

On the Impact of Storage in Residential Power Distribution Systems

Omid Ardakanian
University of Waterloo

Catherine Rosenberg
University of Waterloo

S. Keshav
University of Waterloo

ABSTRACT

It is anticipated that energy storage will be incorporated into the distribution network component of the future smart grid to allow desirable features such as distributed generation integration and reduction in the peak demand. There is, therefore, an urgent need to understand the impact of storage on distribution system planning. In this paper, we focus on the effect of storage on the loading of neighbourhood pole-top transformers. We apply a probabilistic sizing technique originally developed for sizing buffers and communication links in telecommunications networks to jointly size storage and transformers in the distribution network. This allows us to compute the potential gains from transformer upgrade deferral due to the addition of storage. We validate our results through numerical simulation using measurements of home load in a testbed of 20 homes and demonstrate that our guidelines allow local distribution companies to defer transformer upgrades without reducing reliability.

1. INTRODUCTION

It is widely believed that the future electrical grid will have significant amount of storage. Storage may be added at one of several locations: it may be installed near generators to even out variations in generation, in the transmission network to even out peak transmission loads, or at substations and feeders in distribution networks to absorb variations in electrical demand (also referred to as *load*) [2]. In this paper, we study effect of storage on sizing a neighbourhood pole-top transformer.

Customer electrical demand typically exhibits diurnal and seasonal variations. In this situation, if a pole-top transformer is sized to meet only the long-term average demand, there will inevitably be periods when it will be overloaded. An overloaded transformer may overheat, have a shorter life expectancy, and is more likely to fail [1]. Therefore, utilities size transformers for the peak customer demand; this increases the size and the cost of a transformer. However, it is not economical to size a transformer for a peak that may occur only once every ten years.

The dilemma of choosing between transformer overload and underutilization can be resolved by introducing storage adjacent to the transformer. A store that fills up when customer demand is low and partially fulfils demand during periods of heavy demand ensures that, even if the transformer is sized for a demand smaller than the peak, it never exceeds its nameplate rating¹. Of course, this requires guide-

lines to determine the amount of storage that corresponds to a given average and peak customer demand. This is one primary focus of our work.

Storage can also help with transformer upgrade deferral. As customer demands grow over time, both the average and peak demands increase. To meet network reliability standards, local distribution companies (LDCs) must upgrade their infrastructure to meet these increasing demands. This usually requires upgrading pole-top transformers to higher nameplate ratings. However, upgrading a transformer can be costly. Instead, LDCs could leave transformers unchanged and partly serve load using storage. This requires guidelines to determine the amount of storage that would offset a given increase in customer demand, and is a secondary focus of our work.

Distribution transformers are typically sized by electric utilities using an approach similar to that described in Reference [14]. However, utilities lack guidelines for joint sizing of distribution transformers and storage. The novelty of our work is, therefore, to apply probabilistic sizing techniques developed for sizing buffers and links in communication networks to jointly size storage and transformers in the distribution network based on the analogy between a communication access network and a power distribution network (shown in [5]).

We make three specific contributions:

- We present a theoretical foundation for *jointly* sizing pole-top transformers and storage based on methods developed for sizing buffers in telecommunication systems.
- We demonstrate how to map electricity demand to the standard dual leaky-bucket parameters.
- We present a joint sizing guideline for transformers and storage in a residential neighbourhood and show the impact of the aggregate load parameters on this guideline.

Given the dropping costs of storage and its substantial benefit, we believe that it is important to develop guidelines to jointly size pole-top transformers and storage. We believe that this relatively simple framework will be useful to electric utilities to compare different options especially since it is based on a conservative set of assumptions that are in line with today's ways of dimensioning residential transformers.

power output that a transformer can continuously deliver at rated voltage and frequency without exceeding a specified temperature. At this temperature a transformer has its normal lifetime.

