

Energy Sustainability in Cooperating Clouds

Antonio Celesti, Antonio Puliafito, Francesco Tusa and Massimo Villari

Università degli Studi di Messina, Facoltà di Ingegneria, Contrada di Dio, S. Agata, 98166 Messina, Italy

Keywords: Cloud Computing, Federation, Virtual Infrastructure Management, Energy Sustainability, Photovoltaic Systems, Renewable Energy, Energy Optimization.

Abstract: Nowadays, cloud federation is paving the way toward new business scenarios in which it is possible to enforce more flexible energy management strategies than in the past. Considering independent cloud providers, each one is exclusively bounded to the specific energy supplier powering its datacenter. The situation radically change if we consider a federation of cloud providers each one powered by both a conventional energy supplier and a renewable energy generator. In such a context the opportune relocation of computational workload among providers can lead to a global energy sustainability policy for the whole federation. In this work, we investigate the advantages, constrains, and issues for the achievement of such a sustainable environment.

1 INTRODUCTION

Federation is the next frontier of cloud computing. Throughout the federation, different small and medium Cloud providers belonging to different organizations, can join each other to achieve a common goal, usually represented by the optimization of their resources.

The basic idea is that a Cloud provider has not infinite resources. In order to achieve target business scenarios a Cloud provider may need a flexible infrastructure. Federation allows Cloud providers to achieve such a resilient infrastructure asking additional resources to other federation-enabled Cloud Providers. Cloud federation is much more than the mere use of resources provided by a mega-provider.

From a political point of view, the term federation refers to a type of system organization characterized by a joining of partially “self-governing” entities united by a “central government”. In a federation, each self-governing status of the component entities is typically independent and may not be altered by a unilateral decision of the “central government”.

Besides cloud mega-providers, also smaller/medium providers are becoming popular even though the virtualization infrastructures they have deployed in their datacenters cannot directly compete with the bigger counterparts. A way to overcome these resource limitations is represented by the promotion of federation mechanisms among small/medium cloud providers. This allows to pick up the advantages of

other form of economic model considering societies, universities, research centers and organizations that commonly do not fully use the resources of their own physical infrastructures.

In this work, we focus on an innovative sustainable federated cloud scenario in which resources are relocated between cloud providers whose datacenters are partially powered by renewable energy generator systems. The federation is seen as a way for reducing energy costs (**Energy Cost Saving**), but at the same time a possibility to reduce the CO_2 emissions (**Energy Sustainability**). Here we discuss strategies and policies it is possible to apply in federated cloud environments for achieving the just mentioned goals. Specifically our assessment is aimed at the design of an *Energy Manager*, to be included in whichever Virtual Infrastructure Manager as well OpenStack, OpenNebula and Cloud Stack.

The manuscript is organized as follows. Section 2 introduce how an energy sustainability strategy can be applied to a federated cloud environment. The energy consumption of a datacenter is affected by different factors including the Power contribution for the Information Technology (IT) equipment (P_{IT}), the Power contribution for the Electrical (POW) equipments (P_{POW}), and the Power contribution for the Cooling (COOL) equipments (P_{COOL}). To this regards several energy considerations about cloud datacenters are discussed in Section 3. As said before, a decision algorithm for sustainable federated clouds is presented in Section 4. Section 5 discusses related works.

Section 6 summarizes conclusions and lights to the future.

2 CLOUD FEDERATION AND ENERGY SUSTAINABILITY

Federation brings new business opportunities for clouds. In fact, besides the traditional market where cloud providers offer cloud-based services to their clients, federation triggers a new market where cloud providers can buy and/or sell computing/storage capabilities and services from/to other clouds. A cloud provider can decide to lend resources to other clouds when it realizes that its datacenter is under-utilized at given times. Typically, datacenters are under-utilized during the night and over-utilized during the morning. Therefore, as the datacenter cannot be turned off, the cloud provider may decide to turn the problem into a business opportunity.

As federation enables cloud providers to relocate their services on other peers belonging to the system, in our opinion, it is possible to carry out more flexible energy-aware scenarios than the past, when we considered independent non-federated clouds. Two possible alternative energy-aware scenarios are: **Energy Cost Saving** and **Energy Sustainability**.

