Testbed Evaluations of a Controlled-Delivery Power Grid

Roberto Rojas-Cessa, Vinit Sahasrabudhe, Eugene Miglio, Divya Balineni, Justin Kurylo, and Haim Grebel
Department of Electrical and Computer Engineering
New Jersey Institute of Technology
Newark, NJ 07102
rojas@njit.edu

***Abstract*—We present two testbeds and experimental evaluations of the controlled-delivery power (CDP) grid. In the CDP grid, energy is addressable and the amount of power is allocated to users in real time. The energy supplier grants the amount of energy supplied to a user and at the same time, ensures that the total supplied energy is within a fixed cap. Energy amounts are discretized to minimize management complexity and optimized for digital control. One of our testbeds mimics the present grid and the other emulates a grid provisioned with energy storage. In the presented experiments, users request and receive grants through power access points implemented with low-cost computers. The management and control plane of a supplier was implemented with a workstation, and users and supplier communicated through a data network. Alternating current (AC) at 120 V power lines were controlled by the server and delivered power to the access points. The supplier was able to maintain a high satisfaction ratio in terms of requests granted and energy delivered under limited capacity conditions.**

I. Introduction

The concept of a controlled-delivery power (CDP) grid has been proposed [1]–[3] as a potential approach to achieve a finer and more efficient balance between generation of electrical power and the demand of it. In this approach, energy is delivered in discrete (digital) levels to addressed users. By contrast, the present power grid distributes energy in an analog fashion where users are allowed to consume discretionary amounts of power at arbitrary times. This uncontrolled accessibility requires generators to adapt the generation of power to rebalance the grid at times of overproduction and to provide generation safety margins to avoid underproduction.

Everyday experience shows us the robustness of the present power grid; yet, the response of electricity providers (and generators) to increasing demand may still lag due to demand and production fluctuations beyond forecast estimations [4]. A smarter stable grid, where demand is most closely followed by the supply, is highly desirable. The present grids allow users to access and consume large amounts of electricity at any time. Grid overloads must be carefully monitored and in case of an extraordinary event, distribution loops ought to be taken out of the grid. This broad-scale monitoring cannot pinpoint specific users. Close monitoring of the grid's performance may be achieved by deploying (auxiliary) sensing data networks [5]–[14]. At the same time, ensuring working paths, which carry energy all the time, translates into additional management complexity [12], [15]. These works show the adoption of a controlled distribution of power that can be seamlessly coupled with grid monitoring as an objective.

The concept of controlling the distribution of energy through micro-grids has been recently discussed [16]. Approaches to verify users identification before the start of energy transmission in point-to-point communications have been also considered [12]. These approaches, however, require one-to-one connection and therefore, are unscalable to large number of users. Some of these works were motivated by the adoption of alternative-energy applications, where sources and appliances are connected through separate direct current lines, [17]. However, uncontrolled delivery (and consumption) may remain. Elastic loads have been proposed to balance the grid; but through load scheduling by the provider [18].

The CDP grid offers an alternative while performing precise control on energy delivery. In the CDP grid model, users issue requests for power and the provider may fully, or partially, grant these requests. This process enables the provider to determine how to satisfy the requests in accordance with the total demand and energy availability. The CDP grid model also implies the adoption of a controlled supply and a limited capacity. The management of supply under a capped capacity can be made on a user-by-user and a fair basis. It allows the service provider to supplying partial amounts of the energy needed and negotiating with the user for portions of the requested energy.

The amount of delivered power is controlled through the designation of ownership to specific power amounts. Here, ownership means that the supplied power may only be delivered to addressed user(s) [1]. Power is delivered upon matching the user address at an access point at the user's premises to that of packets carrying grants. The address (which may be similar to a network address) is carried by the electrical signal or by an auxiliary data network.

In this power grid, smart power switches ensure the level of power delivered during any given granted period. Demand fluctuations, if present, are known in real time and the energy provider has the option to meet these sporadic demands in full, or negotiate a partial supply. This property is especially noticeable under grids with limited power or under the presence of failure, where energy can be routed to eligible loads rather than being dissipated.