¹The nameplate rating of a transformer is the maximum

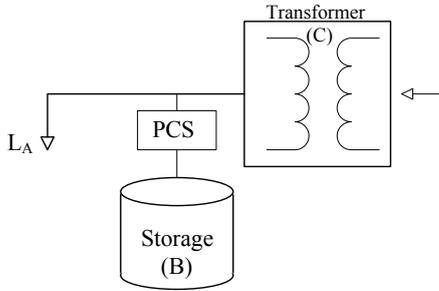


Figure 1: A branch of the electrical grid with the aggregate residential load, L_A , where the capacity of storage is B Watt-hours and the nameplate rating of the transformer is S Volt Amperes. This system is modelled as a fluid queue.

2. BACKGROUND AND ASSUMPTIONS

2.1 Assumptions

A branch of a residential distribution network is shown in Figure 1. It consists of a pole-top transformer with a nameplate rating of S kVA supplying residential demands of n homes, the sum of whose demands creates an active aggregate time-varying load of $L_A(t)$ kW. We assume that all homes are located in a small geographical area so that distribution losses can be neglected. We also assume that generation can always meet the aggregate demand so that it is not a bottlenecked resource.

Let C denote the active power that can be supplied by the transformer. Note that $C = Sf$ where f is the power factor computed for the aggregate load at the transformer. For simplicity, we assume that the power factor is fixed and known *a priori*.

The system contains storage that is characterized by two parameters: its capacity, B Wh, and its charge/discharge rating, *i.e.*, the maximum power at which it can be charged or discharged. Storage is connected in shunt to the distribution feeder through a power control system, denoted by PCS. We assume that storage can be fully charged or discharged and that the charge/discharge process is 100% efficient.

A typical PCS consists of an inverter, a transformer, and a charge controller which controls the charging and discharging of the storage device. Specifically, it charges storage at the rate $C - L_A(t)$ until the storage becomes full whenever the aggregate customer demand is less than the power supplied by a transformer that is loaded at its nameplate rating². Symmetrically, it discharges storage at rate $L_A(t) - C$ until the storage becomes empty whenever the aggregate demand is greater than the power supplied by the transformer that is loaded at its nameplate rating. We assume that the charge/discharge rating of storage is greater than both $C - \min\{L_A(t)\}$ and $\max\{L_A(t)\} - C$.

Whenever storage is depleted and the aggregate customer demand is greater than the power supplied by the transformer loaded at its nameplate rating, the transformer is overloaded, which impacts its lifetime and may even result in an outage. Therefore, a loss of load *might* happen when-

²Since our goal is to maximize the utilization of distribution assets, we assume that PCS charges the store at a rate that requires loading the transformer at its nameplate capacity.

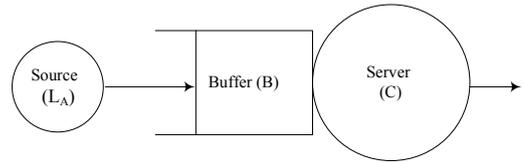


Figure 2: A fluid queue serves traffic generated at the rate of L_A bits/second. The capacity of the buffer is B bits and the service rate of the server is C bits/second.

ever storage is depleted and the aggregate demand exceeds C . We assume, conservatively, that a loss of load *will* happen whenever these two conditions are met.

Grid reliability is typically expressed in terms of the *loss-of-load probability (LOLP)* [13]. The “one-day-in-ten-years” reliability criterion ($LOLP = 2.74 \times 10^{-4}$) is a widely used benchmark. To comply with this reliability criterion, the transformer and storage should be sized such that the loss of load event (*i.e.*, the storage underflow event) occurs less than one-day-in-ten-years. Thus, our plan of attack is to compute the storage underflow probability or to find an upper bound on it for a given transformer and storage sizing and make sure that it is less than LOLP. This upper bound depends on how we model the aggregate load as discussed in Section 4.

