Delivering Smart Load-shedding for Highly-stressed Grids

Noman Bashir*, Zohaib Sharani†, Khushboo Qayyum†, Affan A. Syed†
*Center for Advance Studies in Energy
National University of Science and Technology, Islamabad, Pakistan
†SysNet Lab, National University of Computer and Emerging Sciences, Islamabad, Pakistan
Email: first.last@sysnet.org.pk

***Abstract*—This paper targets the unexplored problem of demand response within the context of power-grids that are allowed to regularly enforce blackouts as a mean to balance supply with demand: *highly-stressed grids*. Currently these utilities use a cyclic and binary (power/no-power) schedule over consumer groups leading to significant wastage of capacity and long hours of no-power. We present here a novel building DLC system, Aashiyana, that can enforce several user-defined low-power states. We evaluate distributed and centralized load-shedding schemes using Aashiyana that can, compared to current load-shedding strategy, reduce the number of homes with *no* power by $> 80\%$ for minor change in the fraction of homes with full-power.**

I. Introduction

Demand response (DR) is a smart-grid technology allowing grid to communicate a demand decrease request to meet supply, against traditional supply-following load behavior, using indirect (pricing) or direct (through some control) signals. Utilities prefer Direct Load Control (DLC) as it gives guaranteed reduction but is difficult for user; pricing signals leave the customer in charge but have an uncertain and possibly time-delayed demand reduction [1]. Pricing signal can also result in secondary peaks due to behavioral shift [2].

We however argue that most DLC work has focused on *over-provisioned* grid systems of developed countries, with a focus on increasing revenue and reliability [3], but remains largely blind to the unique characteristics of *highly-stressed grids* of countries (like Pakistan, Nepal, and India) with a very large and nearly continuous supply-demand gap. As an example, for Pakistan, this gap can be as high as 6GW during summers, but stays around 1.2GW even during the winter months (2011-2012) [4]. The (largely national) utilities in these countries enforce periodic events of controlled blackouts, or load-shedding, to relieve this stress. Existing DLC mechanism, in trying to balance consumer comfort with *some* reduction employ fine-grained (in both time and amount of load-shed) load-control, especially at residential homes [5], [6]. Such DLC allows for control events, like changing HVAC set-points, or possibly for controlling the AC for a few hours a day with over-ride facility [5]. While these mechanism are quite useful in shaving off consumption peaks and preventing peaker plants from running (thus saving money), they are **inadequate in their magnitude as well as flexibility** for managing the large and continuous gaps that exist in highly-stressed grids.

We believe that the consumers in a highly-stressed grid — being acclimatized to frequent blackouts — are much more amenable to aggressive DLC mechanisms and thus willing to accept a wider-range of load-shedding policies. This demand reduction, however, will have to be done through some automated system as users cannot be expected to manually respond to any, potentially large, load reduction signal. *We thus propose instrumenting homes with a system that provides utilities with transitions to several* ***low-power states*** *that map to user-specified appliances.*

In this paper we design and evaluate a novel and practical home-level DLC system solution, *Aashiyana*, that can implement several user-configurable power-states of a home. This system is practical as it can retrofit into the existing wiring scheme of homes; is of low cost while controlling most appliances in a home; provides home consumers a flexible way to describe these lower-power states as a compact disconnectivity matrix requiring one-time configuration.

A question remains regarding incentives for power utilities to promote a proliferation of the Aashiyana DLC system when their current strategy of full blackouts is working? We believe that for national utilities with huge demand-supply gap, possibility of social unrest and potential political backlash (for example, road-blockades and tire burnings [7]) provides an impetus for government to explore alternate solutions. Aashiyana's penetration enables flexible and fine-grained load-shedding policies that will reduce the underload wastage from the current strategy of coarse-grained, group-level shutoff while increasing social comfort within the *same supply-side constraints*; a push for such schemes will thus come top-down for socio-political reasons. A bottom-up push will come as consumer penetration of Aashiyana homes increases, and people observe the increased comfort level of their neighbors.

A serendipitous benefit of our DLC mechanism would be to actually reduce the load on the grid by removing the need for battery backups, extensively used already in countries with a stressed grid. These backup solutions use inefficient battery storage to transition into a single "low-power" state, but have shown to exacerbate the supply-demand gap that leads to greater penetration of battery backup and even greater stress — a death spiral for these grids [8], [9]. Our power-control system will provide an exact substitute for these backups, but with no inefficiency since it gets its (lower) power **directly from the grid**.

