

Energy Packet Networks: Adaptive Energy Management for the Cloud

Invited Paper

Erol Gelenbe
Intelligent Systems & Networks
Dept. of Electrical & Electronic Engineering
Imperial College
London SW7 2BT, UK
e.gelenbe@imperial.ac.uk

ABSTRACT

In the presence of limitations in the availability of energy for data centres, especially in dense urban areas, a novel system that we call an *Energy Packet Network* is discussed as a means to provide energy on demand to Cloud Computing servers. This approach can be useful in the presence of renewable energy sources, and if scarce sources of energy must be shared by multiple computational units whose peak to average power consumption ratio is high. Such a system will use energy storage units to best match and smooth the intermittent supply and the intermittent demand. The analysis of such systems based on queueing networks is suggested and applied to a special case for illustration.

Keywords

Cloud Computing, Energy Packet Networks, Energy Provisioning on Demand, Smart Grid

1. INTRODUCTION

The Cloud offers dynamic provisioning of computing and networking resources to applications that can thus be executed, on demand and in a distributed manner at the best possible cost and quality of service. Middleware can help users dynamically locate and select the most effective Cloud services that meet their needs, and Cloud service providers can compete among themselves to meet the users' needs in the most cost-effective way. While the concentration of Cloud services in data centres that can provide these services cheaply is an attractive option, such concentrated centres have overall energy needs that are often prohibitive [18]. In fact, when a data centre is installed its present and future energy needs have to be planned in advance.

Such installations are already becoming quite difficult to install in large urban areas in Europe such as Paris or London where the computational needs are greatest. Similarly,

CO₂ taxes are already deterring investments in such larger centres [39]. An alternative approach for large urban areas is to distribute computing power over a large number of smaller sites which will be operating asynchronously, and to dispatch *energy on demand* when it is required by computations, rather than to guarantee a high level of power availability all the time.

This leads us to the concept of *Energy Packet Networks* (EPN) which are integrated adaptive electrical energy storage, distribution and consumption systems that we proposed recently [33, 41]. In addition to the conventional scheme for distributing energy based on instantaneous flow of current towards points of energy consumption, EPNs offer the smart request and dispatching of units of energy called *Energy Packets* (EP) to meet the demands of computing consumers such as Cloud users, as well as other electrical devices and appliances. In EPN, smart dynamic generation and storage combine with smart on-demand request and dispatching of electric power. Such systems are particularly well adapted to environments where renewable energy sources are common, and where effective means for storing energy, such as electric cars and uninterruptible power supplies, are available.

In the case of energy provisioning for the Cloud, such a system will include geographically distributed renewable energy sources, conventional backup sources of energy that originate from fossil and nuclear power plants, and a distribution network, together with distributed energy storage facilities and the small or large data centres that create demand and that are provisioned with energy on demand. Many of the consumption sinks may also be coupled at close distance with one or more storage facilities.

The flow of energy in the EPN will be controlled by *Smart Energy Dispatching Centres* (SEDCs) which receive requests from both the consumers (e.g. Cloud servers) and from storage centres that wish to be replenished. The SEDCS then optimise the energy flows by making the best use of renewable energy availability, storage centre state, and existing pricing policies, while satisfying the demands and minimising peak energy flows through buffering and scheduling. SEDCs will obviously be computer control centres which receive information and make dispatching decisions from/to energy switches via data communication networks. The basic unit of energy in an EPN, an EP, can be viewed as a pulse of power that lasts a certain time; it constitutes the basic energy delivery unit of our system, say in KWH. The energy content of such a packet must be small enough to be

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CloudCP 2012: 2nd International Workshop on Cloud Computing Platforms, Bern, Switzerland
Copyright 2012 ACM 978-1-4503-1161-8 ...\$10.00.

close to the smallest energy needs of consumers, e.g. running a particular job or group of jobs at a data centre, and yet large enough to be measurable and billable as a significant and useful quantity.