2.2 The Equivalence Theorem

We apply probabilistic sizing methods developed for dimensioning buffers and links in the context of Internet to jointly size transformers and storage. This is possible because of the Equivalence Theorem (proved in [5]) which demonstrates that a residential distribution network can be accurately modelled as a simple telecommunication network called a *fluid queue*. In a fluid queue, telecommunication sources generate infinitesimal packets that are served by a telecommunication link associated with a buffer. The equivalence theorem indicates that this model can also be used to study a distribution network where electrical sinks consume infinitesimal units of energy (*i.e.* electrons) supplied by a distribution network with a given peak capacity C and associated storage.

The Equivalence Theorem motivates us to model the distribution network described by a transformer with nameplate rating, $C = Sf$ kW, a store of capacity B Wh, and an aggregate time-varying demand $L_A(t)$ kW as a fluid queue with a link of capacity C bits/second, a buffer of size B bits, and an aggregate time-varying source rate of $L_A(t)$ bits/second. Given this model, we are interested in computing the set of B, C pairs corresponding to a desired buffer overflow probability. The direct consequence of the Equivalence Theorem is that the loss of load probability in the distribution network can be approximated by the overflow probability in the dual fluid queue (Figure 2). This is important because the problem of upper bounding the loss probability in a fluid queue (in contrast to the upper bounding an underflow probability) is well-understood [7, 8, 10, 15].

To sum up, in the remainder of the paper we study a simple fluid queueing system. An (infinitesimal) arrival in this system brings $L_A(t)$ traffic to the system at time t and the service rate is C . Our goal is to find an upper bound on the buffer overflow probability of this system. A pair (B, C) is said *feasible* if this upper bound is less than 2.74×10^{-4} .

3. FLUID ANALYSIS

We first present a brief tutorial on the fluid queueing model and then overview the existing upper bound on the stationary probability that the backlog grows beyond a certain level.

3.1 A Fluid Queueing Model

A fluid queueing system is a type of queueing system where an arrival event could occur at any $t \in \mathbb{R}^+$, and the amount of work brought to the queue by an arrival is continuous.

Definition 1. The cumulative input to a queueing system in any interval I , denoted by the function $A(I)$, is defined as the total traffic (also called *work*) that has arrived to the system in I .

Definition 2. The cumulative output from a queueing system in any interval I , denoted by the function $D(I)$, is defined as the total traffic that has departed from the system in I .

Observe that both A and D are continuous, monotonically increasing functions, and $A(I) \geq D(I)$ in any interval I if the buffer is empty at the beginning of this interval. We assume that both A and D are stationary stochastic processes³. The stationarity assumption permits us to extend the domains of A and D as $A(s, t]$ and $D(s, t]$ in $-\infty < s \leq t$.

From the above definitions, it is clear that the backlog of a system at time t (i.e., the amount of work that is in the system at t and is denoted by $Q(t)$) is equal to $A(s, t] - D(s, t]$ if no work was in the system at time s ($s < t$). We say that an interval I is a backlogged interval iff for all $t \in I$ we have $Q(t) > 0$. The server is never idle in a backlogged interval. The following corollary is well-known:

As the difference of two stationary processes, Q is also a stationary process. We can therefore extend the domain of Q to the whole real line. Following convention, the stationary backlog process is denoted by $Q(0)$. Therefore, our goal reduces to finding $\mathbb{P}(Q(0) = B)$ in a finite capacity queueing system. This corresponds to the buffer overflow probability as $Q(t)$ cannot grow beyond the buffer size B .

Computing the loss probability in a finite buffer system is quite difficult. A common practice is to approximate the loss probability in a system with a buffer of size B with the probability of up-crossing the level B in a system with an infinite buffer [17] (as long as the infinite buffer queue is stable, that is, the long-term arrival rate does not exceed the link capacity C). We use this approach in this paper since this condition holds. Our revised goal, therefore, is to find $\mathbb{P}(Q(0) \geq B)$ in an infinite capacity queueing system.