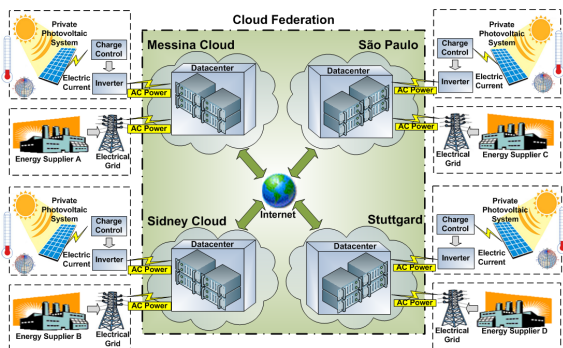


Figure 1: Example of Sustainable Federated Cloud Environment.

The main contribution of this work is to propose a possible approach for the achievement of such an environment. Our approach is based on the following idea: “*moving the computation toward the more sustainable available cloud datacenter*”. This statement is motivated by the following assumptions:

1. Often, the renewable energy generator systems produce more energy than necessary.
2. It is very hard to store the exceeded produced renewable energy (e.g., in batteries).

3. Alternatively, it is becoming very difficult to put the exceeded produced renewable energy in public electric grids. This practice is becoming a problem for Energy suppliers as it implies uncontrolled power surges which are hard to be managed. As this problem is becoming more and more sensitive, the energy suppliers are becoming to be reluctant to absorb energy produced by private renewable energy generator systems.
4. As consequence of 1), 2), and 3) often the exceeded produced renewable energy is wasted.
5. Consequently, it is easier to move the computation toward a datacenter powered by a renewable energy generator system with a high large availability of energy than moving the “green energy” toward the datacenter where the computing has to take place.

If we consider a set of datacenters with these features, a sustainable federated cloud environment can allow to save money, maximizing the use of “green energy” and reducing the level of carbon dioxide.

An example of such a scenario is depicted in Figure 1. The sustainable federated cloud ecosystem includes four cloud providers: Messina, Sidney, São Paulo, and Stuttgart. The electricity suppliers A, B, C, and D are independent companies that provide energy at different costs. In our scenario, the datacenter of each cloud providers is mainly powered by a private renewable energy generator system as primary source of energy. When the primary source of energy is not enough to power the datacenter, the cloud uses the energy of its Electricity Supplier. In addition, each cloud is federated with each other in order to take the advantages of service relocation and resource consolidation. Further details regarding how federate cloud architecture are out of scope of this work. Further details can be found in Section 5.

Each cloud provider joins the federation in order to move its computation in other federated clouds where the production of “green energy” is maximum. In simple terms, we move the computation load towards the more efficient renewable energy generators (in term of produced electricity) maximizing the utilization of the federated clouds in which the workload has been transferred. Considering Figure 1, the four clouds have different latitudes. According to a given period (due to time zone and month), each renewable energy generator system that primarily powers each cloud datacenter can have different energy efficiency compared to each other. These different conditions can depend by different factors according to the adopted source of renewable energy. For example the amount of energy produced by a photovoltaic system depends on the solar radiance which is differ-

ent hour by hour, day by day, and month by month. In addition, the energy production of a photovoltaic system is also affected by the weather and climate conditions. For simplicity, let us consider the time zone, when the time of Messina is 17:00, the time of São Paulo is 13:00. In this situation, in São Paulo the solar radiance of the sun is stronger than the one in Messina. On the other hand, considering both cities in July, the temperature of Messina will be probably higher than the one in São Paulo. Further similar consideration can be made considering the latitude of the two cities. This scenario implies that if a cloud provider wants to enforce energy sustainability policies on its own datacenter, it relocate its services into other federated cloud providers, chosen according to the aforementioned energy consideration. For reasons of Quality of Service (QoS), let us suppose that each cloud provider of the federation has replicated in advance part of its services into other federated clouds. Considering a federation of Infrastructure as Service (IaaS) clouds, service replication means copying Virtual Machines (VMs) disk-image into other federated providers. In this way the cloud providers that wants to apply energy sustainability policies can turn off the blade center hosting its VMs and turn on the copies of these VMs pre-arranged into other federated cloud datacenters, where the renewable energy production is maximum according to temperature, latitude, and time zone.

3 POWER CONSUMPTION CONSIDERATIONS OF A DATACENTER

The first step for the achievement of a Sustainable Cloud Federation is to better understand the main factors affecting the total power consumption of a datacenter. As already introduced these factors are P_{IT} , P_{POW} , and P_{COOL} .

P_{IT} is related to the total power consumption of the IT equipment such as: CPUs, Storage (i.e., Hard Disk, Tapes, Optical Disks, etc.), RAM, Switches and Router, Monitors.