In this paper, we show a testbed of a controlled-delivery power grid where energy is addressable and controlled by the energy distributor. We show two testbeds, one mimicking the real-time supply of the present grid and another with energy storage. The functioning of the CDP grid is demonstrated and the performance is analyzed. In consistency with previous simulation results [3], we show that the satisfaction of energy request approaches 100% as the loop capacity increases through experiments. We show that this power grid does not require maximizing energy generation to satisfy high demands of electricity consumption.

The remainder of this paper is organized as follows. Section II briefly introduces the controlled-delivery power grid. Section III discusses the request model used to emulate user energy expenditure and request. Section IV describes our testbeds of controlled-delivery grid and our experimental evaluations on scenarios with energy storage and without it. Section V presents our conclusions.

II. Controlled-delivery Power Grid

The power grid with controlled delivery aims at supplying discrete and finite amounts of energy, on demand, to users. The approach minimizes the difference between energy generation and supply, facilitates power distribution amongst several grids, and provides intrinsic grid monitoring.

The controlled-delivery power grid carries energy in discrete units. This amount of energy may be controlled by regulating the supply time, current, or a combination of both as provided to a user. To overcome having a grid exposed to discretionary loads, the electrical signal carries the destination address(es) of specific user(s) who are the only one(s) allowed to get access to the transmitted energy. The destination address may be embedded into the electrical signal by encoding the addresses of the destination user (or user premises, as a residential, commercial, or industrial location). Because energy may need to be simultaneously distributed to a single or a number of users, addressing through code division multiple access (CDMA) [19] can be used, where users have their own address [1]. In this paper, however, we use broadcast packets in experimental testbeds as a simplified approach. In this power grid model, the amount of power (e.g., current) is set to discrete levels. The selected level of power destined for a user may be set by using selective current limiters, called smart loads, at the user premises, as Figure 1 shows. The energy supplier performs the selection of a smart load by also embedding the amount of current granted per user in the electrical signal.

To avoid reliance on adjusting the coarse-granularity of electricity generation in a power plant, the energy capacity of distribution loops may be capped and the energy may need to be allocated to specific users. While this requirement may not be possible to achieve by the present grid, the CDP grid is suitable to satisfy by using its capability of addressing and controlling the delivery of energy.

In the CDP grid, the forwarding and aggregation of user demands are performed by gateways that process both information and electrical power. The distribution loop may have a large number of users. Each of them is able to receive addressed energy and paired with controlled smart loads.

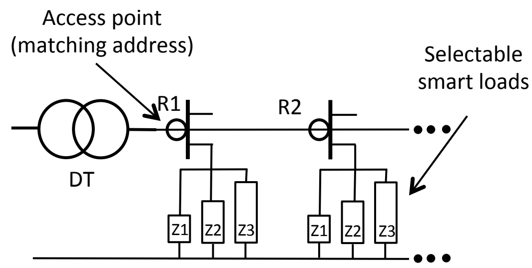

Fig. 1. Smart load selected for each user.

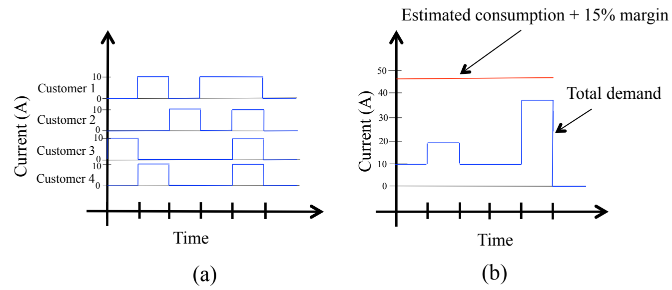

Fig. 2. Example of current demand (a) individual demand of four users, and (b) total demand.

Nodes of the controlled-delivery grid adjust the voltage levels, and forward energy and addresses to the destination end users.