Although we not detail the approach in this paper, our vision is that the SEDCs will use distributed adaptive schemes such as those developed for smart packet network routing [8, 17]. Thus when a request for cloud computing arises, the SEDCs will not only select the computational resource which can do this at lowest energy and/or economic cost, but will also dynamically provision the energy dynamically. Similar ideas have been suggested for energy-aware network routing in recent papers [26, 27].

2. A MODEL FOR EPNS

Although the system we consider is intuitively appealing, its value resides in its potential as a better means of delivering electric power for applications such as data centres and clouds whose energy consumption goes hand in hand with their computational activities. Thus we need to model EPNs and evaluate their performance advantages through system analysis. The paradigm of queueing networks is particularly favorable for studying such systems thus in this section we will develop a model representation for EPNs. The theoretical framework that we will use is a class of stochastic networks [2] known as G-networks [4, 7], which:

- Represent the flow of a discretised commodity (the EPs) and its storage in storage centres (STs) represented by queues,
- Representing the choices that are made regarding the distribution of these flows through the EPN via routing probabilities for the EPs; note that these routing probabilities would normally be determined by the SEDCs and can be modified so as to optimise the system's performance,
- The model can incorporate the flow of data as described in [41] so that control decisions can be taken, for instance to select the demands made by storage units and Cloud servers.

The data flows themselves will be coupled to the decisions to transfer EPs towards the consumers, whose energy consumption process is coupled to their computational service and is represented by a stochastic service process. Thus in this model the EPs constitute the "ordinary customers" of the queueing network, the STs are the queues, the external arrivals of EPs are the energy produced by different sources of energy, and the G-network's "triggers" are the data flows regarding requests made by consumers or by STs whose energy buffers are emptied by the Cloud servers, and are replenished by the renewable or other energy sources.

The system we consider has a set G of energy sources each with an energy generation rate $g(i, t)$ in EP/sec at time t , for $i \in E$, where $g(i, t) \leq G_M(i)$ which is its maximum generation rate. The energy sources are either renewable, in which case $i \in R$, or they are conventional in which case $i \in C$.

The system has a set of S of energy storage centres (ST) each with finite storage capacity $K(j)$, $j \in S$. Each storage centre has an energy conversion efficiency $0 < e_j \leq 1$ at its input so that on average the arrival of $B'(j)$ energy packets to ST j results in the storage of $B(j) = e_j \cdot B'(j)$ EPs.

Furthermore, in addition to its maximum energy storage capacity, it will also have a maximum rate $\theta(j)$ at which it can store energy. It also has an energy loss rate which is β_j per unit time so that if storage is not replenished, the $B(j)$ EPs that S(j) contains will be depleted on average after a time $B(j)/\beta_j$. In addition, a ST will have a finite maximum rate $D(j)$ at which it can deliver energy. Let $d(j, t)$ be the instantaneous delivery rate of the ST at time t and we will have $d(j, t) \leq D(j)$. In fact, when energy is converted back from stored energy to a flow that is dispatched to another storage or consumer centre, there will again be an internal conversion loss, but we will include it within the parameter $e(j)$.

We also have a set of C Cloud computing centres (CC) which are the main energy consumers. The c -th centre $C(c)$ has an a consumption rate of $m(c, t)$ in EPs/sec at time t . Some of these CCs may also have the ability to store energy locally. A Data Centre with its uninterruptible power supply (UPS) is an example of subsystem composed of a CC and a ST.