P_{POW} regards the total power consumption of the Electrical equipment, for example, including: UPS (Uninterruptible Power Supply), PSU (Power Supply Unit), PDU (Power Distribution Unit), Cable (copper wires characterized by an electrical resistance), Lights, Batteries.

P_{COOL} refers cooling equipment including for example: Chiller. responsible for making the GAP among the external (outdoor) and internal (indoor) tempera-

tures, FANs, regarding the Control Room Air Conditioning (CRAC) or to equipment used to discard the heat in the external ambient, Pumps, responsible for moving the refrigerant substance (or water) inside the distribution pipes, Valves, Unit of Control.

The entire cooling system of a datacenter can be referred also as HVAC (i.e., Heating, Ventilating, Air-Conditioning) or HVAC(R) (Heating, Ventilating, Air-Conditioning, and Refrigerating). Consequently, the total power consumption of a datacenter can be defined as:

$$P_{TOT} = P_{IT} + P_{POW} + P_{COOL} \quad (1)$$

Figure 2 shows the total amount of energy consumption of a datacenter. The percentages of the total power spent in a datacenter can be roughly distributed as follows:

$$P_{IT} = 50\%; P_{POW} = 20\%; P_{COOL} = 30\%. \quad (2)$$




Figure 2: Typical Consumption of Energy inside a Datacenter.

P_{IT} and P_{POW} are strongly related to the transistors performances. In fact, currently, they have a physical limits that it is not possible to be overcome. However, recent studies are trying to break such limits and the expectation is that future innovation can bring to more performance equipment from the point of view of the energy consumption. In this direction, a recent and interesting dissertation was conducted in *The Optimist, the Pessimist, and the Global Race to Exascale in 20 Megawatt* (Tolentino and Cameron, 2012). Considering the aforementioned assumption, particular consideration deserves P_{COOL} . We believe that P_{COOL} will have a big role in the energy consumption studies in the ICT field. In fact, at present P_{COOL} is the parameter easier to optimize and how it is described later in this manuscript, cloud federation can help to achieve this goal.

In order to model a sustainable federated cloud environment, we consider the datacenter of each cloud provider as a black box which acts an ideal refrigeration machine also known as the *Carnot Engine*. The performances of this ideal model is only affected by the temperatures: the black box needs to be cooled

as much as depending on the environment where it is placed. Moreover we consider that in such a black box there are both *energy* coming in, and *heat* that has to be discarded in the external environment.

Considering the First Law of Thermodynamic about the conservation of energy, we have the following equation (see Eq. 3) that states any component/device connected to the Electric Grid transform the energy from one form to other ones (Conservation of Energy).

$$P_{in} = W + Q \quad (3)$$

Where P_{in} is the input power over the time (i.e., in hours), W is the energy spent for mechanical works and Q is the energy release as Heat. In a datacenter, P_{in} is the electric power delivered through copper cables, W is the energy for producing movements (Compressor of a Chiller, FANs, Hard Disk with rotors, Optical Readers, etc.). Finally Q is the Heat produced by components, lights, motors, compressors. In a datacenter, if we assume the P_{IT} and/or P_{POW} , the contribution to W is negligible. The Compressor in a chiller (P_{COOL}) catch a lot of energy, but its work is useful for expelling Heat from inside the datacenter to the outside (environment). Using the theoretical analysis of the Thermodynamic model, it is possible to demonstrate that for a Carnot Engine, the Energy Q is linked to the Temperature T .

The measurement of goodness of a datacenter is given by a number called **Power Usage Effectiveness (PUE)**. It is expressed as the ratio from the total amount of energy consumed as input respect to good part of energy used for IT computations. Values of PUE equals to 1 correspond to an energy efficiency of 100%.

$$PUE = \frac{P_{TOT}}{P_{IT}} \quad (4)$$

The increasing of the PUE value, corresponds to a greater weight of either P_{POW} or P_{COOL} contributions (or both). Typical PUE values for a datacenter are greater than 1 and corresponds to 2-2.5.

Looking at Figure 3, it is possible to know how a datacenter cooling system works. Although the Figure refers to an actual datacenter existing in Messina, it may represent the general installation of a working environment. The Figure shows two graphs (in the top and bottom part of the DC) highlighting how the temperature ranges with both Free Cooling and HVAC Plants respect to the real distance from the datacenter (*HeatPath*). Looking at the Free Cooling situation in particular, it is possible to remark that the temperature of the BladeCenter ($T_{chassis}$) should be guaranteed according to the external temperature (T_{env}). This can be accomplished only when the climate conditions of a site allow this configuration. HVAC and Free Cool-




Figure 3: The representation of a Datacenter: it is the real installation of the DC in Messina. The APC CUBE uses the HVAC cooling.

ing performances are both environment *Temperature* dependent.

