Managing Contingencies in Smart Grids via the Internet of Things

Stefano Ciavarella, Jhi-Young Joo, *Member, IEEE*, and Simone Silvestri, *Member, IEEE*

***Abstract*—This paper proposes a framework for contingency management using smart loads, which are realized through the emerging paradigm of the Internet of things. The framework involves the system operator, the load serving entities (LSEs), and the end-users with smart home management systems that automatically control adjustable loads. The system operator uses an efficient linear equation solver to quickly calculate the load curtailment needed at each bus to relieve congested lines after a contingency. Given this curtailment request, an LSE calculates a power allowance for each of its end-use customers to maximize the aggregate user utility. This large-scale NP-hard problem is approximated to a convex optimization for efficient computation. A smart home management system determines the appliances allowed to be used in order to maximize the user's utility within the power allowance given by the LSE. Since the user's utility depends on the near-future usage of the appliances, the framework provides the Welch-based reactive appliance prediction (WRAP) algorithm to predict the user behavior and maximize utility. The proposed framework is validated using the New England 39-bus test system. The results show that power system components at risk can be quickly alleviated by adjusting a large number of small smart loads. Additionally, WRAP accurately predicts the users' future behavior, minimizing the impact on the aggregate users' utility.**

***Index Terms*—Smart grid, Internet of things, contingency management, energy management.**

I. Introduction

CONTINGENCIES resulting in cascading failures are critical issues in power systems operation. These events occur at a low probability but can evolve into large-scale blackouts. A report by the U.S. Executive Office of the President estimates that between 2003 and 2012, 679 large scale power outages occurred in the U.S., each affecting at least 50,000 customers [1]–[3]. The economical impact of cascading failures is also significant, as the costs range from 18 to 33 billion dollars per year [1].

Manuscript received September 13, 2015; revised December 18, 2015; accepted February 9, 2016. Date of publication March 3, 2016; date of current version June 17, 2016. This work was supported in part by the Defense Threat Reduction Agency under Grant HDTRA1-10-1-0085 and in part by the National Science Foundation under Grant CNS-1545037. Paper no. TSG-01125-2015.

S. Ciavarella is with the Department of Computer Science, Sapienza University of Rome, Rome 00100, Italy (e-mail: ciavarella@di.uniroma1.it).

J.-Y. Joo is with the Department of Electrical and Computer Engineering, Missouri University of Science and Technology, Rolla, MO 65409 USA (e-mail: joojh@mst.edu).

S. Silvestri is with the Department of Computer Science, Missouri University of Science and Technology, Rolla, MO 65409 USA (e-mail: silvestris@mst.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSG.2016.2529579

The cause of such cascading failures is often the relative primitiveness of *contingency management*. On one hand, management often relies on the judgments of human operators, who decide on possible countermeasures based on their experience. On the other hand, such operators can only see the high-voltage transmission levels of the grid, with little outlook on the adjustable loads at the end-users' level.

This work particularly examines contingency cases where one or more of the system components are unexpectedly down but the system balance is still achieved due to the strict reliability criteria. Even when the system restores its balance after a contingency, however, there is a risk of cascading failures as with the 2003 North American blackout case [1]. In fact, other lines can approach their maximum limits, and eventually drive the system over the critical point beyond a stable state. Therefore, precautionary measures to avoid cascading failures are necessary.

This work proposes a novel framework to alleviate this type of risks by adjusting a large number of end-use loads. The assumption is that, for such emergency cases, curtailing some non-critical loads to prevent cascading failures yields greater aggregate utility than leaving the lights on and having cascading failures later. Therefore, in this work, all controlled loads, which exclude critical loads such as life-support devices, are assumed to be curtailable within a short time period, e.g., 1 hour. The end-user loads are controlled by smart devices, realized through the emerging paradigm of the Internet of Things [4], [5]. According to this paradigm, smart devices are equipped with communication, computation and storage capabilities, and they are connected to a *smart home management system* through wireless access points. Each smart device controls an appliance such as a space heater, an air conditioner, a refrigerator, etc.

The framework comprehensively involves the system operator, the load serving entities (LSEs), and the end-users' smart systems. The system operator prevents cascading failures by completing the following tasks after achieving the stable, but still potentially risky, state of the power system: 1) identify the components that violate the predetermined reliability criteria, and if there is any, 2) calculate the load adjustment at different locations (i.e., buses) to alleviate the additional stress in those particular components. Since it is assumed that the system is in balance, the total amount of load curtailment at each bus is efficiently calculated by a novel approach using the linearized network equation.