The contribution of this paper are the following. We present the design and implementation of *Aashiyana*: a novel home-level DLC system that can retrofit into wiring system to enforce different power-states for a home, with the set of

978-1-4673-8289-2/15/$31.00 ©2015 IEEE

appliances allowed in each state defined by individual home-owner (Section III). This provides a first, to the best of our knowledge, practical DLC mechanism for the dynamics of load-shedding of a highly-stressed grid. We propose two different DLC strategies that can be implemented now over a smart-grid and leverage any penetration level of Aashiyana to improve the social utility without requiring increased supply (Section V) Finally, we build a custom simulator to model a stressed grid and evaluate our DLC algorithms using the flexibility offered by Aashiyana to show that for 90% penetration we can decrease the fraction of homes with *no power* by $> 80\%$, without significantly decreasing (in some case actually increasing) the fraction of homes with full-power (Section VI).

II. Background and Related Work

In this section we first provide an overview of implementing load-shedding schedules, and then provide delta for our work to existing literature.

A. Background to Load-shedding in Pakistan

In Pakistan, power sector is primarily state owned, with some private sector responsible for generation. These generation units (GENCOs) work with the National Power Construction Corporation (NPCC), a central body which monitors the power grid, to asses the overall demand and operate power plants accordingly. NPCC, based on a survey done during a time with no supply shortage, decides on the allocated split of power generation across the ten major (state owned) distribution companies (DISCOs).

Pakistan has become, due to socio-political reasons, a country which faces a nearly year round supply-demand imbalance situation [4]. Whenever demand exceeds supply, segments of consumers (forming zones) are shutdown to match supply. The task of enforcing blackouts is accomplished primarily by DISCOs and in extreme cases, the NPCC. NPCC uses a generation estimate for each day, based on their dispatch requests, to allocate each DISCO a certain budget (based on a ratio previously determined) which they have to enforce regardless of the actual demand.

DISCOs have an *estimate* of consumption from their allocated area. If this estimate is greater than their budget, they initiate a load-shedding schedule across collection of feeders clustered into zones. They implement a time-disjoint blackout schedule across these groups, and increase the number of hours until they meet their quota. While DISCOs can under/over-use from their quota, the over all grid should remain stable. Otherwise, if instability due to overuse is neigh, the NPCC has the nuclear option of shutting down the 132kV line to the offending DISCO(s) and restore balance to the grid. In the rest of the paper, to simplify the analysis, we consider the case of a DISCO with a particular quota as representative of this more complex load-shedding strategy.

B. Related Work

The vast majority of research in demand response has considered the issue to shifting peak demand to allow for flattening such peaks. There is, to the best of our knowledge, no work that evaluates how such algorithms will work when applied to a grid where there is a continuous demand-supply gap. Our work proposes a fine-granularity DLC mechanism that is practical and leverages the conditioning of consumers to full blackouts in countries with highly-stressed grids. We divide our related work which focus on two different areas: grid-assisted DLC mechanisms and home-consumption changing systems.

Utilities have, for a long period of time, experimented with DLC mechanisms to share peak loads [5], [6]. These systems in the US and Canada, respectively, give consumers rebates for installing equipment to control a specific high power device (like A/C or heating units), for few hours a day. This granularity of control and load-reduction, even with high penetration, can not meet a *sustained* supply-demand gap in *giga-watts* that is typical for the highly stressed grids of countries our work targets. Perhaps the most similar work to ours is [3], [10]. Keshav and Rosenberg [3] propose a smart-grid where consumers contend to turn on appliances (proactive) or close the last appliance turned on that induces instability (reactive). iDR [10] proposes a theoretical framework that generates a signal for DR with the appropriate amount of reduction required from a consumer, such that the overall utility is maximized. Both these work do not have a practical system designed to implement their DR mechanisms, and specifically do not consider the large and continuous demand-supply gap of highly stressed grids.

The second area of research focuses on designing appliance or home level power reduction — while considering the comfort and ease of home owners [11], [12], [13], [14]. Both Yupik [11] and n-Plug [12] propose adding smart plugs to a few (deferable) appliances, whose usages patterns are monitored to present appropriate slack when a grid-stress event (demand peak) is indicated, using prices in Yupik and frequency in n-Plug. Smartcap [14] uses programmable switches or smart-appliances to control background appliances using a least-slack-first algorithm to flatten any peak. Srikantha et al. [13] evaluate how peaks can be flattened if the elastic component is allowed to be programmatically controlled. All these work seek to flatten load such that user prefernces are minimally affected; they fail to leverage (as they do not consider the problem domain) the experience of consumers in highly stressed grid that consequently face full blackouts for > 12hrs a day [4]. Furthermore, these are all primarily peak-shifting algorithms; for a highly-stressed grid there is *a permanent peak*.

To summarize, to the best of our knowledge, no work has yet considered a) the practical concerns of cost and usability for large scale DLC in homes and b) leveraged the readiness of consumers in highly-stressed grids to several levels of fixed power budgets.