In the controlled-delivery power grid, data is coupled to the power lines. A control node, such as a distributor or a substation of the grid: a) finds the requested energy levels as issued by users (or local distribution points) and assigns the power coming from the generation plants to supply those requests, b) finds routing information about where to forward the energy, and c) attaches the destination address and the amount of current for the supplied power for secure and guaranteed delivery.

A. Distribution Points

Distribution points may transfer energy in two possible directions: 1) from the generator to the user, and 2) from the distribution loop towards other grids (energy re-route). These distribution points also perform energy conversions, including voltage step-up and -down, as they may also be used to interconnect different distribution segments, from generators to user supply. Forwarding of energy from one CDP to others CDP grids is out of the scope of this paper.

Figure 2 shows an example of delivery of discrete energy to users, considering 1-level requests, as used in the CDP grid. We compare it with the present power grid whose capacity is capped at the level of the possible highest demand plus a safety margin. In this example, there are four users, each of them issues a random request for an amount of current. The grid can only provide discrete levels of current, which are 1, 2, 4, or 8 Amps. Figure 2(a) shows the individual requests and Figure 2(b) shows the total current supplied, in discrete levels, within the loop capacity.

Round-robin Selection of Requests. As the distribution loop has limited capacity, the aggregated energy demand may exceed it. In such scenarios, the distributor selects requesting users in a round-robin schedule for satisfying their discrete-amount requests. The selection performed by the distribution point is described as follows:

Phase 1. Each user issues an energy request, if any, to the distribution point in an allowable discrete amount.

Phase 2. The distribution point grants a request if the amount of energy remaining is larger than the requested level (full supply), or if the remaining energy is equal to smaller level or energy (partial supply). Once the remaining amount of energy is zero or smaller than the smallest level, no more user requests are granted, and the distribution point waits until the next cycle starts. Energy is supplied in the following time slot after a request is granted.

III. Preliminary Evaluation of Energy Distribution

We evaluated the efficiency of the management mechanisms for the controlled-delivery power grid [3]. For this, we modeled a distribution loop and the management mechanisms in Matlab for computer simulation. The energy requests of a user were modeled as a two-state (ON-OFF) modulated Markov process. The energy request that is currently OFF becomes ON with probability q, or remains OFF with a probability $1-q$ for one hour, which is a characteristic time window. The request changes its state to OFF with a probability p, or continues in ON state with probability $1-p$. The ratio of satisfied requests, defined as requests whose demand is fully or partially satisfied was then assessed. In addressing a request the program first fulfilled the lower power needs of all users and moved on to fulfill the next level using the round-robin rule. Specifically, the distribution node keeps track of the users served and in the order in which they are served. We considered 1000 users in each loop. The demand from all users in the loop was averaged for 1 year (365 days) and the value obtained served as the average capacity of the loop. Each day was split into three 8-hour intervals, with each interval having a different average request-burst size, with values of 4, 6, and 3 hours, respectively. This approach is very similar to the way power grid providers estimate production levels [20]. The results showed about 99% satisfaction ratio for issue energy requests under scenarios with different number of discrete levels of energy and different loop capacities [3].

IV. Testbeds of the Controlled Delivery Grid

For an experimental evaluation of the efficiency of the power distribution scheduling scheme, we implemented two different testbeds: 1) one with distributed energy storage, where users are implemented by laptops and the energy storage is their battery, and 2) a testbed where the users have no energy storage so that the energy consumption occurs for random periods of time, this testbed was implemented with AC light bulbs, where the lights must be turned on for periods of time (or continuous time slots) to mimic their actual use.

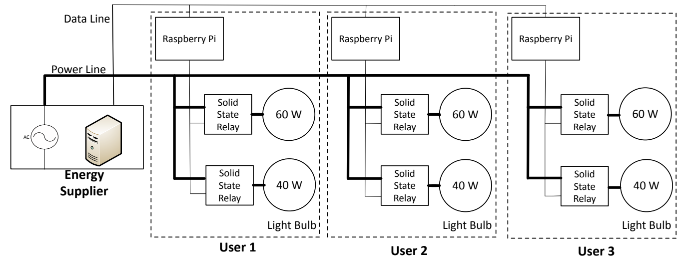

Fig. 3. Diagram of testbed for real-time distribution of energy.

