A Digital Approach to Energy Networks; Allocation and Distribution of Energy Requests

Camila Fukuda, Henrique Pita, Roberto Rojas-Cessa and Haim Grebel
ECE Department,
The New Jersey Institute of Technology,
Newark, NJ 07102, USA
grebel@njit.edu

Ahmed Mohamed
Grove School of Engineering, Department of Electrical Engineering
City College of the City University of New York,
New York, N.Y. 10031, USA

***Abstract*— In the current grid, power is available at all times, to all users, indiscriminately. This makes the grid vulnerable to sporadic demands and much effort has been invested to mitigate their effect. We offer here a digital approach to power distribution: an energy-on-demand approach in which the user initiates an energy request to the server of the energy provider before receiving the energy. Considering a micro-grid with a mix of generators (sustainable and other sources), the server optimizes the entire power network before granting the energy requests, fully or partially. The energy is packetized and is routed to the user's address by an array of switches. For example, in an office building, the energy provider may queue energy requests by some air-condition units and grant these requests later. During recovery from a blackout, pockets of instability may be isolated by their unusual energy demands. In its simplest form, this network can be realized by overlaying an auxiliary (control, or, data) network on top of an energy delivery network and coupling the two through an array of addressable digital power switches. In assessing this approach, we are concentrating in this paper on the management of energy requests by using statistical models. An energy network with a limited channel capacity and the optimal path for energy flow in a standard IEEE 39 bus are considered.**

Keywords—Energy Networks; Energy management and Distribution; The Digital Grid

I. The Digital Approach to Energy Networks

At the present time, loads are determining the level of current consumed by the grid. A large effort has been invested in sensing the change in frequency and phase of this (analog) power grid; otherwise, blackouts may occur. Our approach uses a demand-supply management model [1-4]: (1) Users (or loads) issue energy requests (energy equals power over time). (2) The service provider optimizes the energy allocation prior to the energy dissemination. (3) The service provider then allots the energy to selected users on a cycle-by cycle basis. The cycle duration may vary and depends the grid's reaction time. While this seems to pose a heavy burden on the grid's communication, the protocol has a stabilizing effect on a limited and intermittent energy supply. Energy storage elements become integral part of the network and their inclusion is simpler than interfacing them with the currently deployed power grid. Owing to the request-grant allocation protocol, energy demands and energy consumptions may be closely monitored and safety margins can be optimized (and minimized) for any given moment, thus increasing the overall power network's efficiency. The advantage of such approach can be demonstrated during the recovery of power networks from a blackout. The issues that brought the grid down still exist and would be unknown in the present power grid framework. However, by analyzing all energy requests, the energy provider can isolate pockets of vulnerabilities before the energy is distributed. The approach is wholly and considers every aspect of the network: power generation, distribution, and usage.

Fusing data and energy [5-8] in discrete formats dramatically reduces management complexity because, in principle, the energy (power delivered over time) can be directed to specific users. Packetizing energy is a new concept that is uncommon to the power community but not to the information network community. It does not violate Kirchhoff's laws as proven by numerous information networks and we propose to realize it within the micro-grid framework. The digitation of time (through allocation of energy at varying time slots), or the digitation of power (to be delivered by discrete current levels while keeping the voltage constant) are but two possible energy digitization approaches. The simplest adaptation of such Energy Networks is by interfacing a power network with an auxiliary data network that opens and closes power switches along the energy routs to addressed users (Figure 1(a)). The energy supplier selects the appropriate smart load through data fused to the energy packet. Figure 1(b) shows the synergistic operation of data and energy where a data network provides the management and control of the power network.

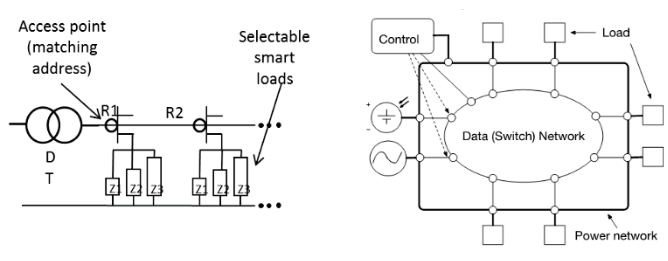

(a) (b)

Figure 1. (a) Smart limiters and (b) the simplest framework of two overlaid networks: the data (switch) network, which is coupled to the power network through controlled smart loads. Also shown are AC and solar power sources.