III. Aashiyana: A practical DLC system

We propose to explore flexible and fine-grained load-shedding strategies by *enforcing multiple power states inside a home*. To achieve such control in a practical setting given the economic conditions of the afflicted countries, the system implementing these power-state should have a low-costs around (\$250-\$300). This price range corresponds to the cost of a battery backup that people are already comfortable purchasing and that provides similar (to Aashiyana) functionality to their owners. This is a challenge since making individual devices

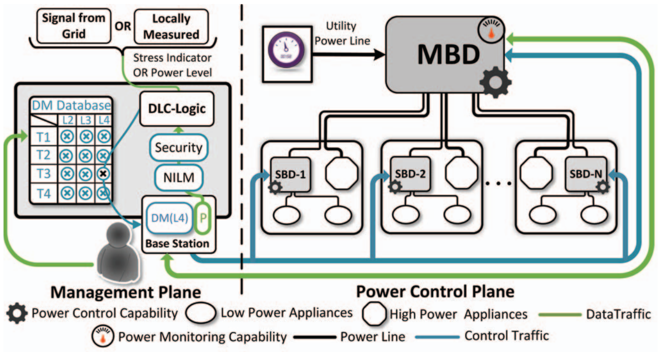

Fig. 1: Aashiyana Architecture

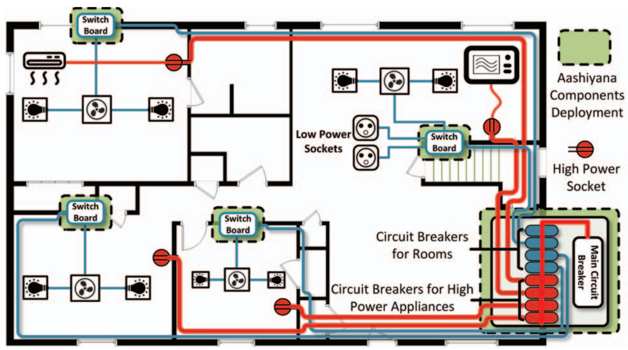

Fig. 2: Home Wiring Scheme

intelligent (using intelligent plugs or power strips [12], [11], [15]) is quite an expensive proposition. While ensuring a good-UI and configuration ease are also essential for this system, we do not explore them in this current manuscript.

Aashiyana is meant to be a user configurable system to enforce demand response that is acceptable to both utilities as well as consumers. We next describe the major design decisions for Ashiyana and then the architectural components of our demand-management solution.

A. Design Decisions for Aashiyana

A first design decision was to select the location and granularity of appliance control within a home within our cost constraint. For granularity we decided to restrict control to the level of switching devices on-or-off. Instrumenting any greater level of smartness would require significantly larger cost and configuration complexity.

We decide on locating our two control components at the main distribution board and at the level of switch boards installed at each home, in light of the traditional wiring structure for Pakistan (Figure 2). These locations provides us sufficient control as the high-power device sockets (separate in each room for A/C, Refrigerators) are accessible from the main distribution board, while individual sockets as well as fixed appliances like fan, and lights for each are accessible from a switch box. The choice of these location, and our custom board design, leads to a price point of around \$300 for a four room home (details in Section IV).

We next decide to restrict the power consumption of our homes to **five levels**. *Level 5 and Level 1* represent the current two modes of unrestricted power consumption and full disconnection, respectively. *Level 4-2* represent power consumption that is 75%, 50%, and 25% of full rated capacity. We restrict ourselves to just three configurable levels as a user will have to supply, for each level and possibly even different time-of-day or week, a matrix of control points that will be disconnected (which we call a *Disconnectivity Matrix (DM)*).

Finally, we also decide on using existing building automation and IoT frameworks, like [16], [17] to enable a ease of application development and a robust rendezvous mechanism. A custom system to provide redirection and web service accessibility requires significant development overhead, thus tilting us in favor of our final decision.

To better explain the architecture of Aashiyana, we split it into two planes (Figure 1): Management Plane and the Power Control Plane. While the home management plane implements the logic to trigger different power-levels, the power control plane enforces these states. We assume an external grid management plane that is (optionally) responsible for providing the demand-reduction signal to homes that implement Aashiyana. We detail the components in these planes next.

B. Power Control Plane

The power control plane controls the on/off state of appliances at the level of switch boards and main distribution board. Thus, we need two different control components, the Main Board Device (MBD), and the Switch Board Device (SBD).

The MBD is primarily responsible for controlling all heavy appliances from the main distribution box where, as per current wiring strategy, each high power sockets is connected via separate (higher rating) wires and circuit breakers (c.f. Figure 2). The MBD can similarly control power supply to every room, as a single wire goes to the switch box of each room, and distributed to individual appliances from there. A final purpose of the MBD is to monitor the power consumption at each room level to provide monitoring ability to prevent overuse at room level[1].

The SBD is located inside the switch board for each room, which terminates the direct line coming from the main distribution box. It is responsible, much like MBD, to control the wires distributing from this switch board. As shown in Figure 2, some wires go to hard-wired devices like fans and lights, while the rest go to individual sockets. We thus will have to *infer* devices connected to the sockets, while the hard-wired devices can be one-time configured.