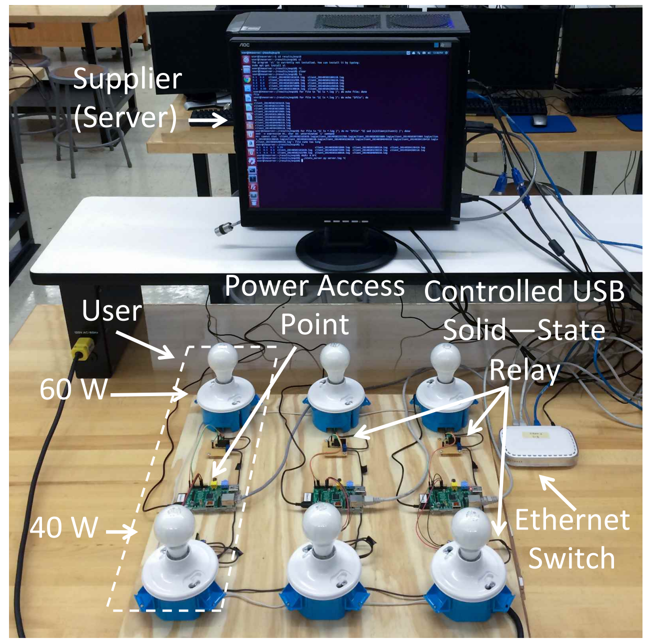

Fig. 4. Diagram of testbed for real-time distribution of energy.

A. Testbed for Distribution of Energy in Real Time

We refer to real-time distribution of electrical power to users (or applications) that demand a continuous supply of energy. Most of today's electrical appliances require this form of supply. In such cases, energy packets must be provided continuously in our proposed controlled-delivery power grid for as long as they are needed, as otherwise interruptions would diminish the use of such applications. To emulate this scenario in a testbed, the used light bulbs require continuous energy supply to function. In general, a user's load in the testbed would be implemented as two light bulbs, 60 and 40 W each, and the user would request 0, 40, 60, and 100 W as energy, selected at random with a uniform distribution. The random function (mimicking a user) and the PAP are implemented in a low-cost computer, Raspberry Pi [21]. A user generates continuous requests for the period of time the light bulb is kept ON. These requests are modeled as a two-state Markov modulated process with average ON state β and average OFF time α. Figure 3 shows the electrical diagram of our testbed and Figure 4 shows the implementation of this testbed.

1) *Request Selection Scheme:* Users and the energy provider follow a protocol to request and grant energy, respectively. As this testbed works with discrete levels of energy, requests and grants are issued in the same discrete amounts. In our testbed, time is slotted. Energy is granted for each time slot (t_s). The time slot duration is $t_s \geq \frac{L}{R}$, where L is the size of a grant packet and R is the transmission rate of the network link carrying the grant packet toward the users' PAP. In our case, we select a time slot duration of 0.5 s while $L = 500$B and $R = 100$ Mb/s. Each time slot, a user sends a grant for a discrete amount of energy to the energy provider with a given probability (as described in Section III). To keep the appliance (a light bulb in our testbed) in ON state, the user continuously issues requests until the appliance is to be turned OFF. Appliances that are OFF issue no requests. The service provider broadcasts a grant every time slot. A grant includes the IP address of each user and the amount of energy granted per user. Therefore, to keep an appliance ON, requests must also be issued and granted each time slot.

The selection of users requesting energy follows a round-robin schedule. There is a pointer that indicates the user with the highest priority, and the grants continue to be issued to those users who follow in the fixed order, as long as the user in turn has issued a request. However, if a user is being granted energy at a given time slot and the user continues to issue requests in following time slots, the provider continues to grant those requests to keep the user (and the corresponding electrical appliance) ON. This service model provides continuous grants for the period of time that a user requests. However, the service may remain OFF even if a user issue several requests (depending on the round-robin schedule and on the energy demand by other users) but once the user requests is granted, it will remain ON until the user stops requesting more energy. In this way, the round-robin order is applied to users that are currently OFF and issue requests.