3.2 An upper bound on $\mathbb{P}(Q(0) \geq B)$

The upper bound on the probability that the backlog in a stationary infinite buffer grows beyond level B was derived by Kesidis and Konstantopoulos [10] for a work-conserving queueing system that has a constant service rate, C , and a deterministically shaped arrival process. We present this bound next.

Definition 3. Suppose $A(I)$ is a cumulative input function as defined in Section 3.1. We say that A is α -smooth or equivalently α is an *arrival curve* if, for any finite interval,

³In fact A is not stationary due to time of day dependency. However, our analysis applies eventually to the peak day of year or busy hours modelled as a stationary process.

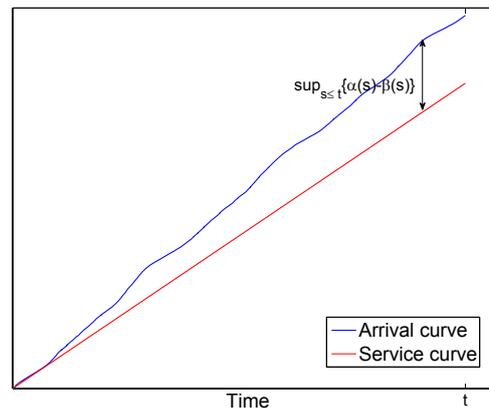


Figure 3: Arrival and service curves.

I , we have $A(I) \leq \alpha(|I|)$ where $|I|$ is the length of interval I .

In the following, we will consider a (π, ρ, σ) arrival curve defined as $g(t) := \min\{\sigma + \rho t, \pi t\}$ with $\rho < \pi$. This means that we upper bound the arrival process by a deterministically shaped process characterized by the three parameters (π, ρ, σ) . Note that this is a standard dual leaky-bucket model [9].

Definition 4. Suppose $D(I)$ is a cumulative output function as defined in Section 3.1. We say that the system offers a *service curve* β to a flow if for any finite backlogged interval I the output of the queue is at least $\beta(|I|)$, i.e., $D(I) \geq \beta(|I|)$.

Thus, an upper bound on the backlog is given by [11]:

$$Q(t) \leq \sup_{s \leq t} \{\alpha(s) - \beta(s)\} \quad (1)$$

As can be seen in Figure 3, this is the maximum vertical distance between α and β .

Now, consider the case where the cumulative input function to an infinite queueing system is constrained by $g(t)$ and the server offers a constant service rate C where $\rho < C < \pi$. Kesidis and Konstantopoulos [10] derived an upper bound for the stationary probability that the buffer level exceeds B under the assumption $\frac{(\pi-C)\sigma}{(\pi-\rho)} \geq B$:

$$\mathbb{P}(Q(0) \geq B) \leq \frac{\sigma - \frac{\pi-\rho}{\pi-C} B}{\frac{C}{\rho} \sigma - B} \quad (2)$$

This bound allows the sizing and provisioning of links and buffers so that the quality of service requirement is guaranteed, i.e., the loss probability is less than a threshold $\epsilon = \text{LOLP}$. This is because the minimum link bandwidth C (measured in bits per second) is a function of the buffer size B (measured in bits) for a given quality of service requirement ϵ .

4. METHODOLOGY FOR JOINT SIZING OF STORAGE AND TRANSFORMER

We now show how utilities can use the bound in (2) to compute transformer sizing and storage capacity while meeting network reliability requirements. Recall that we approximate the storage underflow probability in the grid (which we interpret as the probability of a failure in the grid) by the stationary buffer overflow probability in the dual queueing system. Thus, the bound in (2) is simply the LOLP, allowing

us to rewrite (2) as follows:

$$B \geq \frac{\sigma(1 - LOLP \frac{Sf}{\rho})}{\frac{\pi - \rho}{\pi - Sf} - LOLP} \quad (3)$$

This allows us to jointly size the transformer nameplate rating $S = C/f$ and the storage size B (for $B \neq 0$).