Both these components communicate their data (power, state) to the home management plane through some IoT-based communication technology. We describe this plane next.

C. Management Plane

This plane consists of a DLC-logic module as well as a Base-station component that enables the communication using any IoT technology (802.15.4, Z-wave, power-line) used by the MBD and SBD.

We first save user preferences, in the form of a database of disconnectivity matrix (DM), for each power-state at this plane. We also collect consumption data from the control plane and provide this information to users for easy selection of

[1]NILM techniques can be used to increase the observation granularity

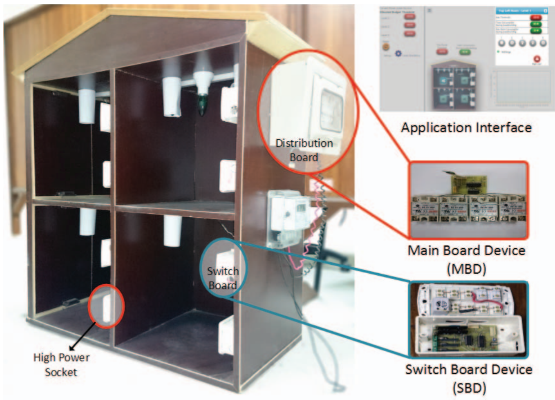

Fig. 3: Aashiyana Prototype Home

a DM that meets a given power budget. A second function is to appropriately respond to a grid-stress signal. Thus, any DLC algorithm is implemented through the selection of the appropriate power-state by the DLC-logic present within this plane. Once the power level is selected the appropriate DM is used to send commands to the control plane in order to switch off power to selected points. We believe that tamper-detection and prevention strategies will also be implemented at this central location.

The demand-reduction process initiated by the management plane requires an indication of grid-stress. While this detection can occur in a fully distributed manner at each home (by, for example, sensing frequency [12]), we expect the utilities to have some demand-response logic based on the current supply-demand gap. This logic is represented by an optional Grid Management Plane (not shown in Figure 1).

Using the Aashiyana system, the utilities now have the option to specifying five power-states (as described in Section III-A), instead of the current two states of uninhibited consumption or a full blackout. We *reiterate* that these different power states are *acceptable for consumers where black-outs* are a regular occurrence, but might not be equally palatable for consumers where such events are unthinkable.

IV. Aashiyana: Implementation and Evaluation

In this section we present our working prototype of an Aashiyana system. We replicate the wiring of a typical home by enlisting a professional electrician, over a prototype and scaled-down four room home (Figure 3). This home has a main distribution board with circuit breakers for each room and high power appliances. We use Current Cost's power meter to obtain the home-level energy consumption data.

A. Hardware Devices

The **Main Board Device (MBD)** located inside the distribution box of the home consists of three major components; a communication module, a processing unit, and controlling relays. We had wi-fi, PLC, and RF as candidates for communication technologies. Keeping in mind our cost constraints, we opted for a CC2500 RF module, having cost of \$3.95, for this purpose. We use Msp430 launchpad (\$2.80 each) for our processing needs. We chose Solid State Relay (SSR) due to its longer life and noiseless operation despite a slightly higher cost. Four high rating (25A @ 220V, \$7.94 each) SSRs are used to control the high power sockets.

The **Switch Board Device (SBD)** is similar to the MBD with a primary difference in number of relays and their ratings. The SBD has five outlets with three of them for fixed appliances (fans and lights) and two power sockets. We used 1A rating SSR (\$5.18) for the three (known rating) appliances and 5A ones (\$14.56) for power sockets with variable load.

The **Base Station** is responsible for implementing the control-decisions of the management plane and send them to the MBD and SBDs using the communication protocol over the CC2500 RF chip. This chip is coupled with a micro-controller, currently an mbed, but is being replaced with our msp430 option for cost savings.

B. Software Components

We use Microsoft Research's Lab-of-Things (LoT), an open source platform that enables seamless integration of home-automation systems into a coherent framework by requiring vendor specific drivers. It provides a cloud-based rendezvous mechanism and a relayering feature, allowing NAT traversal capability so that applications built are accessible from anywhere. We wrote a LoT application for Aashiyana system which implements the disconnectivity matrix when a decision is made in response to stress signals. Currently, we use a dedicated laptop as a home hub but target a low-cost tablet supporting .NET framework for around \$65.

We custom wrote the device drivers for the MBD and SBD. The device drivers allow the application to view each control point that can be mapped to a particular disconnectivity matrix. These device driver talk with the base station through a serial port, and the base station uses its communication stack to coordinate with the SBDs and MBD. Our current implementation thus, besides implementing different power states, also provides LoT roles to allow on/off control of devices in a home.