Hence, considering a sustainable federated cloud environment the PUE of two different cloud datacenters with the same equipment, but placed in two different regions, may assume different values. The environment temperature of each region affects the resulting PUE depending on the adopted cooling techniques.

4 DECISION ALGORITHM

Considering both the scenario previously described in Section 2 and the considerations already pointed out in this work, in the following we are going to introduce a simple algorithm designed to address the problem of establishing the best site on which a given service has to be deployed for minimizing both energy consumption and costs, on the strength of the actual environmental data collected on each.

We assume a given workload has to be executed in our four-sites federated scenario in a given time: we know how many resources will be needed to accomplish that task and how many free computational slots will be available on each site. For each one, we also assume to know the availability of the instantaneous amount of electrical power produced from the photo voltaic equipment, the amount of electrical power absorbed from the HVAC(R) (or free cooling system) and that one used for achieving the computation (there will be also other contribution that will be explained in the following).

If we want to find a method for optimizing computation respect to costs, we have to identify an analytic

approach able to put together all the parameters characterizing the scenario. Since energy providers apply their fares on the actual amount of consumed kWh, our analytic model will take into account the energy contributions as mean value of the electrical power over a time interval. In order to have a good snapshot of the energy production/absorption when the computing element placement have to be performed, in our case this interval corresponds to one hour.

Looking at the scenario depicted in Figure 1, and paying attention to the aforementioned considerations, we might assume the energy consumption of each site is represented from the Eq. 5

$$E_{GRID}(T, G, s) = [E_{IT}(s) + E_{COOL}(T) + E_{POW} - E_{PV}(T, G)] \quad (5)$$

where each energy term represents the produced/absorbed medium power in a time interval of one hour and is expressed in kWh. $E_{GRID}(T, G)$ is the amount of power grid energy needed from a given site and is related to the other energy contribution terms appearing in the formula. According to this latter, it will depend on the external temperature T where the site operates, the G factor describing the availability of the renewable green energy source (e.g. sun radiation, wind intensity, etc.) and the s parameter associated to the number of computational unit allocated (e.g. virtual machines) on that site.

In this specific work, we are assuming each datacenter can rely on a green energy source. For the sake of simplicity, we can restrict our considerations to a scenario where each site takes advantage of just renewable photovoltaic energy. As consequence, we can associate the G factor to the sun radiation factor of the energy plant of a given site. According to the above mentioned assumption, The remaining terms of the expression in order are: E_{IT} , a non-constant factor associated to the energy needed from a data center to perform the computation for a given service that depends on the number s of computational slot to be allocated; $E_{COOL}(T)$ an energy factor associated to the HVAC(R) system (or free cooling system) that depends on the external temperature T characterizing the area where the site is working (it is related to the measured PUE for that site); E_{POW} is the constant factor associated to the energy consumption of the power supply equipment of the data center; finally the last term $E_{PV}(T, G)$ is a function of both external temperature T and sun radiation factor G , related to the mean energy amount made available from the photo voltaic system of that site in the time interval of one hour.

If we consider that each site retrieves its electrical power from a different provider, we might estimate the costs needed in terms of electrical power to

achieve a workload execution according to the applied fares for kWh. With this assumption and starting from the Eq. 5, you can obtain a new expression 6 related to the energy expenses associated to each site functioning:

$$C_{GRID}(T, G, s, c) = E_{GRID}(T, G, s) \cdot c \quad (6)$$

In order to cut costs, the organization of our scenario relying on four different computational sites, can select the one for which the cost values associated to expression 6 is minimum for a given set of parameters (i.e. temperature T , sun radiation G , needed computational slots s and fare f). To achieve this goal, we can consider data related to those parameters is collected from a sensor network on each site periodically. Depending on these obtained values a table is built considering the mean values of the retrieved parameters.

Taking into account such an approach, this particular module of the VIM operating on the datacenters, will take care of computing the costs for each site relying on the physical measurements collected from the available sensors and stored in the table. An example is reported in Table 1. A simple algorithm implemented within the *Energy Manager* will be able to choose the most convenient site (in terms of energy consumption) where allocate the computation associated to a given service.