This work was funded in part by NSF Grant CNS-1641033.

An alternative energy source (e.g., a solar panel) may be incorporated into the distribution loop as yet another address (see path allocation below). The data network manages the communications between users and the distribution point, energy delivery, and the management of the two energy sources at a much higher frequency than the frequency used to transmit the electrical power itself.

The focus of this paper is on optimization of energy delivery when using this new approach. We ran statistical models to answer the following question: how many request cycles are needed before a customer is satisfied under limited power conditions? Answering this question will help us figuring out the optimal cap for energy networks.

II. Latency and Storage Issues

For simulation purposes users are requesting energy, randomly. Each user has an address so a specific user's behavior can be followed. We assign a probability to users who request energy to turn their appliance ON (meaning they start with an appliance OFF) and another probability for those users who have their equipment already ON and wish to continue to do so. In this way we can simulate the user behavior at any given cycle. When the energy requests pass a certain cap, some of the requests will be sent to a queue. Specifically, the micro-grid operates as the currently deployed grid if it can satisfy all requests without sending any to the queue. An example is provided in Figure 2: two random numbers are generated for each user. For those users that were OFF in the previous round, we check whether the randomly generated number, p_{req}, is smaller than a given request probability, $p_{request}$. If yes, then a new random number is generated for the actual energy request. The second randomly generated number is used for those users whose equipment is already ON. If the randomly generated number, p_{on}, is larger than p_{stay_on}, then their new request will be 0 (they will be turned OFF). Otherwise, they will remain ON with their previous energy request. In this way we minimize a succession of turning the power ON and OFF for those users who consume relatively large amount of power. Unsatisfied energy requests are sent to the queue. (As a note, one could generate a single random probability number and compare it to the $p_{request}$ and p_{on} for the two groups involved as in step 2 of Figure 2: the group with its power ON in the previous step and the group with its power OFF in the previous step. Both approaches yielded similar results). Finally, we point out that the process is not entirely Markovian; the queue memorizes the size and the order of the requested energy until satisfied or until the request is dropped. Specifically, a 2-state Markovian chain is an adequate analytical model for ON and OFF states as long as the queue is empty. A three-state analytical Markovian chain has a 6% discrepancy with the numerical results because the queue has a selection rule and does not accept the energy requests randomly. Specifically, we considered two examples: satisfy the large energy requests first (hence sending the reaming small amounts to the queue) and separately, satisfy small energy requests first (hence sending the large energy requests to the queue).

```
Step 1: (ON=logical(e_requested)

preq=rand(1,numel(e_requested));                     %Generate number to compare with p_request
pon=rand(1,numel(e_requested));                      %Generate a number to compare with
p_stay_on

Step 2:
e_requested(not(ON) & preq<p_request)=rand(1,numel(e_requested(e_requested==0 &
preq<p_request)));                                   %Generate a random request when turning on
e_requested(ON & pon>p_stay_on)=0;                   %Turn off
```

Figure 2. An example of simulating steps in MatLab for the Digital Power Networks.

Currently, micro-grids are limited by the overall generated power. Predictive models and years of data collection aid the utilities in forecasting the level of service during the hour of a day and the month for the entire grid. In addition, phase sensors monitor the frequency at which the grid operates. We have a modest goal in mind and consider only a subset of the entire grid. Yet, and contrary to current grid design we are moving one step further by adjusting the allotted energy according to actual energy requests. As a leading example, consider an office building with many air-condition (A/C) units. Each unit places a request for energy. Since starting an A/C motor requires as much as 8 times more current than its steady states consumption, the energy server may delay some requests for a marginal service disruption. In our simulations we set an energy cap for the total energy available per request cycle (round). As the two probabilities become larger, more users (or for that matter, addresses) turn their equipment ON and more users whose equipment is already ON remain ON. That situation puts an unsustainable burden on the power network - the power network cannot satisfy all users at the same time and unsatisfied requests ought to be sent to the queue. In the current analog grid, the system will become overloaded and fail.

A finite probability to stay in the queue provides us with the freedom to let the user choose a limited time frame during which energy is still needed. Thus, if that probability to stay in the queue is small, then the energy request will eventually be dropped from the queue and the number of remaining requests could be satisfied more quickly. If, on the other hand, the probability of staying in the queue is large (probability 1 for staying in the queue until the request is satisfied), then the waiting period may be prolonged. If the energy cap is large enough and all requests are satisfied, then there will be no queued requests. In this case, the Energy Network behaves similarly the current power grid, yet with a direct knowledge of the grid's status during the request cycle and full control over the energy flow (Figure 3,4).