The communication stack on base-station first discovers all the devices in a home and converts a disconnectivity matrix into a command to switch off corresponding relays at each control device. This command is sent as a broadcast message, where the relay state is represented by bits in a byte, and all devices identified in the message payload individually. We require an acknowledgment from every device for reliability. The ack also contains the current state of the relays on that device allowing the LoT application to have an accurate view of the home-power state.

C. Practical Evaluation of Aashiyana

We now evaluate the practicality of our system by measuring its power-level switching delay, and power consumption.

The power-level switching delay has to be on the order of reaction delays that prevent grid instability. A power grid requires an instantaneous matching of supply with demand: regulating reserves[2], however, are typically employed to manage the time until load-following reserves can come online. A typical grid has around *five minutes* of this *stability time*, but can also reach up to an hour in some cases [18].

[2]Regulating Reserve is the capacity of generators to supply energy within an economic dispatch interval in response to the grid frequency variations.

The average delays for our system to respond to a stress signal is around 30 msec; even while including typical Internet latency of around 200msec our response time is more than *two-orders-of magnitude* faster that the grid stability time. We also evaluated the reliability of communication and achieved 94 % Packet Reception Rate (PRR).

Finally, we observed the power consumption of our prototype system which consisted of four rooms. The power cost of our Aashiyana prototype, if all devices are operating, will be an additional 5 watts and much less when a lower power-level is enforced. The detailed evaluations of the Aashiyana System can be found at a companion tech-report[19].

V. DLC ALGORITHMS

We propose two different algorithms, one central and the other distributed which enforce reduction at an hour-long granularity, and we ensure (unless impossible) that customer are not successively put into a *any* load-shedding state. We also limit Aashiyana homes to L2, leaving L1 (full shutoff) as a final resort, using an emergency signal. We only use additive backoff; any increase (multiplicative or additive) to meet extra supply made available due to reduction were deemed too prone to oscillations. With a grid trying to *meet* its demand, it thus appears that only decrease in demand, returning to full level at the end of an hour, is the appropriate response.

It is important to note here that utility companies, predominantly state owned in countries requiring load-shedding, will be motivated to decrease the discomfort (and thus potential for unrest) by promoting the deployment of Aashiyana-like systems. Consumers, however, will be motivated when they see their neighbors with the system installed having greater utility. While they can buy and install local backups (that have fixed and recurring cost), our system can be deployed at a similar cost but fewer power units consumed (no losses of a battery backup), for the same effective utility.

A. Distributed DLC Algorithm

In our distributed algorithm, the power throttling level is stochastically generated inside the DLC logic. However, due to a variable Aashiyana Penetration (AP), the grid utility is still involved to fully shut-off non-Aashiyana homes. The only external information required is the stress level on the grid, communicated either by the utility or by locally sensing supply frequency [12]. We define this *stress level* as $sl = \frac{D-S}{D}$; where S is the current supply while D is the actual (unfulfilled) demand. Intuitively this is the fraction of actual demand not currently being met, and thus the fractional amount to be shed.

We propose an iterative algorithm, where in the first iteration an Aashiyana homes run a stochastic algorithm (Algorithm 1) to determine their new power state. This state is selected as a function of the sl and a *Distribution Profile (DP)* characterizing the relative percentage of homes that we would want to stay in L4, L3, and L2. We then completely shutoff non-Aashiyana homes, if this load reduction does match supply to demand. We start the second iteration[3] only if the demand is still not met, and repeat this algorithm. This demand reduction continues until demand meets supply.

We make homes already in back-off to use their original (iteration 1) stress level, reducing the probability that all consumer are put in a back-off state. All non-backoff homes (in any iteration) use the current sl value that is lower due to demand reduction in earlier iterations.

Algorithm 1 Distributed DLC algorithm

Require: sl, DP

$\forall h_i \in \mathbb{H}_a$ ▷ $\mathbb{H}_a$ set of homes with Aashiyana
if $(LS_{lh} == true\ AND\ Emergency == false)$ **then**
  $exit$ ▷ LS_{lh} Loadshedding in Last Hour
$r_i \leftarrow rand(1, 100)$
if $(DLCDone == false)$ **then**
  $sl_i \leftarrow sl, sl_{init.i} \leftarrow sl$
  if $(r_i < sl_i)$ **then**
    $SetToDifferentPowerLevels(h_i, r_i, DP)$
    $DLCDone == true$
else$(sl_i \leftarrow sl_{init.i})$
  if $(((r_i < sl_i)\ OR\ Emergencey)\ AND\ CurrentLevel \neq L2)$ **then**
    $SetToLowerLevels(h_i)$
  At the end of hour $DLCDone == false$

B. Centralized DLC Algorithm

Our centralized algorithm (Algorithm 2) assumes that the utilities have a full view of the current consumption of each home. The utilities already have feeder-level groups established in which they cyclically implement their current load-shedding. In our scheme, they now refine this process by picking a group and computing the savings by shutting-off all non-Aashiyana homes. The utility then computes the energy savings by reducing the consumption level of Aashiyana homes, given the the demand-supply gap (Δgap) is still positive, They do so by considering highest consumption homes first, to any level (chosen equally-likely) below their current consumption[4]. They choose another home only if the $\Delta gap > 0$, and if cycling through all Aashiyana home still doesn't satisfy demand, the next group is selected and the process continues until demand is met. Once this decision is made centrally, the control decisions are then communicated directly (and at once) to every home: Aashiyana homes are communicated the selected level that bypasses the DLC logic shown in Figure 1, while other homes are completely shut-off by their smart-meters.