If two different sites are characterized from the same cost values in a given time, the algorithm implemented within the *Energy Manager* will evaluate also the amount of photo voltaic energy made available from each site, finally preferring the one where this factor is the greatest. In some situations where energy sustainability is preponderant on cost optimization, the same algorithm could be applied in a different way: the site(s) where the photo voltaic energy production is the greatest is first chosen and, in the case of more sites producing the same photo voltaic energy amount, from the retrieved set, the site where the overall costs are the lowest will be finally selected for allocating the services.

The algorithm could be implemented creating a complementary module for the VIM (*Energy Manager*) that retrieves needed data from the sensor network available in each computational site, computing the associated values of energy contributions (considering the instantaneous electrical power values retrieved in that moment) and storing them on the table until the next data refresh (after one hour). We designed the *Energy Manager* having in mind the possibility to include it in whichever Virtual Infrastructure Manager as well OpenStack, OpenNebula and Cloud Stack. When one of these VIMs has to allocate a new set of s virtual machines for a given

Table 1: Data retrieved from 4 different sites that will be given as input for the VIM *Energy Manager*.

Site	Temperature (T [°C])	Sun Radiation (G [MJ/m ²])	Energy Grid Fare (c [\$])	Photo Voltaic Energy E_{PV} [kWh]	Slots (s)	PUE	Costs [\$]
Site 1	35	20	0.08	100	120	3.7	10
Site 2	30	18	0.09	150	90	3.2	13
Site 3	18	14	0.07	80	70	1.2	10
Site 4	23	15	0.08	80	75	2.5	15

service, together with snapshot of physical resource availability of each site (reported in Table 1), it will also invoke the *Energy Manager* to retrieve information on the most convenient site to which deploy the allocation either in terms of cost minimization or energy sustainability. Since each site can offer a limited number of computational slots, if the virtual machine number needed for a given service are greater than the maximum availability for a site, the load will be split across more sites still considering the satisfaction of the same requirements (cost minimization or energy sustainability).

Looking at Table 1, if we suppose executing a service that needs 100 virtual machines, in the first case of cost optimization, the site selected by the *Energy Manager* will be one between *Site 1* or *Site 3* (as they guarantee the lowest energy costs: respectively 10 \$ and 13 \$). The final choice will lead to *Site 1* since his photo voltaic energy availability (100 kWh) is greater than the one offered by *Site 3* (80 kWh). The E_{PV} availability in this situation is preponderant: although the temperature in *Site 3* is 18 °C and allows free cooling as refrigeration methods, the C_{GRID} costs coming from Eq. 6 are still more convenient on *Site 1* where the “free cost” energy is offered by the photo voltaic equipment. Furthermore, *Site 1* has the availability of $s = 120$ computational slots and is able to directly satisfy the requested service demand. Otherwise the computation would be split among *Site 1* and *Site 3*.

Still looking at the same table in the alternative situation we have mentioned before (energy sustainability optimization), the first set of selected sites will be formed by *Site 3* and *Site 4* (either offering the same amount of $E_{PV} = 80$ kWh). This time, the final choice will fall on *Site 3* as it is able to offer lower grid energy costs than *Site 4* (10 \$ against 15 \$). Differently from the previous case, the available computational slots of this site is lower than the needed ones. In this case, 70 of the requested VMs will be deployed on *Site 3* while the remaining ones on *Site 4*.

As reported in the table, the HVAC(R) system of each data center is characterized in terms of efficiency through the *PUE* (the values in the table refers to a mean value of the coefficient in the time interval of

one hour). High *PUE* values are associated to better refrigerator systems that allows to use less electrical power to push out heat from the data center. The *PUE* values are tightly related the $E_{COOL}(T)$ values contributing in Eq. 5. The best efficiency in terms of energy spent for cooling is achieved on the site 3, where thanks to the low external environmental temperature, it is possible to use the free cooling technique thus reaching a $PUE = 1.2$.

5 RELATED WORKS

In the section hereby, the early part analyzes works falling into energy saving and green energy topics aimed at datacenter. While in latter part several works dealing with cloud and federation are reported.

The work we highlight just below show as the problematic we are trying to address is an *hot* topic indeed. Many works dealing with datacenters and sustainability exist in the scientific literature, however our contribution tries to give an answer in the are of cooperating clouds.

The work in (Wang et al., 2011) highlights an innovative cooling strategy that leverages thermal storage to cut the electricity bill for cooling. The authors claimed the system does not cause servers in a datacenter to overheat. They worked on Computational Fluid Dynamics (CFD) to consider the realistic thermal dynamics in a datacenter with 1120 servers.