Note that the joint sizing depends on choice of π , ρ , and σ , which are chosen by the utility to properly model the aggregate customer demand $L_A(t)$ and a significant contribution of our work is in the careful choice of these parameters. Given a neighborhood characterized by its aggregate load demand $L_A(t)$, a utility has to estimate the three parameters π , σ , and ρ so that $L_A(t) \leq \min\{\sigma + \rho t, \pi t\}$ for all t . In a telecommunication network, π is the peak traffic rate, ρ is a bound on the long term average traffic rate, and σ is the maximum burst tolerance which is defined as the buffer size required for a loss-free transmission when the link rate is ρ . This gives us the following insight into choosing these parameters in the context of the distribution grid.

Let $S(B)$ be the minimum value of S for which inequality (3) is satisfied for the target LOLP. Note that utilities already have guidelines to size residential transformers when there is no storage, i.e., they have guidelines to compute $S(0)$. Therefore, we suggest that a reasonable value for π is $\pi = S(0)f$ because $S(0)$ has been selected to ensure that the target LOLP is met in the peak day of year⁴ when there is no storage.

We define ρ as the long-term mean demand; i.e., the mean demand over a time duration which is long enough to capture diurnal and seasonal variations in aggregate customer demand. Utilities that already instrument their distribution network to obtain long-term measurements can use these measurements to directly compute ρ .

In cases where such long-term measurements are not available, a practical alternative would be to estimate ρ as the average customer demand during either the peak day, the peak week, or the peak month of year because each of these estimators is a strict upper bound on the long-term mean demand.

The parameter σ is defined as the storage capacity required so that a transformer with the nameplate rating of ρ/f (i.e. the long-term average demand) is not overloaded. We propose two different approaches for computing σ . The first approach is based on a numerical simulation of the charge level of storage for different storage capacities given the measured aggregate load profile on the peak day of year to determine whether the transformer is loaded higher than its nameplate rating, ρ/f . We set σ to be the minimum storage capacity that results in no overloading. We use the load profile of the peak day because if the storage capacity is enough to prevent the transformer overloading in the peak day it would be enough to prevent the transformer overloading in any other day.

The second approach upper bounds σ by $(\pi - \rho)T$ where T is defined as the *sum* of the length of time periods in the peak day of year in which the aggregate load is greater than ρ . This approach is more conservative than the first approach and results in a larger storage size.

The choices of ρ and σ are not independent. A utility that

⁴The peak day is either in summer or in winter depending on the geographical region. We assume that utilities have accurate estimates for the load profile of the peak day.

chooses ρ conservatively can use a less conservative choice of σ because the closer the estimated long-term demand is to the estimated peak-demand, the less the variability in the model of the aggregate customer demand.

The approach described above can be used for transformer upgrade deferral as follows. Suppose that the LDC can estimate the future value of π , ρ , and σ by analyzing historical trends in these values. This allows the LDC to compute the B required meet the LOLP when S is fixed. This is precisely the amount of storage required to accommodate increases in customer demand without changing the transformer, as required.

We end this section with two technical notes. First, the careful reader will note that we have modelled a power transformer as a server with a *constant* service rate, $C = Sf$. However, a power transformer might be loaded higher than its nameplate rating for a limited time, for example, during a demand peak period. Therefore, it is more accurate to model the power transformer as a server which offers a service curve $\beta(t) = Ct$ (i.e., the service rate can be greater than C). We prove in Appendix A of the extended version of this paper (Reference [6]) that even with this model the stationary overflow probability of the fluid queue with the same deterministically shaped arrival curve is bounded by the same expression derived for a system with a constant service rate. Second, the bound in (3) does not hold when $B = 0$. In Appendix B of [6], we show that B goes to 0 when $S(B)$ goes to π/f , so that (3) is indeed continuous at this limit.