This process is repeated hourly, with already reduced groups chosen only if all others groups have been made to back-off first.

Algorithm 2 Centralized DLC algorithm

$G_i \leftarrow \mathbb{G}\ (RoundRobinSelection)$
for $(\forall\ h_i\ \in\ \mathbb{G}_i\ AND\ \forall\ h_i\ \in\ \mathbb{H}_{na})$ **do** ▷ $\mathbb{H}_{na}$ set of non-Aashiyana homes
  $SetToPowerLevelOne(h_i)$
if $(\Delta gap\ >\ 0)$ **then**
  for $(\forall\ h_i\ \in\ \mathbb{G}_i\ AND\ \forall\ h_i\ \in\ \mathbb{H}_a)$ **do**
    $SetToLowerLevels(h_i)$
    if $(\Delta gap\ \leq\ 0)$ **then**
      $Stop$
if $(\Delta gap\ >\ 0)$ **then**
  $goto\ step\ 1$

VI. EVALUATING AASHIYANA-ENABLED DLC

We now evaluate the benefits of our proposed large-scale DLC algorithms using Aashiyana, employing a custom event-driven simulator implemented in C++. We model the

[3]each iteration is set to be 1 sec, sufficient for an Aashiyana response and grid stability.

[4]Thus if consumption of a home is above 75% of its meter rating, either of L4, L3, or L2 can selected.

individual appliance level data as a distribution by using the appliance consumption information from the UK-DALE [20] and REDD [21] dataset. We apply the Kernel Density Estimation (KDE) technique on this power consumption data to obtain the cumulative distribution function (CDF) estimate, modeling a devices consumption in a stochastic manner. We then use inverse transform sampling to obtain a stochastic power consumption value for every device, averaged over a full hour.

We assume three class of homes (A, B and C) with 7, 10, and 13 set of appliances respectively. We implement the five power levels in the simulator as a function of different maximum rating for each class — class A has 500W, class B has 750W, and class C has 1000W. Power level 1 is full shutoff (no appliance) while power level 5 allows all appliances of that home type to be turned on. We setup a disconnectivity matrix for each class of home for every intermediate level to attain 75%, 50%, and 25% consumption of their maximum rating.

The power consumption of different appliances in a home vary every hour based on their observed CDF. A Central body agent checks the supply-demand gap every second and generates a stress signal to reduce the gap based on the selected algorithm. On reception of the DLC signal, the DLC logic implements the selected algorithm. Our zone definition consists of ten feeders in the grid. There are a total of 40,000 home agents distributed uniformly across these zones. We also plan to open source our simulator to build and run large-scale simulation studies for smart-grid power systems.

We assume, for our evaluations, lossless transmission and distribution in the grid. We assume the utility have a list of all homes, whether they have Aashiyana installed or not. We also assume instantaneous reception of DLC signal and that its response is implemented within a second.

A. Traditional Load-shedding and Under-load Wastage

We first show the running of our simulator with traditional load-shedding strategy. Figure 4 shows the scenario where we have a 20% demand supply gap. We can see that there are several locations where the demand satisfied by implementing cyclic load-shedding, leads to something below the available generation capacity. In these scenarios, the governor control on generator side leads to *less power being generated than the capacity online*. This is what we call under-load wastage (ULW), and it not only produces energy less efficiently across our ensemble of generators, but also allows fewer people to have power than is possible. The Figure 4 (inset) shows that (for hour 18) our algorithms can be efficient by reducing ULW from 1.12MW to just 0.2MW. This reclaimed energy allows not only a decrease in the number of home with no power (L1) by 80% but also *an increase* by 6% in the fraction of homes with full power (L5).