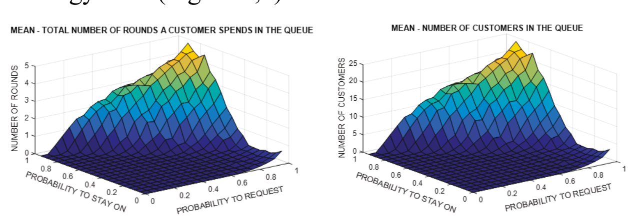

(a) (b)

Figure 3. (a) Mean numbers of rounds and (b) users (customers) waiting in the queue when the probability of waiting in the queue is 0.1. Here, we consider a total number of users to be 500. The channel capacity (energy cap) was set to 150 units of energy and each user could ask for up to one unit of energy.

When the probability of staying in queue increases to 0.5, the waiting period (in number of rounds) and the number of users waiting in the queue would obviously increase, as shown in Figure 4. While the probability to stay in the queue increased 5 times, the maximum number of cycles in the queue has not increased as much; the Energy Network was able to distribute the requests quite efficiently.

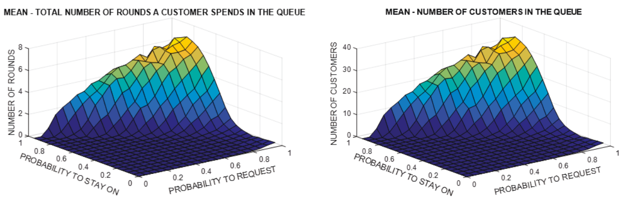

(a) (b)

Figure 4. (a) Mean of the total number of rounds that a user spends in the queue and (b) the mean number users waiting in the queue, when the probability of waiting in the queue is 0.5. This means that more users are waiting in the queue compared to the case presented in Figure 3. The total number of users is 500 and the cap is set to 150 units, where each user may ask for up to one unit of energy.

There are two approaches to decide whose request is served first. In the one used to generate Figures 3-4 we satisfy the smaller energy requests first and send the overflow of requests to the queue. In the one used to generate Table 1 and Figure 5 we satisfied the largest energy requests first and send the overflow requests to the queue. The two approaches would result in different queue time and number of requests waiting to be satisfied. Satisfying the smaller energy requests first, would results in smaller number of queued and large energy requests albeit with prolonged waiting periods. Our simulations take into account random energy requests for each customer (Figure 5). This is a simple case of a power grid with capacity of 250 energy units and 500 customers. Each round presents one cycle of energy requests (say every 0.5 second). For most cases, the Energy Network can satisfy all users except for the case where the probability to stay ON approaches 1. Thus, for most scenarios (probabilities) the situation is very similar to the present grid. However, and unlike the present grid, overflow of users' demands is placed in a queue and would not overwhelm the entire grid, thus avoiding blackouts.

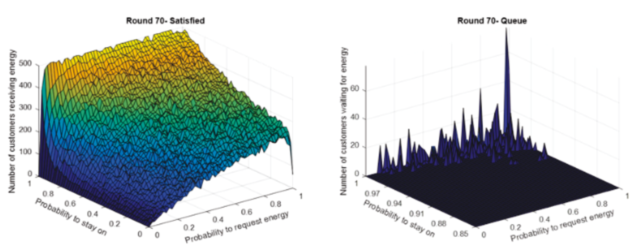

(a) (b)

Figure 5. (a) In most cases the digital grid accommodates all energy requests. When the demand is large (probability to stay ON is close to 1) some users are sent to the queue. (b) Number of users waiting in the queue. The scale is focused on the large probability range for staying ON.