5. RESULTS

We now validate the joint sizing guidelines proposed in Section 4. In practice, LDCs can estimate the three parameters (π, ρ, σ) using their existing measurements and by following the prescribed methodology. However, we do not have access to LDC meter readings and field data. To mitigate this to some extent, we deployed our own measurement nodes in a residential neighborhood consisting of 19 houses and one home-based small business as described in [4]. Each measurement node consists of a Current Cost Envi device [3] and a netbook. The Envi device measures the active power consumption (in Watts) of a house every six seconds and stores it locally in flash memory. A script on the netbook queries the device every six seconds to obtain an XML file that it stores on disk. This is uploaded using a secure connection to a server in our laboratory once a day.

A guideline provided by a utility in our region classifies homes into a number of different classes based on a few simple parameters including the living area size and the nature of the heating and cooling systems, which constitute the major loads in our geographical area. These parameters are used to compute a ‘unit value’ that represents the load expected from that home. The guideline maps the total unit value of a neighbourhood to a transformer size when there is no storage, i.e., it gives $S(0)$.

We asked the participants of our study to tell us their home’s unit value. We used this to compute the total unit value of this neighborhood. The transformer size that is recommended by the guideline for this neighborhood is thus computed as 100 kVA. Therefore, assuming that the power factor corresponding to the aggregate residential load is equal to $f=0.95$, we set the value of π to 95 kW.

Furthermore, based on our measurements over a period

of six months including the peak month, peak week, and peak day of year, we compute the three estimators of the sustained mean load; these are 25.5 kW, 30.6 kW, and 34.9 kW respectively. Then, using the two approaches for computing σ introduced in the previous section, we obtain six different (ρ, σ) pairs. These allow us to compute the set of B and corresponding S values that meet the LOLP; we call this a *trade-off curve*. Each point on each trade-off curve corresponds to a (S, B) pair computed by substituting π , ρ , and σ in (3).

Figure 4 illustrates the impact of the aggregate load parameters on the trade-off curves. As expected, the second approach to compute an upper bound on σ results in more conservative trade-off curves (that is, higher values of B for a given value of C). Moreover, it turns out that the farther the estimated ρ is from the actual sustained mean load the lower the trade-off curve.

Recall that our sizing guidelines assume that the aggregate customer demand is stationary. In practice, the aggregate demand has distinct diurnal and seasonal variations and therefore is far from stationary. We therefore perform numerical simulation to validate the degree to which the (S, B) pairs obtained from a trade-off curve are admissible, that is, they satisfy the reliability criterion of the grid ($LOLP = 2.74 \times 10^{-4}$). This numerical simulation uses the sum of the loads of the 20 homes measured over the period of six months to compute the transformer overloading duration for every (S, B) pair.

Our simulations show that the transformer overloading duration is zero for all (S, B) pairs except in two cases. The first case corresponds to (S, B) pairs when S is less than 32 kVA, ρ is estimated by the average demand of the peak month, and σ is computed using the first approach. The second case corresponds to (S, B) pairs when S is less than 35 kVA, ρ is estimated by the average demand of the peak week, and σ is also computed using the first approach. In both cases, the grid reliability criterion *may* be violated if overloading of the transformer leads to the loss of load. This suggests that we should compute σ using the second approach especially when the transformer is sized for a value that is very close to the long-term mean load. When using this latter approach, despite the assumptions made in our work, the sizing guidelines result in *no* loss of reliability.

6. RELATED WORK

The lines of work closest to ours are by Le Boudec et al. [12] and Wang et al. [16]. In [12] min-plus system theory is used to size the battery, and schedule its operation such that it can be guaranteed that the inflexible load is always satisfied. In [16] network calculus concepts are used to jointly size storage, solar photovoltaic panels, and wind turbines for a specific loss of power supply probability when intermittent renewable generation is supplemented with energy storage. However, both of these papers are different from ours as they do not consider transformer capacity limitations.