When an LSE is notified of a load curtailment amount at its load bus, it solves a mixed integer linear optimization

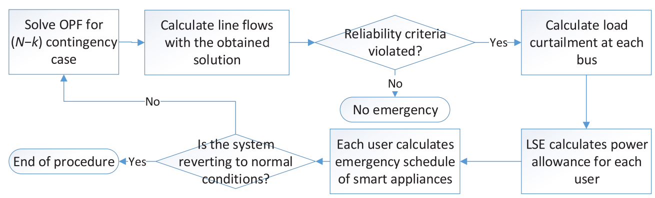

Fig. 1. The flowchart of the contingency management framework.

problem that maximizes the aggregate utility, i.e., the sum of its end-users' *utility*. In this work *utility* is defined as a quantifiable measure of user satisfaction from using a certain appliance. To improve scalability of the LSE's problem, an approximated convex problem of the mixed integer optimization is solved, using an efficient heuristic based on regression techniques.

The solution to the LSE's problem is the individual load curtailments of the LSE's users. The smart home management system then calculates an *emergency schedule*, which defines the best set of appliances that the user is allowed to use. This schedule minimizes the impact of the curtailment on the user's habits, while satisfying the power allowance requested by the LSE. The calculation of the emergency schedule requires the knowledge of the future user interaction with appliances. To predict this interaction, the framework uses the WRAP (Welch-based Reactive Appliance Prediction) algorithm. WRAP uses smart devices to monitor the user habits and to predict the appliance usage following the contingency.

After load curtailments, the system operator evaluates the system condition, and if the system is not reverting to the normal condition, the procedure is repeated from the beginning. The flowchart of the framework is shown in Fig. 1.

Extensive simulations are performed on the IEEE New England 39-bus test system, using real power consumption datasets, to validate the benefits of this framework. The results show that the proposed method is effective in calculating the load curtailments needed after a contingency, ensuring that the lines operate within their capacity margins. Additionally, the WRAP algorithm achieves highly accurate predictions of the appliance usage. Finally, the regression-based heuristic performed by the LSEs closely matches the results achieved by solving the original mixed integer programming problem. As a result, the proposed framework is effective in keeping the system stable during contingencies, preventing cascading failures while maximizing the aggregate user utility.

The main contributions of this paper are the following:

- A comprehensive framework for contingency management using smart appliances based on the paradigm of the Internet of Things;
- A novel and efficient method to enable system operators to calculate the load curtailment needed in order to keep the system in a safe state after a contingency;
- An efficient heuristic to distribute such curtailments across end-users;
- The Welch-based Reactive Appliance Prediction (WRAP) algorithm to predict utilization of each appliance by a user;
- Extensive simulations on realistic system settings to validate the proposed approach in managing contingencies while maximizing users' utility.

II. Related Works

There have been efforts in managing contingencies with adjustable demand. Reference [6] considers dispatch of load curtailment at the system-level operation alongside with generation, based on the bids submitted by the customers. Presumably these customers represent the load serving entities, but the dynamics or attributes of the individual load models are missing. In [7], demand response is used in place of spinning reserves to restore the frequency post contingencies. Reference [8] uses demand response for efficient use of transformers during contingencies, and reference [9] corrects voltages adjusting post contingency demands. The last three works are novel in terms of using demand to manage contingencies, but have different purposes from this work, where we alleviate congestions post contingencies.

There are also many works on predicting users' power consumption by appliance. Reference [10] proposes an algorithm to identify the individual consumptions of residential appliances. Reference [10] predicts electrical heating and cooling power by large groups of customers. In reference [11], user discomfort is minimized with a Q-learning algorithm, which computes the optimal set of appliances to switch off during the system peaks. However, this approach assumes that the importance of an appliance to a user is known in advance. This work takes a step further from the literature on users' demand and utility, and proposes an effective algorithm that exploits smart appliances to *predict* and maximize the users' utility using historical data.

Another novel aspect of the proposed work is in modeling the system comprehensively from the power transmission grid all throughout the end-users equipped with smart home management systems. The objective is not only to manage contingencies at the system level, which existing literature has studied extensively, but to maximize the aggregate utility of all users, when a certain amount of capacity is requested from the system operator.

The paper is organized as follows. Section III describes the problem of the system operator, while Sections IV and V address the users' and LSE's problems, respectively. The simulation results are presented in Section VI, and Section VII concludes the paper.

III. The Problem of a System Operator

The objective of the system operator in general is to keep the system reliable at the least cost. After one or more lines failed, the power flows in other lines can approach their limits. Therefore, in this work the system operator's objective is to find the load curtailment that can alleviate the line flows of these additional lines at risk to prevent cascading failures.

The relationship between the active power flows in the lines and the active power injected into each bus can be linearized with a power transfer distribution factor (PTDF) matrix H.