Having the ability to schedule the delivered energy makes us wonder if adding a battery to the ENERGY NETWORK would decrease the waiting time for a waiting/demanding user [9-10]. The short answer is NO - even with a very large battery, the slow charge/discharge cycles may not cope with the fast and randomly distributed demand fluctuations. As we have shown earlier [3] one requires a minimum time of a least 1 minute to meaningfully charge a laptop battery. However, if every customer is equipped with a super-capacitor/battery system, and if all of these energy storage units are at the service of the entire grid (namely, they may deliver energy to other users, as well), then one may imagine that not only we may achieve fast charge/discharge cycles but also a more regulated grid. Such approach, which we dub 'cloud energy storage,' has the effect of increasing the energy cap for the power network. While fast charging and discharging of large amounts of electrical energy make super-capacitors ideal for short-term energy storage the amount of energy stored is rather limited with today's technology. *The Energy Network takes a wholly approach to storage and delivery of energy on an energy cycle basis.*

When the energy cap is increased to 2/3 of the maximum users' requests the digital grid can accommodate the random energy requests without the need for a battery and without sending users' requests to the queue. Our simulations showed that batteries will be charged and will stay charged all the time, independently of the probabilities of staying ON or OFF.

In Table 1 we provide a snapshot of one time slot (one request cycle, or one round) of a 10-user digital micro-grid. These 10 users may tap into an energy storage while their energy grant is pending. Each user may consume no more than one unit of energy (but could consume less). Let us consider an energy cap on the energy consumption of 3 energy units. The probability to switch from OFF to ON and to continue to request a service are both 0.3. These switching probabilities are fairly low but the energy cap is also very low; it is 30% of the maximum consumed energy by all users. The maximum stored energy in the battery is set to 10% of the maximum total energy (1 unit in total) that may be consumed in one cycle. Let us consider Users 2 and 7 in Table 1. The users request various amounts of energy at some point in time. Because the energy cap is only 3 units of energy, the system cannot accommodate all requests and some users are put on hold (namely, their requests are queued). Priority in this case was given the largest energy requests in reverse order thus the smallest request was queued. In order to avoid long delays in the energy supply, User 2 taps into a battery resource.

Users:	1	2	3	4	5	6	7	8	9	10	Total
Request	0.4974	0.4869	0	0.5473	0	0	0.5221	0	0.9519	0	3.0056
Grant	0.4974	0	0	0.5473	0	0	0.5221	0	0.9519	0	2.5187
Queued	-	1	-	-	-	-	-	-	-	-	1
Storage	0	0.4869	0	0	0	0	0	0	0	0	0.4869

Table 1. A snapshot of requested, granted, queued, and stored energy for a 10-user micro-grid network.

The energy storage may be physically situated on the user premise or shared amongst all users (cloud energy storage). If there is an excess of energy in the system, then the energy storage is charged. Fast charge/discharge of the energy storage is required because the requests are varying for each round.

III. Optimization Scenarios – Energy Storage

Here, the energy storage acts as a secondary energy source, and the power network provides for the primary source of energy. The energy storage is treated as an addressed user when extra energy, left by the optimization process, and energy is routed in its direction. The energy storage stores up to 10% of grid's capacity. Let us also consider that only designated users, which typically is set as 10% of the total number of users, are allowed to tap the energy storage. When the primary source cannot offer energy to the users, they enter the queue and try to tap into the secondary source. These users do not leave the queue because there is no guarantee that the storage has energy during the next round. Below, we have used an optimization approach that minimizes the wait in the queue.

There are several ways to handle the energy requests, and thus, the optimization of the energy flow. One family of optimization algorithms is the genetic algorithm. It is based on a learnt process (namely, collecting data for several request cycles), yet allows for random events to happen.

A. Genetic Algorithms

Genetic Algorithms optimize the energy allocation by using a few simple rules:

1. Selection: select parents from population (energy users) for the next generation of solutions (children). In our case we fit the incoming small energy requests first, moving on to the larger requests until we reach the channel capacity limit. Unsatisfied large energy requests are sent to the queue. Obviously, instead, one can accommodate the largest energy requests first, or any combination of the above. In a large pool of users and limiting the maximum demand to one energy unit the selection process does not significantly change the outcome.

2. Crossover: just like in biology, the characteristics of the parents in each generation are "mixed" into several possible solutions (population/children).

3. Mutation: random changes happen in the characteristics of the parents to generate the children.

The genetic algorithms are learnt algorithms; for random processes they reach optimization after several rounds and thus, the process was run for 100 or 50 cycles to reach an averaged solution and run for 20 more cycles to reach a higher degree of optimization. Time wise, it translated to a longer computing times. It took 1 sec to handle 500 customers to reach a non-optimized solution on a laptop computer. It took 283 sec for the genetic algorithm to reach a much better optimized energy distribution among 500 users while using the same laptop machine.