A Workload Distribution for Internet datacenters is proposed in (Abbasi et al., 2010), where the server provisioning algorithm is aware of the temperature distribution in a DC. The authors try to find a way where the utilization constraints (in term of *capacity* and *performance* constrains) are satisfied and energy consumption is minimized.

Modeling a thermal behavior of a datacenter is a challenging work due to the high number of physical parameters need to be considered. An interesting model along with a close-loop control system is described in (Zhou and Wang, 2011). The authors assessed a datacenter with many CRACs. The *inlet* temperature of many racks is investigated for accom-

plishing the Partition in Zone of a datacenter for an efficient decentralized control.

5.1 Cloud and Federation

In this paragraph, we provide an overview of currently existing solutions in the field of Cloud Federation, taking into account initiatives born in academia and major research projects. Most of the work in the field concerns the study of architectural models able to efficiently support the collaboration between different cloud providers focusing on various aspects of the federation.

In our previous work (Celesti et al., 2010) we describe an architectural solution for federation by means of a Cross-Cloud Federation Manager (CCFM), a software component in charge of executing the three main functionalities required for a federation. In particular, the component explicitly manages: i) the discovery phase in which information about other clouds are received and sent, ii) the match-making phase performing the best choice of the provider according to some utility measure and iii) the authentication phase creating a secure channel between the federated clouds.

In (Buyya et al., 2010), the authors propose a more articulated model for federation composed of three main components. A Cloud Coordinator manages a specific cloud and acts as interface for the external clouds by exposing well-defined cloud operations. The Cloud Exchange component implements the functionality of a registry by storing all necessary information characterizing cloud providers together with demands and offers for computational resources.

The dissertation in (Kiani et al., 2012) describes the *large-scale context provisioning*. The authors remarked that the adoption of context-aware applications and services has proved elusive so far, due to multi-faceted challenges in cloud computing area. Indeed existing context aware systems are not ideally placed to meet the domain objectives, and facilitate their use in the emerging cloud computing scenarios. The use of a predominant focus upon designing for static topologies of the interacting distributed components. Presumptions of a single administrative domain or authority and context provisioning within a single administrative, geographic or network domain.

6 CONCLUSIONS

Nowadays, a sensitive problem is finding the right combination between high performance datacenter and energy sustainability. In this work, considering a

scenario of cloud federation, we proposed a methodology for enabling sustainable cooperating clouds. Considering photovoltaic energy generation systems, our approach is based on an energy and temperature-driven strategies in which the computation workload of a cloud is moved toward the most efficient sustainable federated cloud. According to such a strategy and considering a federated CLEVER-based scenario, we defined an algorithm for the management of VM allocation according to energy and temperature-driven policies. In future works, we plan to consider also heterogeneous cooperating clouds.

REFERENCES

- Abbasi, Z., Varsamopoulos, G., and Gupta, S. K. S. (2010). Thermal aware server provisioning and workload distribution for internet data centers. In *HPDC*, pages 130–141.
- Buyya, R., Ranjan, R., and Calheiros, R. N. (2010). Intercloud: Utility-oriented federation of cloud computing environments for scaling of application services. In *Proceedings of the 10th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP 2010)*, pages 21–23. Springer.
- Celesti, A., Tusa, F., Villari, M., and Puliafito, A. (2010). Three-phase cross-cloud federation model: The cloud sso authentication. In *Proceedings of the 2010 Second International Conference on Advances in Future Internet*, AFIN '10, pages 94–101, Washington, DC, USA. IEEE Computer Society.
- Kiani, L., Anjum, A., Bessis, N., and Hill, R. (2012). Large-scale context provisioning. In *2012 Sixth International Conference on Complex, Intelligent, and Software Intensive Systems, CISIS 2012*.
- Tolentino, M. E. and Cameron, K. W. (2012). The optimist, the pessimist, and the global race to exascale in 20 megawatts. *IEEE Computer*, 45(1):95–97.
- Wang, Y., Wang, X., and Zhang, Y. (2011). Leveraging thermal storage to cut the electricity bill for datacenter cooling. In *Proceedings of the 4th Workshop on Power-Aware Computing and Systems, HotPower '11*, pages 8:1–8:5, New York, NY, USA. ACM.
- Zhou, R. and Wang, Z. (2011). Modeling and control for cooling management of data centers with hot aisle containment. In *ASME 2011 International Mechanical Engineering Congress & Exposition*.