B. Distributed Algorithm

Having established that for a static supply case, our distributed algorithm allows for increased social comfort by significantly reducing the number of homes with no-power (L1) for a slight decrease in homes with full-power (L5), we now consider a case where the supply demand gap is a consistent percentage of actual demand. For this purpose, we

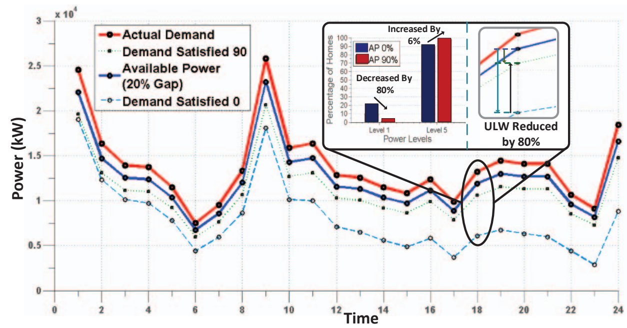

Fig. 4: Observing under-load wastage (ULW) reduced by using Aashiyana-based DLC algorithms

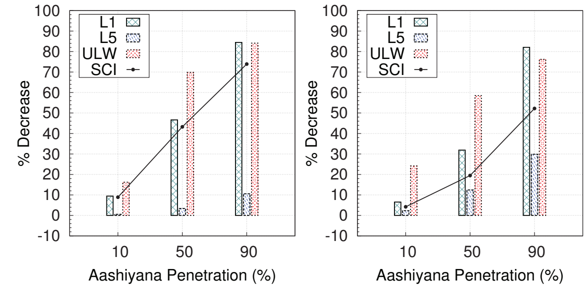

(a) 20% gap (b) 40% gap

Fig. 5: Distributed Algorithm: Change (from no AP) in distribution of home-levels with varying demand-supply gaps.

vary this gap as 10-40%, and observe the percentage decrease, over a 24hr period and compared to no Aashiyana penetration (AP) (i.e. traditional load-shedding), in homes that are in L1 and the concomitant decrease for L5 at different AP levels.

Figure 5 show the result of averaging 10 simulation runs for each experimental setting (% gap, AP level). As is quite evident, the fractional decrease of homes in L1 *always* greater than (by more than 100%) the corresponding decrease in L5. This difference, corresponding to increase in social comfort, is understandably greatest at the highest AP level with the **social comfort index** (SCI[5]) ≈ 80 percentage points for 90% AP, thus clearly indicating the benefit of wide-scale adoption. We notice that as the demand-gap increases, the improvement in SCI decreases. This is so since large gap will necessarily demand a lots of homes to be load-shed, thus necessitating more homes to go below L5.

C. Centralized algorithm

Our centralized algorithm has a holistic view of the energy consumption status. We therefore expect that, while the trends will be similar to those for our distributed algorithm, it will be more efficient by reducing the amount of ULW and should thus help increase (SCI).

Figure 6 shows the results for our centralized algorithm averaged over 10 simulation runs. We observe that most improvements are when we have the highest Aashiyana penetration, and that it decreases with increasing gap. Further, we can confirm that with a more informed strategy, the centralized approach obtains greater ULW savings, thus we have 99% less wastage (0.019 MW vs 3.3 MW) for 90% AP at 20% gap.

[5]SCI is defined as the magnitude of difference between the fractional decrease in L1 and fractional decrease in L5.

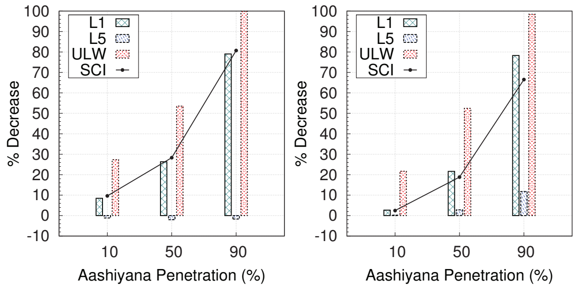

(a) 20% gap (b) 40% gap

Fig. 6: Centralized Algorithm: Change (from no AP) in distribution of home-levels with varying demand-supply gaps.

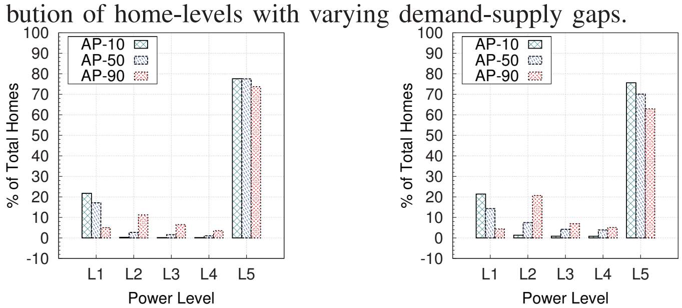

(a) Centralized Algorithm (b) Distributed Algorithm

Fig. 7: % of homes in all levels for a 30% Gap.

Moreover, it is interesting to observe that in certain cases, the fraction of homes in L5 actually increases (negative decrease). Such a scenario is purely win-win since now, for the *same supply-demand gap*, we have not only decreased the fraction of homes in L1, but were also able to *increase* the fraction of homes with unrestricted power!

D. Discussion

We note that between two algorithms, the centralized approach is consistently better at reducing the ULW, and provides better SCI at higher AP. However, the SCI values for the distributed case are slightly better at lower penetration. We believe this is so because SCI show distribution for only two possible levels. Figure 7 shows this distribution for homes for the same gap, but after different algorithms have run. We observe that the distributed algorithm results in distributing savings to lift nodes from no-power (L1) into a low-power state. The centralized algorithm instead, results in more homes with full-power and fewer are taken out from L1.