The code is part of optimization package of Matlab. We considered the following experiment setup: number of users – 500; channel capacity – 200 units (maximum 1 unit per user); storage capacity of 20 energy units; the probability of a request to stay in the queue equals to 1; the probability to change status in the queue is equal to 0; and the number of preferred users that can tap into energy storage elements was 40. The results are shown in Figure 6.

Since the requests are randomly changing from one energy request cycle to another, it is better to present both the mean and its standard deviation. In Figure 6 we show 500 users in a power network with a capacity of 200 energy units. The mean value well represents the situation at low probabilities. When the network reaches its full capacity, the variation in the number of users sent to the queue obviously increases.

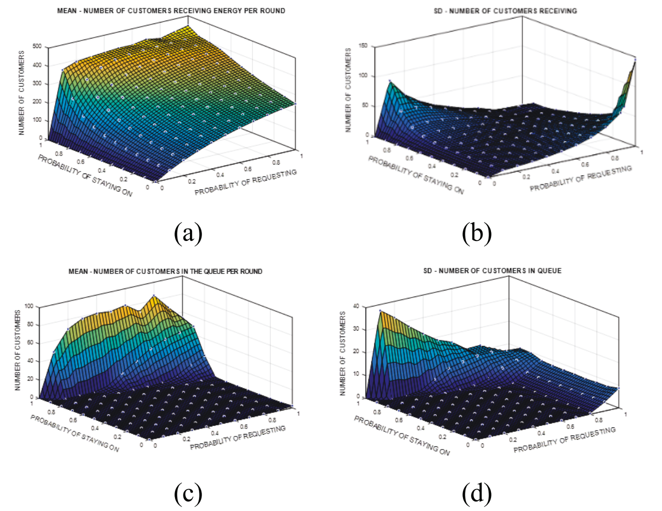

Figure 6. (a) Mean number of users receiving energy as a function of probability to stay ON and as a function of the probability of energy requests. (c) Standard deviation of number of customers receiving energy. (b) Mean number of requests in the queue and its standard deviation (d). In obtaining the mean values, the program was running for 20 simulations and collected information during 50 cycles (rounds) each.

A direct comparison for cases with and without the energy storage was made. We found out that the impact of the battery on the digital power network is relatively small; the energy storage element holds relatively small amount of energy and needs to charge part of the time. A larger impact is achieved if the battery is interfaced with a solar panel of equal energy capacity. The charging of the battery is not made on the expense of the grid and what extra solar energy left is added to the network energy capacity (Table 2). An optimized solution is consider in Section 4.

IV. Optimization Scenarios - Solar Energy with Battery

There are several ways that extra energy from a storage element and/or sustainable sources may be utilized by the power network. Solar energy may be added to the overall network resources as is done today albeit with the caveat that its availability is not certain (that is the reason why we assign a finite probability to its energy supply). In the example (Table 2 below) we interfaced a solar panel with a battery. The reason being that while the capacity of the solar energy source is relatively small, it can augment the 'standard' micro-grid in times of need. We allow for only special customers to tap into this resource. The simulations conditions are: number of simulations was 50; number of cycles (rounds of time slots) – 50; number of customers was 500; number of special customers – 50; energy cap (or channel capacity) – 100; the probability to stay in the queue was p=1; the probability to change your queue status was p=0; battery capacity was 10 energy units. Each user may request up to 1 energy unit. The comparison was made for a specific probability to stay ON was p=0.5 and the probability to request energy if the user was at OFF state as, p=0.5. The cap for the solar energy source was 10 energy units.

	WITH SOLAR ENERGY AND NO OPTIMIZATION	WITH SOLAR ENERGY AND OPTIMIZATION
Energy distributed per round:	99.53	99.978
Energy requested per round:	189.80	148.23
Number of customers in the queue per round:	92.29	96.73
Number of customers that received energy per round:	203.10	200.356
Number of customers that requested energy per round:	295.4	297.11
Number of customers that were satisfied in the queue per round:	4.83	47.73
Energy distributed by the solar energy system per round:	2.417	2.5571
Energy available in the battery of the solar energy system per round:	0.897	0.0020
Solar energy produced per round:	2.443	2.5571
Number of customers that requested solar energy per round:	9.312	49.9740
Number of customers that received solar energy per round:	2.476	6.49
Total energy delivered (Solar+Grid) per round:	101.9	102.55
Number of rounds a customer in the queue waits to be satisfied:	10.56	2.0856
Number of rounds a customer spends in the queue	9.229	9.57
Wait time to receive energy from the solar system:	0.2033	6.3794

Table 2. A power network interfaced with a power storage): the comparison is made with and without optimization). The main advantage for the optimized solution is the ability to accommodate more users in the queue and the decrease in the queueing time. The time is measured in cycles (rounds, or time slots). 'Energy requested per round' includes new and queued energy requests – it becomes smaller for the overall optimized solution.