$HP = F$ where P is a vector of active power injection at each node except the slack bus, and F is a vector of active power flow in each line [12], [13]. Therefore, demand curtailment ΔP_D, defined as a vector with the adjustable demand buses as its components, that yields the line power flow difference ΔF can be calculated by solving

$$H_C C_D \Delta P_D = \Delta F_C \tag{1}$$

where ΔF_C is the line adjustment vector with only the congested lines selected from F, whose length is equal to N_L, number of congested lines. C_D is a bus-demand connection matrix with a dimension (the number of total buses in the system N_B)-by-(the number of demand buses N_D), whose element is 1 when the bus (row) is a demand bus (column) and 0 otherwise. H_C is extracted from H with only the rows of the congested lines, thus N_L-by-N_B. Usually since $N_D > N_L$, the solution to this equation can be obtained as $\Delta P_D = H^+ \Delta F$ where H^+ is the Moore-Penrose pseudoinverse of H.

This calculation gives a solution with the minimal norm among the many solutions of (1). Therefore, the sum of the solution, or the total system load adjustments, can be negative, resulting in the need for more generation to balance the supply and demand. In order to avoid this, the solution is sought so that the load adjustments in the system sums up to zero, i.e.,

$$\sum_{n=1}^{N_D} \Delta P_D(n) = 0. \tag{2}$$

Concatenating (1) and (2) yields an augmented power flow equation

$$\begin{bmatrix} H \\ \mathbf{1}_{N_D}^T \end{bmatrix} \Delta P_D = \begin{bmatrix} \Delta F \\ 0 \end{bmatrix}, \text{ or } \tilde{H} \Delta P_D = \Delta \tilde{F} \tag{3}$$

where $\mathbf{1}_n$ denotes an n-length column vector that has 1 as all its elements. The solution can be obtained in the same way by solving $\Delta P_D = \tilde{H}^+ \Delta \tilde{F}$.

The solution to (3) can include negative load adjustments, which means that some load buses need to *increase* their consumption. If the system operator decides that this is unreasonable or if it is technically infeasible, then the system operator can take the nonnegative solution $\Delta P_D|_+$ where $\Delta P_D(l)|_+ = \max[0, \Delta P_D(l)]$ for all l, and resolve the power flows with this adjusted solution. It should be noted that the power flows with this solution may result in power flow adjustment smaller than the target ΔF_C. However, since the relationship between the load curtailment and the power flows is linear, the load curtailment solution can be simply scaled by the factor of the desirable power flow adjustment.

IV. The Problem of a User

This section discusses the problem solved by the smart home energy management system of a user. As a contingency occurs, the system operator sends each LSE the load curtailment $\Delta P_D(l)$ for each bus l. Let $M^l_{\max}$ be the power allowance that the LSE is allotted for Bus l. Then the LSE distributes $M^l_{\max}$ across its users, calculating the individual power allowance $M^l_1, \ldots, M^l_{N_l}$ for each of its N_l users so that $\sum M^l_i = M^l_{\max}$. Finally, the users' smart home management systems schedule the smart appliances to be used. Since most variables and parameters for an LSE's problem are defined in the users' problem, to improve readability, the problem solved by the users is first presented, followed by how an LSE distributes $M^l_{\max}$ in Section V.

In the proposed framework, a smart home has n smart devices $d_1, \ldots, d_n$. For each appliance, the framework defines a time-dependent *importance factor*, according to the user's usage preference and patterns, which may vary depending on the time of day, the season of year, and the user's habits. The goal is to use smart devices to learn the importance factors of the appliances for each user u during normal system conditions. To calculate such factors, the framework considers *time slots* $\tau_1, \tau_2, \ldots$ of arbitrary length, set as one hour in this work. Let $\lambda^u_{i,j} \in [0, 1]$ be the fraction of time that user u uses the appliance i in time slot j. The importance factor $\gamma^u_{i,j}$ of appliance i is defined as:

$$\gamma^u_{i,j} = \frac{\lambda^u_{i,j}}{\sum_{h=1}^{n} \lambda^u_{h,j}}. \tag{4}$$

Therefore, $\gamma^u_{i,j}$ represents the relative usage time of appliance i with respect to the other appliances during time slot j.

The approach assumes that the importance factor $\gamma^u_{i,j}$ measures the contribution of appliance i to the utility of user u during the time slot τ_j. Therefore, given a set of appliances A, the utility that results from using such appliances in time slot j for user u is $\sum_{d_i \in A} \gamma^u_{i,j}$. As soon as the LSE informs the user u's smart home management system of the new power allowance M_u, an *emergency schedule* that determines the set of appliances that can be used is calculated so that it maximizes the *user utility*.