To our mind, since distributed mechanisms are robust to individual failures, and their low complexity allows quick implementation, exploring a distributed approach to manage stress is a promising direction. It is also worth exploring if increasing the number of levels is beneficial; we do caution that more levels decrease the user-friendliness of the system since the user has to think about and configure a disconnectivity matrix for each state.

VII. Conclusions

We present here a novel and practical DLC system, Aashiyana, that enables several different low-power states for homes within the context of highly stressed grids. We design and implement this with practical incentives for the utilities (decreasing social unrest) as well as consumers (low-cost, lower hours with no-power, greater utility), *all without* having to increase the supply side equation. We propose two types of algorithms that utilize this ability of guaranteed budget reduction at different levels that allow for more efficient reduction in gap with reduced amount of underload wastage. We show that, compared to current load-shedding strategy, for the same supply-demand gap, we can reduce homes with *no* power by $> 80\%$ while not significantly impacting the fraction of homes with full power.

References

[1] M. A. Piette, G. Ghatikar, S. Kiliccote, D. S. Watson, E. Koch, and D. Hennage, "Design and operation of an open, interoperable automated demand response infrastructure for commercial buildings," *Journal of Computing Science and Information Engineering*, vol. 9, no. 2, 2009.

[2] A. H. M. Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments.," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 120–133, 2010.

[3] S.Keshav and C.Rosenberg, "Direct adaptive control of electricity demand," tech. rep., University of Waterloo, 2010.

[4] I. N. Kessides, "Chaos in power: Pakistan's electricity crisis," *Energy Policy*, vol. 55, 2013. Special section: Long Run Transitions to Sustainable Economic Structures in the European Union and Beyond.

[5] S. C. Edison, "Summer discount plan." web site https://www.sce.com/wps/portal/home/residential/rebates-savings/summer-discount-plan/.

[6] EnerSource, "Peak saver." web site http://www.enersource.com/energy-savings-tips/Pages/peak_saver.aspx.

[7] "Dawn news - violent protests in lahore, multan against loadshedding." website http://www.dawn.com/news/1121213.

[8] D. P. Seetharam, A. Agrawal, T. Ganu, J. Hazra, V. Rajaraman, and R. Kunnath, "Hidden costs of power cuts and battery backups," ACM e-Energy '13.

[9] Z. Sharani, K. Qayyum, N. Bashir, and A. Syed, "Softups: Eliminating the need and cost of battery backups in the developing world," ACM e-Energy '14.

[10] V. Chandan, T. Ganu, T. K. Wijaya, M. Minou, G. Stamoulis, G. Thanos, and D. P. Seetharam, "idr: Consumer and grid friendly demand response system," ACM e-Energy '14.

[11] T. Bapat, N. Sengupta, S. K. Ghai, V. Arya, Y. B. Shrinivasan, and D. Seetharam, "User-sensitive scheduling of home appliances," in *ACM SIGCOMM Workshop on Green Networking, 2011*.

[12] T. Ganu, J. Hazra, D. P. Seetharam, S. A. Husain, V. Arya, L. C. D. Silva, R. Kunnath, and S. Kalyanaraman, "nplug: a smart plug for alleviating peak loads," in *ACM e-Energy'12*.

[13] P. Srikantha, C. Rosenberg, and S. Keshav, "An analysis of peak demand reductions due to elasticity of domestic appliances," ACM e-Energy '12.

[14] S. K. Barker, A. K. Mishra, D. E. Irwin, P. J. Shenoy, and J. R. Albrecht, "Smartcap: Flattening peak electricity demand in smart homes.," in *IEEE PerCom, 12*.

[15] "Aeotec smart strip." website http://aeotec.com/z-wave-power-strip.

[16] C. Dixon, R. Mahajan, S. Agarwal, A. J. Brush, B. Lee, S. Saroiu, and P. Bahl, "An operating system for the home," USENIX NSDI'12.

[17] Citris, "Openbas: Open-source building automation systems for smaller commercial buildings." website http://goo.gl/7XV3fb.

[18] M. M. Erik Ela and B. Kirby, "Operating reserves and variable generation," electric power-plants, load management, electric power production, Colorado, aug 2011.

[19] N. Bashir, Z. Sharani, K. Qayyum, and A. Syed, "Enabling Practical Demand Response in Highly-Stressed Grids using Aashiyana," *ArXiv e-prints*, Apr. 2015.

[20] J. Kelly and W. Knottenbelt, "UK-DALE: A dataset recording UK Domestic Appliance-Level Electricity demand and whole-house demand," *ArXiv e-prints*, Apr. 2014.

[21] J. Z. Kolter and M. J. Johnson, "Redd: A public data set for energy disaggregation research," in *SIGKDD, 11*.