V. Optimization Scenarios – The Path of Energy Flow

Optimization of the power flow should not only include time but also space (path). It is difficult to divert energy in the currently deployed grid (albeit it is possible). With distributed energy resources (DER) in mind we have considered a test bus system composed of a mixed generators sustainable sources and users. In Figure 7 below, we present the statistics of connecting several alternative sources to several users using the IEEE 39 test bus system. The bus is made of sources (green nodes), users who receive energy (orange nodes), users who do not ask for energy (light yellow nodes), users who are in queue (red nodes), path-through users that do not tap into the energy flowing through them (blue nodes) and energy flow paths (light blue arrows). The program searches for the minimal path (the Dijkstra's method) to determine which source will be used to which user. There was no limit on the source or the path capacity, however, there was a limit of 5 users per source. Extra users were directed to another source nearby. The total energy delivered to the users could not be larger than the global system capacity. For simplicity, an average energy loss of 6% per path was considered.

The statistics changes at each request cycle (round, or time slot) and we present here a snapshot of a randomly chosen cycle. For simplicity, we included the probability of connecting the nodes but not their associated loss. As we can see from the figure, the scenario is rather complex; some nodes forward energy, others consume it, and some nodes play several roles, such as generators and consumers. Such scenarios may be important for a system with sustainable generators with storage elements; since the energy from such sources may change in time, the conventional grid and energy storage elements would supplement (even momentarily) the needed energy.

VI. Summary

We have proposed a request-grant protocol to control energy networks. It was shown that this protocol handles well energy demands by sending over-the-cap requests to a queue. This paper outlined how to handle such queued requests. While requiring an auxiliary power switching array, the Energy Network concept nonetheless mitigate power fluctuations and incorporates sustainable sources in a seamless fashion. Adaptation of such approach to the smart homes and to very large Internet of Things (IoT) is an exciting possibility.

References

[1] Xu, Y., Rojas-Cessa, R., & Grebel, H. (2012, November). 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 (pp. 615-618). IEEE.

[2] Rojas-Cessa, R., Xu, Y. and Grebel, H. (2013, August). Management of a smart grid with controlled-delivery of discrete levels of energy. In Electrical Power & Energy Conference (EPEC), 2013 IEEE (pp. 1-6). IEEE.

[3] Rojas-Cessa, R., Sahasrabudhe, V., Miglio, E., Balineni, D., Kurylo, J. and Grebel, H. (2014, November). Testbed evaluations of a controlled-delivery power grid. In Smart Grid Communications (SmartGridComm), 2014 IEEE International Conference on (pp. 206-211). IEEE.

[4] 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.

[5] F. Blaabjerg, R. Teodorescu, M. Liserre, and A. V. Timbus, “Overview of control and grid synchronization for distributed power generation systems,” Industrial Electronics, IEEE Transactions on, vol. 53, no. 5, pp. 1398–1409, 2006.

[6] S. H. Strogatz, “Exploring complex networks,” Nature, vol. 410, no. 6825, pp. 268–276, 2001.

[7] A. Mohamed and O. Mohammed, “Real-time energy management scheme for hybrid renewable energy systems in smart grid applications,” Electric Power Systems Research, vol. 96, pp. 133–143, 2013.

[8] Abe, R., Taoka, H., & McQuilkin, D. (2011). Digital grid: Communicative electrical grids of the future. IEEE Transactions on Smart Grid, 2(2), 399-410.

[9] Mohd, A., Ortjohann, E., Schmelter, A., Hamsic, N., and Morton, D. (2008, June). Challenges in integrating distributed energy storage systems into future smart grid. In 2008 IEEE international symposium on industrial electronics (pp. 1627-1632). IEEE.

[10] Joseph, A., & Shahidehpour, M. (2006, October). Battery storage systems in electric power systems. In 2006 IEEE Power Engineering Society General Meeting (pp. 8-pp). IEEE.

