smart Grids are dealing with smoothing energy peak demand (Tang et al., 2013), improving the effective utilization of intermittent and fluctuating renewables, and improving the scalability of the proposed solutions (Fang et al., 2012). Our goal here is not to contribute to the development and technical advances of Smart Grids. It only consists in utilizing this emerging and innovative technology.

6 CONCLUSION AND FUTURE WORK

Most of the work on energy-efficient management of Clouds does not consider networking resources as they go beyond the limits of the data centers. Thus, they are not entirely managed by one provider and their consumption is not accessible to the cloud manager. Yet, the intrinsic energy consumption of the Google request includes transporting the request to a data center, computing the request and transporting the answer back to the user. Taking into account only the computing costs leads to proposing optimization algorithms that do not optimize the overall energy consumption, but only the consumption of one data center. Indeed, network resources account for a non-negligible part in the energy equation (Orgerie et al., 2014; Orgerie et al., 2011). Making Smart Grids and Cloud managers collaborate could solve these issues and lead to important energy savings.

Our future work will focus on exploiting this energy saving potential resulting from the interconnection of Smart Grids and Clouds. This challenging goal will require to propose implementable distributed reconfiguration mechanisms and task scheduling algorithms for the allocation of Cloud resources (both networking and computing resources) taking into account energy efficiency and utilization of renewable energies. It will also led us to design an efficient management of electricity distribution and utilization in the context of distributed energy sources linked by Smart Grids and used to power a distributed Cloud infrastructure. Working towards this objective should both increase the knowledge about these two complex systems – namely Cloud infrastructures and Smart Grids – and remove technological barriers towards their interconnection for coping with the big power issue.

REFERENCES

- Baliga, J., Ayre, R., Hinton, K., and Tucker, R. (2011). Green Cloud Computing: Balancing Energy in Pro-