These centres are interconnected by an energy distribution or transport network (EDN) represented by a graph, so that link (u, v) of the directed graph represents a power line that has an energy transport capacity $C(u, v)$ which is the maximum amount of power that can be transferred instantaneously from node u to node v . In addition the link will have an efficiency $0 < c_{uv} \leq 1$ which is the fraction of energy introduced into the link that actually reaches the destination. The nodes of the EDN may be production nodes, consumption, storage nodes, or they may also be *transduction nodes* which can have many inputs and outputs. A transduction node u does not generate or store energy but dispatches it from one or more nodes to one or more other nodes; it has a transduction power capacity $T(u)$ so that for any successor node v we have $\sum_v C(u, v) \geq T(u)$ and for any predecessor v we have $\sum_v C(v, u) \leq T(u)$. Thus the incoming link capacities to a transduction node cannot exceed its own capacity, while the transduction node's outgoing links need to have a total capacity that exceeds its own capacity. The transduction node u will also have an efficiency $0 < t(u) \leq 1$ so that a fraction $t(u)$ of the power that it receives is wasted.

Each CC will send its energy requests to some Smart Dispatching Centre (SDC). SDCs are facilities that are interconnected to system components via a computer-communications network. Each SDC keeps track of the energy needs and requests in an area and assigns flows from the STs and EGs to the CCs. The SDCs also send requests to the EGs so that they may replenish the STs.

The SDCs' role is to satisfy the requests of the CCs and to make sure that STs have a standby capacity to meet unexpected needs. The SDCs will typically use pricing policies help the STs replenish their power at the best price, and also when energy from photovoltaic, wind or from other renewable sources is more readily available. The SDCs also attempt to maintain a flow of EPs across the EDN which is as low as possible so that energy traffic peaks are avoided and the EDN avoids saturation. Indeed one of the overarching objectives of the energy packet system is to be able to operate reliably with the lowest overall load being carried by the EDN.

Since the system as a whole depends on constant sensing, monitoring, communication and decision, the computer servers and network equipment will also constantly consume

energy and this needs to be included in the model. In particular we will assume that every unit in the system receives a constant flow of energy from some of the energy sources (e.g. generators) for this purpose. The assumption then is that each of the units in the system, whether it be a ST or a CC, is a connected to the generators to receive a flow of energy independently of its own requests based on its consumption, to assure that the information processing and communication systems can operate in a non-stop manner. Of course the generators themselves will be similarly monitored and connected to power sources (possibly the local units) but this aspect is neglected in the model. The purpose of these continuously running power supply from the generators to the ST and CC is also to allow them to receive some flow of energy over and above what they may request in order to allow their systems to operate all of the time on standby.

3. G-NETWORK MODEL OF THE EPN

Based on the previous presentation, a queueing network analysis of the system that we have described can be constructed based on a stochastic representation of energy production, storage and consumption. The only element that we will not include in the model are the transduction nodes, but these can be included in future studies. We represent by the probability π_{uv} , the fraction of energy leaving node u which is directed towards node v , while $p(u, v) = \pi(u, v)c(u, v)$ is the fraction of the energy that leaves node u and *actually arrives* at node v . Furthermore, we also represent the effect of the SDC by its effect on the manner in which energy is actually being requested and dispatched, so that $q(v, u)$ will be the probability that when node $v \in S \cup C$ consumes or dispatches energy to some other node, then it requests energy from some other node $u \in G \cup S$.

The time behaviour of the generators will be represented by Poisson flows of rate γ_g for $g \in G$, while the time dependent energy consumption of consumers will be represented by a rate parameter μ_k for $k \in C$. The leakage or loss rate of a storage unit will be represented by an exponential distribution of parameter β_j , while the storage units' instantaneous energy output cannot exceed some given rate δ_j for $j \in S$. Each of the generators g will also attempt to provide a flow ϕ_g of standby energy to the storage and consumer centres, as indicated above, although the net standby energy flow it supplies will be $\phi_g \rho_g$ if ρ_g is the probability that the generator actually is able to supply energy, and will obviously be strictly less than one in the case of renewable and intermittent sources. As a result of these definitions, we can write equations that represent the equilibrium behaviour of the whole system using G-network theory as follows, where the main approximation concerns the formula that we will use to represent the finite capacity of the storage units. Since the storage units have a maximum energy delivery rate γ_j we will also define the parameter D_j as the actual energy delivery rate of the ST j :

$$D_j = \max[\delta_j, \sum_{k \in C} \rho_k \mu_k q_{kj} \pi_{jk} + \sum_{i \in S} \rho_i \delta_i q_{ij} \pi_{ji}] \quad (1)$$