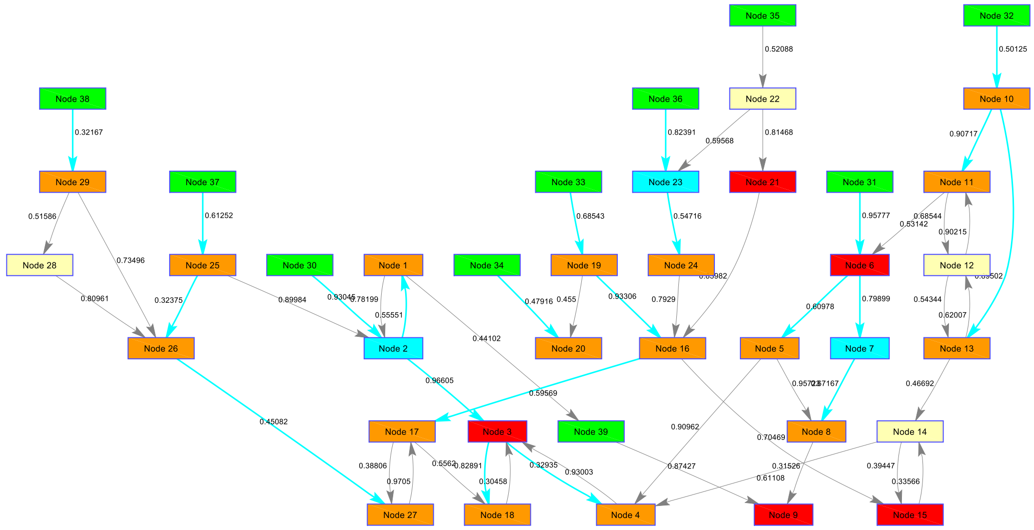

(a)

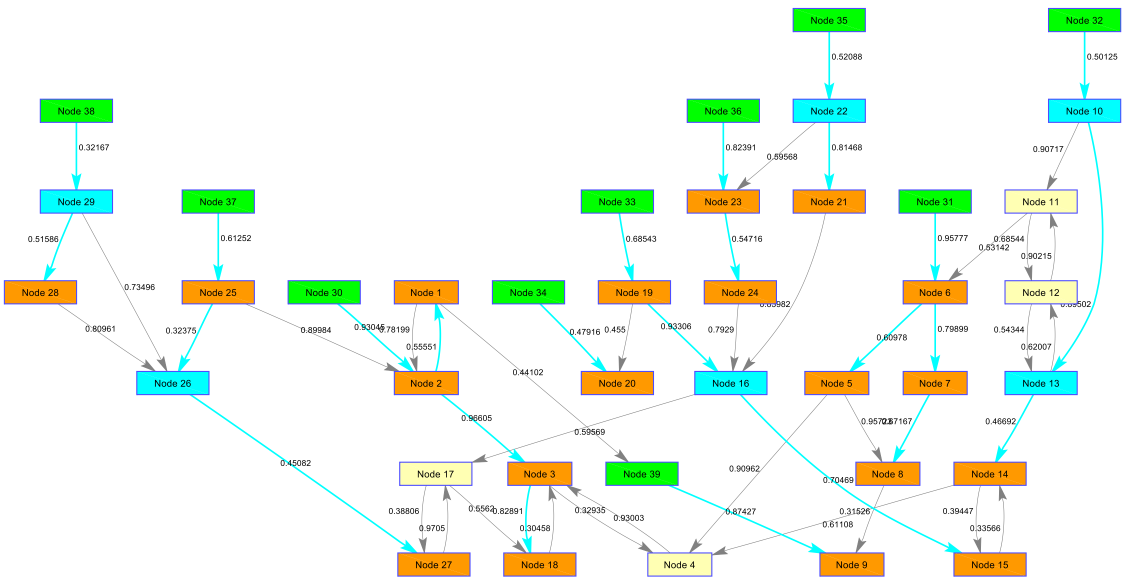

(b)

Figure 7. (a) Simulating energy distribution in the IEEE 39 bus system. This is a snapshot at some particular round. Each number assigned to a dark gray arrow is the probability of a connected path. Red nodes: Users waiting in the queue; Green nodes: Sources; Light Yellow nodes: Users not requesting energy; Orange nodes: Users receiving energy; Blue arrows: Energy Path; Blue nodes: Users where energy is flowing through without tapping into it. (b) A situation where no node is in the queue