- cessing, Storage, and Transport. *Proceedings of the IEEE*, 99(1):149–167.
- Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., and Warfield, A. (2003). Xen and the Art of Virtualization. In *ACM Symposium on Operating Systems Principles (SOSP)*, pages 164–177.
- Betti, G., Amaldi, E., Capone, A., and Ercolani, G. (2013). Cost-aware optimization models for communication networks with renewable energy sources. In *IEEE Conference on Computer Communications Workshops (INFOCOM Workshops)*, pages 25–30.
- Bolla, R., Bruschi, R., Davoli, F., and Cucchietti, F. (2011). Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures. *IEEE Communications Surveys and Tutorials*, 13(2):223–244.
- Chen, X. and Phillips, C. (2013). An evolutionary based dynamic energy management framework for ip-over-dwdm core networks. In Pierson, J.-M., Da Costa, G., and Dittmann, L., editors, *Energy Efficiency in Large Scale Distributed Systems*, Lecture Notes in Computer Science, pages 233–247. Springer Berlin Heidelberg.
- Christensen, K., Reviriego, P., Nordman, B., Bennett, M., Mostowfi, M., and Maestro, J. (2010). IEEE 802.3az: the road to Energy Efficient Ethernet. *IEEE Communications Magazine*, 48(11):50–56.
- Elson, J. and Howell, J. (2008). Handling Flash Crowds from Your Garage. In *USENIX Annual Technical Conference, ATC'08*, pages 171–184.
- Fan, X., Weber, W.-D., and Barroso, L. A. (2007). Power provisioning for a warehouse-sized computer. In *ACM International symposium on Computer architecture (ISCA)*, pages 13–23.
- Fang, X., Misra, S., Xue, G., and Yang, D. (2012). Smart Grid: The New and Improved Power Grid: A Survey. *IEEE Communications Surveys Tutorials*, 14(4):944–980.
- Feng, X., Peterson, W., Yang, F., Wickramasekara, G., and Finney, J. (2009). Smarter grids are more efficient. ABB review.
- Figuerola, S., Lemay, M., Reijs, V., Savoie, M., and St. Arnaud, B. (2009). Converged Optical Network Infrastructures in Support of Future Internet and Grid Services Using IaaS to Reduce GHG Emissions. *Journal of Lightwave Technology*, 27(12):1941–1946.
- Goiri, I., Le, K., Haque, M., Beauchea, R., Nguyen, T., Guitart, J., Torres, J., and Bianchini, R. (2011). GreenSlot: Scheduling energy consumption in green datacenters. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 1–11.
- Greenpeace (2011). How dirty is your data? Greenpeace report.
- Gunaratne, C., Christensen, K., and Suen, S. (2006). Ethernet Adaptive Link Rate (ALR): Analysis Of A Buffer Threshold Policy. In *IEEE Global Telecommunications Conference (GLOBECOM)*, pages 1–6.
- Hledik, R. (2009). How green is the smart grid? *The Electricity Journal*, 22(3):29–41.
- Jain, R. and Paul, S. (2013). Network virtualization and software defined networking for cloud computing: a survey. *IEEE Communications Magazine*, 51(11):24–31.
- Katz, R. H. (2009). Tech Titans Building Boom. *IEEE Spectrum*, 46(2):40–54.
- Koomey, J. (2011). Growth in Data Center Electricity Use 2005 to 2010. Analytics Press.
- Lara, A., Kolasani, A., and Ramamurthy, B. (2014). Network Innovation using OpenFlow: A Survey. *IEEE Communications Surveys Tutorials*, 16(1):493–512.
- Lefèvre, L. and Orgerie, A.-C. (2010). Designing and Evaluating an Energy Efficient Cloud. *The Journal of Supercomputing*, 51(3):352–373.
- Mills, M. (2013). The Cloud Begins With Coal - Big Data, Big Networks, Big Infrastructure, and Big Power. Technical report, Digital Power Group.
- Miyoshi, A., Lefurgy, C., Van Hensbergen, E., Rajamony, R., and Rajkumar, R. (2002). Critical power slope: understanding the runtime effects of frequency scaling. In *ACM International conference on Supercomputing (ICS)*, pages 35–44.
- Orgerie, A.-C., de Assunção, M. D., and Lefèvre, L. (2014). A Survey on Techniques for Improving the Energy Efficiency of Large Scale Distributed Systems. *ACM Computing Surveys*, 46(4).
- Orgerie, A.-C., Lefèvre, L., and Guérin-Lassous, I. (2011). On the Energy Efficiency of Centralized and Decentralized Management for Reservation-Based Networks. In *IEEE GLOBECOM*.
- Pratt, R., Balducci, P., Gerkenmeyer, C., Katipamula, S., Kintner-Meyer, M., Sanquist, T., Schneider, K., and Secrest, T. (2010). The Smart Grid: An Estimation of the Energy and CO2 Benefits. Technical report, Pacific Northwest National Laboratory (PNNL), U.S. Department of Energy.
- Snowdon, D., Ruocco, S., and Heiser, G. (2005). Power Management and Dynamic Voltage Scaling: Myths and Facts. In *Workshop on Power Aware Real-time Computing*.
- Talaber, R., Brey, T., and Lamers, L. (2009). Using Virtualization to Improve Data Center Efficiency. Technical report, The Green Grid.
- Tang, S., Huang, Q., Li, X.-Y., and Wu, D. (2013). Smoothing the energy consumption: Peak demand reduction in smart grid. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 1133–1141.
- Wang, L., Tao, J., Kunze, M., Castellanos, A., Kramer, D., and Karl, W. (2008a). Scientific Cloud Computing: Early Definition and Experience. In *IEEE International Conference on High Performance Computing and Communications (HPCC)*, pages 825–830.
- Wang, Y., Keller, E., Biskeborn, B., van der Merwe, J., and Rexford, J. (2008b). Virtual routers on the move: live router migration as a network-management primitive. *ACM SIGCOMM Computer Communication Review*, 38(4):231242.