Applying G-network theory we have the following system of non-linear relations which describe the flow of energy into the storage centres and the consuming centres into which

energy arrives on demand:

$$\Lambda_j = \sum_{g \in G} \phi_g \rho_g s_{gj} c_{gj} + \rho_j D_j \left[\sum_{g \in G} \rho_g q_{ji} p_{ij} + \sum_{i \in S} \rho_i q_{ji} p_{ij} \right], \quad j \in S \quad (2)$$

$$\Lambda_k = \sum_{g \in G} \phi_g \rho_g s_{gk} c_{gk} + \rho_k \mu_k \left[\sum_{g \in G} \rho_g q_{kg} p_{gk} + \sum_{j \in S} \rho_j q_{kj} p_{jk} \right], \quad k \in C \quad (3)$$

where s_{gj} is the fraction of standby energy that a generator g supplies to the ST or CC j , while $c(g, j)$ is the fraction of this energy that actually gets to the destination. The unknown terms ρ_i , $i \in S \cup G$ represent the probability that the corresponding storage centres and generators have energy to provide, while ρ_k , $k \in C$ represents the probability that a consumer centre is busy consuming energy. These quantities are expressed as functions of the Λ_i , $i \in S \cup C$:

$$r_k = \frac{\Lambda_k}{\mu_k}, \quad k \in C, \quad (4)$$

$$r_j = \frac{\Lambda_j}{\beta_j + D_j}, \quad j \in S, \quad (5)$$

and for any unit that has finite capacity, such as a storage centre that can at most store K_j EPs, we have:

$$\rho_j = r_j \frac{1 - r_j^{K_j}}{1 - r_j^{K_j+1}} \quad (6)$$

Note that (6) is an approximation that is exact when we deal with a single queue with Poisson arrivals and exponential service times. We use it as an approximation in the networked systems that we are considering. Furthermore the case with $K_j = 1$ also will apply to a data centre that is limited to storing the energy that it is instantaneously consuming. For example during a machine cycle, all the energy that will be used in the cycle is already stored in the circuits of the computer which is drawing energy from its power supply as it computes successive cycles. When $K_j = 1$ we have

$$\rho_j = \frac{r_j}{r_j + 1} \quad (7)$$

and

$$r_j = \frac{\rho_j}{1 - \rho_j} \quad (8)$$

Since γ_g represents the nominal power generating rate of generator g at some given time, where g can be either a non-renewable or renewable source of energy, it will of course depend on the period of time being considered. For a non-renewable source will generally be possible to vary γ_g within a lower and upper bound, though changes may have to be carried out slowly over time. For a renewable source γ_g will generally depend on the instantaneous conditions (e.g. wind speeds, levels of ambient lighting for photovoltaic sources) but there are also ways to limit its value if it is deemed to be too high. We can therefore obtain the probability that the generator is able to satisfy all of its demands, given by:

$$\rho_g = \min \left[1, \frac{\gamma_g}{\phi_g + \sum_{k \in C} \rho_k \mu_k q_{kg} \pi_{gk} + \sum_{j \in S} \rho_j \delta_j q_{jg} \pi_{gj}} \right] \quad (9)$$

If the CCs also have unlimited local storage capacity, then from (4) we have:

$$\rho_k = \frac{\sum_{g \in G} \phi_g \mu_k \rho_g s_{gk} c_{gk}}{1 - \sum_{g \in G} \rho_g q_{kg} p_{gk} - \sum_{j \in S} \rho_j q_{kj} p_{jk}}, \quad k \in C \quad (10)$$