7. CONCLUSION

The introduction of storage into the distribution grid creates the challenge of jointly sizing pole-top transformers and their associated storage. We propose a novel approach to this problem drawing on an analogy to sizing telecommunication networks. We present trade-off curves obtained for a neigh-

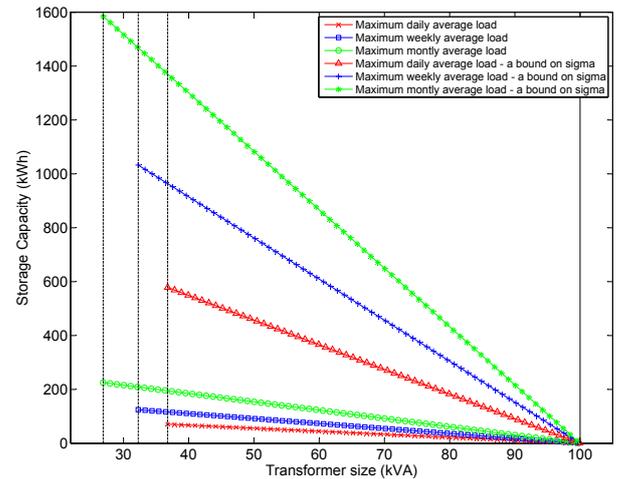


Figure 4: Trade-off curves. The lower three curves are obtained using the first approach for choosing σ and the rest are obtained using the second approach.

borhood in which we deployed measurement nodes which measure the active power consumption and showed, using numerical simulations, that, despite the many assumptions of our work, with the appropriate choice of load parameters, our guidelines do not result in the loss of network reliability.

8. REFERENCES

- [1] IEEE Guide for Loading Mineral-Oil-Immersed Transformers. *IEEE Std C57.91-1995*, 1996.
- [2] Handbook of Energy Storage for Transmission and Distribution Applications. *EPRI-DOE*, 2003.
- [3] Current cost. <http://www.currentcost.com/>, January 2011.
- [4] O. Ardakanian, S. Keshav, and C. Rosenberg. Markovian models for home electricity consumption. In *Proc. ACM SIGCOMM Green Networking Workshop*, 2011.
- [5] O. Ardakanian, S. Keshav, and C. Rosenberg. On the use of teletraffic theory in power distribution systems. In *Proc. e-Energy*, May 2012.
- [6] O. Ardakanian, C. Rosenberg, and S. Keshav. On the impact of storage in residential power distribution systems. Technical Report CS-2012-08, University of Waterloo, May 2012.
- [7] C.-S. Chang, Y.-m. Chiu, and W. T. Song. On the performance of multiplexing independent regulated inputs. In *Proceedings of the 2001 ACM SIGMETRICS*, pages 184–193, 2001.
- [8] R. Cruz. A calculus for network delay. I. *Information Theory, IEEE Transactions on*, 37(1):114–131, Jan. 1991.
- [9] S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley, 1997.
- [10] G. Kesidis and T. Konstantopoulos. Extremal shape-controlled traffic patterns in high-speed networks. *Communications, IEEE Transactions on*, 48(5):813–819, may 2000.
- [11] J.-Y. Le Boudec and P. Thiran. *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*. Lecture Notes in Computer Science. Springer, 2001.
- [12] J.-Y. Le Boudec and D.-C. Tomozei. A demand-response calculus with perfect batteries. In *WoNeCa*, 2012.
- [13] A. Meier. *Electric Power Systems: A Conceptual Introduction*. Wiley-IEEE Press, 2006.
- [14] K. Schneider and R. Hoad. Initial transformer sizing for single-phase residential load. *Power Delivery, IEEE Transactions on*, 7(4):2074–2081, oct 1992.
- [15] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. *Communications, IEEE Transactions on*, 51(5):735–740, may 2003.
- [16] K. Wang, S. Low, and C. Lin. How stochastic network calculus concepts help green the power grid. In *IEEE Smart Grid Communications 2011*, pages 55–60, oct. 2011.
- [17] Y. Ying, F. Guillemin, R. Mazumdar, and C. Rosenberg. Buffer overflow asymptotics for multiplexed regulated traffic. *Perform. Eval.*, 65:555–572, July 2008.