2) *Experimental Results:* A major advantage of the proposed power grid is that it can control the amount of delivered energy. To demonstrate this, we tested a loop with different power capacities $C = \{100, 140, 200, 240, 280\}$ W. For example, when $C = 140$W, the loop would allow turning on up to three light bulbs (e.g., one 60-W bulb and two 40-W bulbs, or three 40-W bulbs) out of the six available in our testbed.

Figures 5(a) and 5(b) show the *ratio of satisfied requests* and the *ratio of supplied energy* in reference to the requested energy for each of the tested loop capacities. The results show that as the level of requests (i.e., request load) remains constant, satisfaction ratios increase as the loop capacity increases. We categorize request satisfaction ratio and energy satisfaction ratio as the number of request granted, although no necessarily granting the amount of energy requested. For example, if a user requests 100 W and the provider grants 40 W, the request is considered granted. In contrast, the energy satisfaction ratio is defined as the amount of energy requested over the amount of energy supplied. As the loop capacity increases, the satisfaction ratios, for both requests and amounts of energy, increase.

Figure 6 shows a sample (0.99 load) of requests and grants issued during an experiment performed during eight hours. The distribution shows that users issued requests most of the time. Furthermore, as requests for 40, 60, and 100 W are issued with same probability, the numbers of requests for all three are similar. However, as the capacity is limited, most grants are 40 W, and 60 W is also largely granted. Specially, we can note that a large number of 100-W requests are granted 40 W instead.

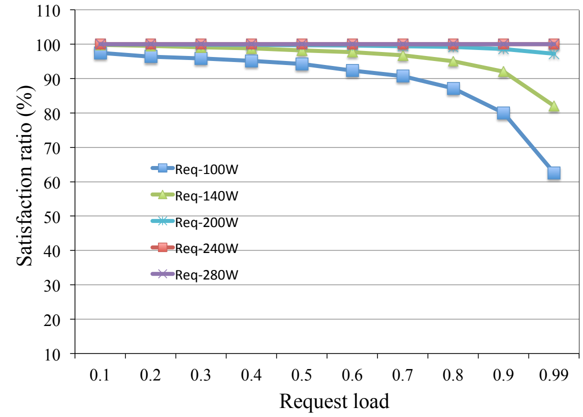

(a) Satisfaction ratio of issued requests.

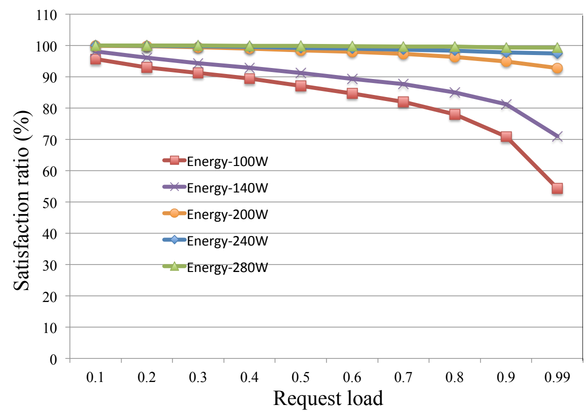

(b) Satisfaction ratio of requested energy.

Fig. 5. Ratio of satisfied requests (those requests that have been fully satisfied) and satisfied amount of energy (energy requested over energy granted).

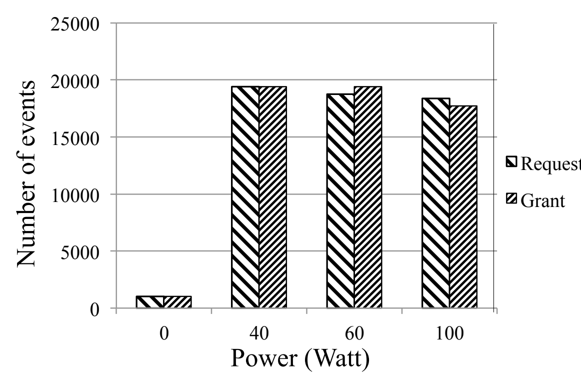

Fig. 6. Example of distribution of requests and grants.