In the special case where the STs have infinite storage capacity $K_j = \infty$ so that $\rho_j = r_j$, and the STs' loss rate is negligible $\beta_j = 0$ we also have:

$$\rho_j = \frac{\sum_{g \in G} \frac{\phi_g}{D_j} \rho_g s_{gj} c_{gj}}{1 - \sum_{g \in G} \rho_g q_{ji} p_{ij} - \sum_{i \in S} \rho_i q_{ji} p_{ij}}, \quad j \in S \quad (11)$$

In order to have enough energy for all the needs of the system, these equations would have to satisfy the constraint:

$$\sum_{g \in G} \gamma_g \geq \sum_{g \in G} [\phi_g + \sum_{k \in C} \rho_k \mu_k q_{kg} p_{gk} + \sum_{j \in S} \rho_j \delta_j q_{jg} p_{gj}] \quad (12)$$

so that the total energy that reaches the storage centres and the consumers do meet the rate at which they make their demands. Better still, ideally we would also like all the consumers' energy needs to be satisfied:

$$\rho_k = 1, \quad k \in C \quad (13)$$

4. A SPECIAL CLASS OF EPNS

A special class of EPNS that has intuitive appeal would have the renewable energy sources feed an overwhelmingly large part of their energy production directly into storage units except for a small fraction that is sent to the data centres just to keep them awake in case the storage centres are down. This has the advantage of buffering the fluctuations of renewable energy sources, but will also lead to potentially higher storage and energy conversion losses. In this case after some analysis we can obtain results for two cases of interest.

Assume that the C data centres that are providing processing for the Cloud are all identical and have the same statistically identical computational load. Suppose that they all have unlimited local energy storage. Suppose also that the ensemble of S storage centres are identical with unlimited capacity. and let l be the fraction of power that is lost in transit through the energy transmission network. Assume also that the storage centres do not have any conversion losses. After some calculations we can show from the preceding analysis that the probability that any one Cloud service centre has enough energy to meet its computational load is given by:

$$\rho_c = \frac{\gamma_g G (1-l)}{C \mu_c} [(1-s) + s(1-l)] \quad (14)$$

where s is the fraction of energy that is sent directly to the storage centres.

On the other hand, under similar conditions, if the Cloud service centres do not have the ability to store energy locally so that there may be energy wastage from the energy they receive but do not absolutely need to run their computations (for instance, they may use the energy to cool down their equipment at a higher rate or to keep their machines running at a higher rate than is needed, then with the same level of energy production γ_g the probability that one of these data centres is unable to meet its computational load will be:

$$\rho_c^* = \frac{\frac{G}{C} \gamma_g (1-l) [(1-s) + sl]}{\mu_c + \frac{G}{C} \gamma_g (1-l) [(1-s) + sl]} \quad (15)$$

so that quite obviously $\rho_c > \rho_c^*$. This simply states that in addition to shared energy storage it would also be useful to have local energy storage at the data centres. However this analysis has not included the effect of energy loss at the storage units, nor of loss due to reconversion, although this can be included in the loss factor l . Thus the results will be mitigated when one considers these additional losses in the storage units.

5. CONCLUSIONS

This paper discusses a novel system that we call an *Energy Packet Network* to store renewable or cheap electric energy and deliver it on demand to computational systems in the Cloud. This approach can have great value whenever scarce or valuable sources of energy must be shared by multiple computational units whose peak to average power consumption ratio is high so that storage can smooth the intermittent supply and the intermittent demand. A method for the analysis of such systems based on queueing networks is suggested. The analysis is applied to a special case, indicating that it would be advantageous to have local energy storage at the data centres as well as energy storage units that are shared among many different data centres. In future work we plan to show how SEDCs can use distributed adaptive schemes such as those developed for smart packet network routing [8, 17] so that when a request for cloud computing arises, the SEDCs will select the computational resource which can do this at lowest energy and/or economic cost, and dynamically provision the energy dynamically.

Acknowledgements

The author gratefully acknowledges the support for this research from the Fit4Green European Union FP7 Project co-funded under ICT Theme: FP7-ICT-2009-4.

6. REFERENCES

- [1] D. W. Bunn, E. Farmer (Eds.) "Comparative Models for Electric Load Forecasting", John Wiley & Sons, 1985.
- [2] E. Gelenbe and A. Stafylopatis "Global behaviour of homogeneous random neural systems", *Applied Mathematical Modelling*, 15 (10):534-541, 1991.
- [3] R. Ramanathan, R. Engle, C.W. Granger, F. Vahid-Araghi and C. Brace "Short-run forecasts of electricity loads and peaks", *International Journal of Forecasting*, 13(2):161-174, 1997.
- [4] E. Gelenbe, A. Labed "G-networks with multiple classes of signals and positive customers", *European Journal of Operations Research*, 108(2):293-305, 1998.
- [5] G. Winkler, C. Meisenbach, M. Hable, and P. Meier, "Intelligent energy management of electrical power systems with distributed feeding on the basis of forecasts of demand and generation", In CIRED 2001.
- [6] D. J. Zack "Overview of wind energy generation forecasting", Tech. Rep., TrueWind Solutions, LLC, 2003.
- [7] E. Gelenbe and J.-M. Fourneau "G-Networks with resets", *Performance Evaluation*, 49:179-192, 2002.
- [8] E. Gelenbe "Cognitive Packet Network" *U.S. Patent No. 6804201 B1*, Oct. 12, 2004.