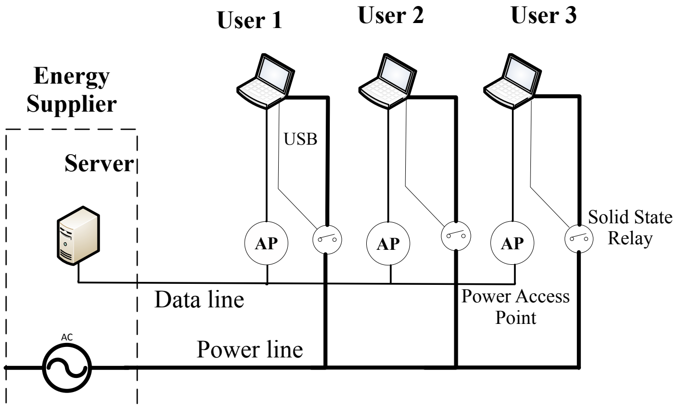

Fig. 7. Diagram of testbed for real-time distribution of energy.

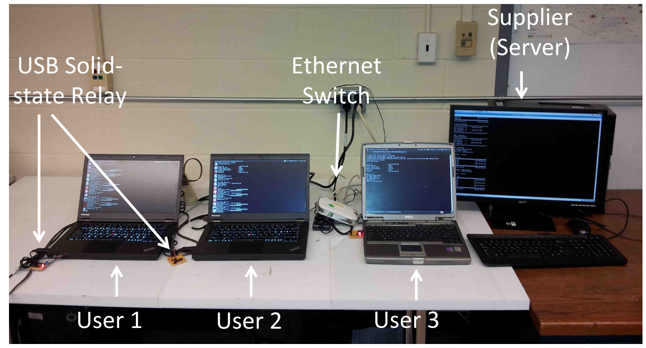

Fig. 8. Diagram of testbed for real-time distribution of energy.

B. Testbed with Distributed Energy Storage

A way maximize the utilization of limited loop capacity, small portions of energy can be provided to users with energy storage. In this way, energy delivery can be scaled up in time and in the number of users for a distribution loop as energy storage would collect the granted small portions of energy, which can be used continuously by different user appliances once the amount of the energy stored is sufficient for that. We implement energy storage by using batteries for laptops mimicking real users. A laptop would issue energy requests at the rate of one per time slot, and the energy provider may issue grants each time slot. However, rather than requesting energy at random, as described in the testbed with no energy storage, laptops request energy after their battery reaches a level of energy in it. Figure 7 shows the diagram of the experimental setup. Figure 8 shows a photograph of the actual setup with laptops connected through a controllable-delivery grid.

1) Implementation of the Testbed with Energy Storage: In this testbed, each laptop communicates with the energy provider through a Gibabit-Ethernet network. The time slot in this testbed must be set long enough to allow the battery to store charge. Therefore, we set a time slot to 60 seconds. The provider issues one grant to each user every time slot. This relaxed timing causes very small traffic load in the data network as the energy provider broadcasts a grant every time slot. Each laptop includes a software implementation of the PAP, which issues energy requests to the provider and checks for destination address and amount of energy granted. To provide access to the power line, a laptop is paired with a solid-state switch, which is controlled by the PAP by using a signal sent through laptop's USB port. Figure 7 shows the diagram of this testbed and Figure 8 shows a picture of the implemented testbed.

2) Energy Requests and Grants: The energy requests and grants in this testbed are different from those used without energy storage. We used two thresholds: a high threshold, Th_h, and a low threshold Th_l. If the level of energy of the battery, B_L is above Th_h ($B_L > Th_h$), a laptop issues no request of energy. If $Th_l < B_L \leq Th_h$ then, a laptop issues a request for energy during half the duration of a time slot. If $B_L \leq Th_l$ then the laptop issues a request for energy for a complete time slot.

Table I summarizes the different energy amounts, in terms of the amount of time with access to the power line, used by each of the laptops. The laptops attempt to keep the energy levels to 70% and request moderate amounts when the battery's level approaches it.

TABLE I
REQUEST SELECTION

Battery level P_b	Request Frequency	Requested Energy
$P_b \geq 70\%$	0 Req/t_s	0 W
$70 > P_b \geq 50\%$	0.5 Req/t_s	0.5 W
$P_b < 50$	1 Req/t_s	1 W

The amount of energy requested to the supplier depends on the frequency in which requests are issued, the amount of power (which is converted into a period of time a laptop and its battery connect to the power line) requested in each request, and the number of users. It is expected that as the battery energy level approaches the target energy level in the battery (e.g., 70%), a laptop may request less energy until it stops. As the battery's energy depletes, the laptop starts requesting energy again. In the request scheme, we target a battery energy level of 70% charge, which is reached after stabilizing. It should be noted that the target energy level of the request scheme is to keep the battery charged at that level, rather than making energy consumption highly efficient. A scheme can be tailored towards this target but that is out of the scope of this paper.

The grants in this testbed follow a round-robin schedule, whether they are for requested for 0.5 or 1 W, one every time slot. Since there is an energy storage device, there is no priority on requests or users that have been granted, and each time slot, a set of user grants are selected anew. Therefore, it is possible to provide energy to the battery of the laptops intermittently. The use of energy storage simplifies the grant scheduling scheme.

3) Experimental Results: The experimental results are shown in Table II. These results show a high satisfaction ratio in the number of requests (RSR) and the energy satisfaction ratio (ESR). Users 1 and 2 have fast-charging batteries so they need fewer time slots to achieve the 70% level goal.

They issue fewer requests than User 3. The RSR is 100% in most cases. As for User 3, the battery is less efficient and it needs a larger number of time slots to achieve this goal, and yet the RSR is 98.84%. As for the ESR, these users spend energy during the experiment. The time in which they detect battery levels below 70% accounts for the difference to achieve 100%. However, the achieved ESR is over 99% for User 1 and about 99% for the other two users. The charge of the batteries of all three users eventually reaches and stays around 70%, using only 2 W per minute. It should be noticed that without the controlled-delivery grid, the power line would have 3 W available anytime. As the number of users increases, the efficiency of the controlled-delivery grid increases.

TABLE II
RESULTS

Loop Capacity (W)	Target Charge	User Identification					
		1		2		3	
		RSR	ESR	RSR	ESR	RSR	ESR
2	70%	100	99.72	100	98.17	98.84	98.45

V. CONCLUSIONS

We show the operation of the controlled delivery grid under on two testbeds: one with energy storage, emulated with laptops and their batteries, and another without energy storage, shown by the supply to incandescent light bulbs. The latter testbed mimics the delivery experienced in the present grid; continuous power delivery. However, the testbed uses of a request-grant protocol, enabled by a data network. We show the experimental performance of the two testbeds in terms of the capability of satisfying users' requests, or satisfaction ratio. The results show high satisfaction ratios for loop capacities that are smaller than the maximum energy requested, showing the advantages of a grid with fine energy allocation. Our results are consistent with the theoretical evaluations presented before [3].

REFERENCES

[1] Y. Xu, R. Rojas-Cessa, and H. Grebel, "Allocation of discrete energy on a cloud-computing datacenter using a digital power grid," in *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*. IEEE, 2012, pp. 615–618.

[2] R. Rojas-Cessa, Y. Xu, and H. Grebel, "Management of a power grid with controllable delivery of discrete levels of energy," in *IEEE Electrical Power and Energy Conference*. IEEE, 2013, pp. 1–5.

[3] ——, "Management of a smart grid with controlled-delivery of discrete power levels," in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1–6.