- [9] J. W. Taylor "Density forecasting for the efficient balancing of the generation and consumption of electricity", *International Journal of Forecasting*, 22(4):707–724, 2006.
- [10] E. Gelenbe and George Loukas "A self-aware approach to denial of service defence", *Computer Networks*, 51(5):1299–1314, April 2007.
- [11] J.R. Cancelo, A. Espasa, and R. Graffe "Forecasting the electricity load from one day to one week ahead for the spanish system operator", *International Journal of Forecasting* 24 (2):588–602.
- [12] M. Black and G. Strbac "Value of bulk energy storage for managing wind power fluctuations", *IEEE Transactions on Energy Conversion*, 22(1):197–205, Mar. 2007.
- [13] D. Infield, J. Short, C. Home, and L. Freris "Potential for domestic dynamic demand-side management in the UK", *IEEE Power Engineering Society General Meeting*, pp. 1-6, June 2007.
- [14] V. Dordonnat, S. Koopman, M. Ooms, A. Dessertaine and J. Collet "An hourly periodic state space model for modelling French national electricity load", *International Journal of Forecasting*, 24(4):566–587, 2008.
- [15] I. Sanchez "Adaptive combination of forecasts with application to wind energy", *International Journal of Forecasting* 24(4):679–693, 2008.
- [16] J. Manwell, J. McGowan, and A. Rogers "Wind Energy Explained: Theory, Design and Application", Wiley, 2009.
- [17] E. Gelenbe "Steps toward self-aware networks", *Comm. ACM*, 52(7):66–75, 2009.
- [18] A. Berl, E. Gelenbe, M. di Girolamo, G. Giuliani, H. de Meer, M.-Q. Dang, and K. Pentikousis "Energy-efficient Cloud Computing", *The Computer Journal*, Vol. 53 (7), pp. 1045–1051, 2010.
- [19] G. Sakellari and E. Gelenbe, "Demonstrating cognitive packet network resilience to worm attacks", *Proc. ACM Conference on Computer and Communications Security*, pp. 636–638, 2010.
- [20] The MeRegio Project. <http://www.meregio.de/en/>. 2011. Center for Renewable Energy Sources. <http://www.cres.gr/>.
- [21] H. Berthold, M. Boehm, L. Dannecker, F.-J. Rumph, T.B. Pedersen, C. Nyctis, H. Frey, Z. Marinsek, B. Filipic and S. Tselepis "Exploiting renewables by request-based balancing of energy demand and supply", *Proc. 11th IAEE European Conference*, 2010.
- [22] MIRACLE Project 2010. MIRACLE Project Website. MIRACLE Project. <http://www.miracle-project.eu>.
- [23] Nationalgrid UK 2010. Metered half-hourly electricity demands. Nationalgrid UK. <http://www.nationalgrid.com/uk/Electricity/Data/Demand+Data/>.
- [24] NREL 2010. Wind Integration Datasets. NREL. <http://www.nrel.gov/wind/integrationdatasets/>.
- [25] E. Gelenbe and C. Morfopoulou "Routing and G-Networks to optimise energy and Quality of Service in packet networks", *Energy-Efficient Computing and Networking, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, 54, Part 4, 163–173, Springer, 2011.
- [26] E. Gelenbe and T. Mahmoodi, "Energy-Aware Routing Protocol in the Cognitive Packet Network", *International Conference on Smart Grids, Green Communications, and IT Energy-aware Technologies (Energy 2011)*, Venice, Italy, 22–27 May 2011, ISBN : 978 – –1 – –61208 – –006 – –2.
- [27] E. Gelenbe and C. Morfopoulou "A framework for energy aware routing in packet networks", *The Computer Journal* 54 (6): 850–859, 2011.
- [28] Dinorwig power station. First Hydro Company, URL: <http://www.fhc.co.uk/dinorwig.htm>.
- [29] E. Bitar, R. Rajagopal, P. Khargonekar, and K. Poolla "The role of co-located storage for wind power producers in conventional electricity markets", *Proc. American Control Conference (ACC)*, pp. 3886–3891, July 2011.
- [30] K. Chandy, S. Low, U. Topcu, and H. Xu. A simple optimal power flow model with energy storage. In 49th IEEE Conference on Decision and Control (CDC), pages 1051–1057, Dec. 2010.
- [31] D. Gayme and U. Topcu "Optimal power flow with distributed energy storage dynamics", *Proc. American Control Conference*, 2011.
- [32] P. Grünwald, T. Cockerill, M. Contestabile, and P. Pearson "The role of large scale storage in a GB low carbon energy future: Issues and policy challenges", *Energy Policy*, 39(9): 4807–4815, 2011.
- [33] E. Gelenbe "Energy Packet Networks", Keynote Talk, *GreenNets 2011 : The First ICST International Conference on Green Communications and Networking*, Colmar, France, Oct. 5–7, 2011.
- [34] C. Naish, I. McCubbin, O. Edberg, and M. Harfoot "Outlook of energy storage technologies", Technical Report
- [35] H. Oh "Optimal planning to Include Storage Devices in Power Systems", *IEEE Transactions on Power Systems*. 26 (3): 1118–1128, Aug. 2011.
- [36] G. Sinden "Characteristics of the UK wind resource: long-term patterns and relationship to electricity demand", *Energy Policy*, 35 (1):112–127, 2007.
- [37] H.-I. Su and A. E. Gamal "Modeling and analysis of the role of fast-response energy storage in the smart grid", *Proceedings of the Forty-Ninth Annual Allerton Conference on Communication, Control, and Computing*. University of Illinois at Urbana-Champaign, Sept. 2011.
- [38] H.-I. Su and A. E. Gamal "Modeling and analysis of the role of fast-response energy storage in the smart grid", *CoRR*, *abs/1109.3841*, 2011.
- [39] Financial Times, 29/8/2011, p.2
- [40] N. Wade, P. Taylor, P. Lang, and P. Jones "Evaluating the benefits of an electrical energy storage system in a future smart grid", *Energy Policy*, 38 (11): 7180–7188, 2010.
- [41] E. Gelenbe "Energy Packet Networks: Smart energy storage to meet surges in demand", *Proc. 5th International ICST Conference on Simulation Tools and Techniques*, Simutools 2012, March 19–23, 2012, Desenzano, Italy.