[4] S. Pahwa, A. Hodges, C. Scoglio, and S. Wood, "Topological analysis of the power grid and mitigation strategies against cascading failures," in *Systems Conference, 2010 4th Annual IEEE*, april 2010, pp. 272 –276.

[5] R. Gono, S. Rusek, and M. Kratky, "Reliability analysis of distribution networks," in *Electrical Power Quality and Utilisation, 2007. EPQU 2007. 9th International Conference on*, oct. 2007, pp. 1 –5.

[6] H. He, "Toward a smart grid: Integration of computational intelligence into power grid," in *Neural Networks (IJCNN), The 2010 International Joint Conference on*, july 2010, pp. 1 –6.

[7] S. Galli, A. Scaglione, and Z. Wang, "For the grid and through the grid: The role of power line communications in the smart grid," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 998 –1027, june 2011.

[8] W.-H. Liu, K. Liu, and D. Pearson, "Consumer-centric smart grid," in *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, jan. 2011, pp. 1 –6.

[9] K. Budka, J. Deshpande, J. Hobby, Y.-J. Kim, V. Kolesnikov, W. Lee, T. Reddington, M. Thottan, C. White, J.-I. Choi, J. Hong, J. Kim, W. Ko, Y.-W. Nam, and S.-Y. Sohn, "Geri - bell labs smart grid research focus: Economic modeling, networking, and security amp; privacy," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, oct. 2010, pp. 208 –213.

[10] G. Lu, D. De, and W.-Z. Song, "Smartgridlab: A laboratory-based smart grid testbed," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, oct. 2010, pp. 143 –148.

[11] W.-Y. Yu, V.-W. Soo, M.-S. Tsai, and Y.-B. Peng, "Coordinating a society of switch agents for power distribution service restoration in a smart grid," in *Intelligent System Application to Power Systems (ISAP), 2011 16th International Conference on*, sept. 2011, pp. 1 –7.

[12] R. Abe, H. Taoka, and D. McQuilkin, "Digital grid: Communicative electrical grids of the future," *Smart Grid, IEEE Transactions on*, vol. 2, no. 2, pp. 399 –410, june 2011.

[13] F. Bouhafs, M. Mackay, and M. Merabti, "Links to the future: Communication requirements and challenges in the smart grid," *Power and Energy Magazine, IEEE*, vol. 10, no. 1, pp. 24 –32, jan.-feb. 2012.

[14] Z. Fan, P. Kulkarni, S. Gormus, C. Efthymiou, G. Kalogridis, M. Sooriyabandara, Z. Zhu, S. Lambotharan, and W. Chin, "Smart grid communications: Overview of research challenges, solutions, and standardization activities," *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1 –18, 2012.

[15] T. Takuno, M. Koyama, and T. Hikihara, "In-home power distribution systems by circuit switching and power packet dispatching," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, oct. 2010, pp. 427 –430.

[16] M. M. He, E. M. Reutzel, X. Jiang, R. H. Katz, S. R. Sanders, D. E. Culler, and K. Lutz, "An architecture for local energy generation, distribution, and sharing," in *IEEE Energy2030*, Nov. 2008, pp. 1 –6.

[17] T. Takuno, Y. Kitamori, R. Takahashi, and T. Hikihara, "Ac power routing system in home based on demand and supply utilizing distributed power sources," *Energy*, vol. 4, no. 5, pp. 717–726, june 2011.

[18] M. Alizadeh, A. Scaglione, and R. J. Thomas, "From packet to power switching: Digital direct load scheduling," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 6, pp. 1027–1036, 2012.

[19] A. J. Viterbi *et al.*, *CDMA: Principles of spread spectrum communication*. Addison-Wesley Reading, 1995, vol. 129.

[20] S. Shao, T. Zhang, M. Pipattanasomporn, and S. Rahman, "Impact of tou rates on distribution load shapes in a smart grid with phev penetration," in *Transmission and Distribution Conference and Exposition, 2010 IEEE PES*. IEEE, 2010, pp. 1–6.

[21] "Raspberry pi," http://www.raspberrypi.org/.