A. Optimal Emergency Schedule

Using the importance factors defined previously, the framework makes use of the following optimization problem to determine the emergency schedule of a user u. The problem is solved by the smart home management system, which receives a power allowance M_u from the LSE (the calculation of M_u is described in Section V). The description considers a load curtailment during time slot τ. Let $x_i \in \{0, 1\}$ be a decision variable, where $x_i = 1$ if appliance i is allowed to be used during the emergency schedule, and $x_i = 0$ otherwise. Additionally, let $e_1, \ldots, e_n$ be the maximum power rating of the appliances. For simplicity, the following discussion assumes the state of an appliance is either ON or OFF. Then the optimal scheduling problem is:

$$\underset{x_i}{\text{maximize}} \quad \sum_{i=1}^{n} \gamma^u_{i,\tau} x_i \tag{5}$$

$$\text{subject to} \quad \sum_{i=1}^{n} x_i e_i \le M_u \tag{6}$$

where the values of $\gamma^u_{i,\tau}$ are calculated according to (4), and are estimated as described in the next subsection.

This optimization problem is clearly NP-hard. However, since smart homes generally have a limited number of appliances, the problem can be solved in a short time optimally, or through standard heuristics [14]. The resulting schedule is enforced by the smart home management system, which restricts the use to only the selected appliances.

B. Learning Algorithm for Importance Factors

In order to solve (5), the values of time-dependent importance factors $\gamma_{i,\tau}^{u}$'s need to be known. However, since they represent the user's future behavior, they can only be predicted. This section describes the Welch-based Reactive Prediction (WRAP) algorithm, executed by the smart home management system. WRAP makes use of a statistical change detection mechanism based on the Welch's t-test [15] to predict the importance factors.

WRAP is based on the assumption that the fraction of time $\lambda_{i,j}^{u}$, during which user u uses appliance i in time slot τ_j, is distributed over multiple days as a Gaussian random variable. The means and standard deviations of $\lambda_{i,j}^{u}$'s may change over time. The results in Section VI prove that this assumption enables accurate estimation of the importance factors and maximization of user utility.

The smart home management system keeps track of the historical usage of each appliance. In particular, for each time slot τ_j the system calculates the *historical mean* $\mu_{i,j}^{H}$ and the *historical variance* $\sigma_{i,j}^{H}$. At the end of each time slot, these historical values are updated with the newly observed values.

1) Short-Term Change Detection: WRAP adopts a change detection mechanism based on the Welch's t-test [15], to achieve high accuracy and reactivity, i.e., ability to react to changes in the usage pattern. In particular, the idea of detecting short-term changes is to verify if the most recent utilization of an appliance is *unusual* with respect to the historical data.

Consider an emergency period occurring at time period τ_j, hence the importance factors need to be predicted for τ_j. Let $< \mu_{i,j}^{H}, \sigma_{i,j}^{H} >$ be the historical distribution for appliance i during τ_j, and $< \mu_{i}^{W_S}, \sigma_{i}^{W_S} >$ be the distribution over only a recent time window W_S, e.g., the last 60 minutes. In order to detect if there is a change in user behavior, WRAP determines whether the distribution $< \mu_{i,j}^{H}, \sigma_{i,j}^{H} >$ and $< \mu_{i}^{W_S}, \sigma_{i}^{W_S} >$ belong to the same population (*null hypothesis*) or not (*alternative hypothesis*).

The Welch's t-test defines a parameter t, which depends on the two distributions [14]. WRAP performs the test for each appliance d_i and calculate the value of t_i as follows:

$$t_i = \frac{\mu_{i,j}^{H} - \mu_{i}^{W_S}}{\sqrt{\frac{\sigma_{i,j}^{H}}{n_H} + \frac{\sigma_{i}^{W_S}}{n_{W_S}}}} \tag{7}$$

where n_H and n_{W_S} are the numbers of samples used to calculate the historical distributions and the distributions over W_S, respectively. For each t_i it is possible to estimate the degree of freedom ν_i as follows [14]:

$$\nu_i \approx \frac{\left(\frac{\sigma_{i,j}^{H}}{n_H} + \frac{\sigma_{i}^{W_S}}{n_{W_S}}\right)^2}{\frac{\left(\sigma_{i,j}^{H}\right)^2}{{n_H}^2(n_H-1)} + \frac{\left(\sigma_{i}^{W_S}\right)^2}{(n_W)^2(n_{W_S}-1)}} \tag{8}$$

The test can verify if a change has occurred with a given probability. In particular, it is possible to determine if the alternative hypothesis is verified with probability α. To this purpose, given t_i and ν_i of appliance d_i, Student's t distribution tables give the value β_i, such that if $t_i > \beta_i$ then a change has occurred with probability α. WRAP has $O(1)$ complexity, since prediction, change detection and distributions update can be performed in constant time.

2) Prediction by the WRAP Algorithm: WRAP exploits the property of Gaussian random variables that the minimum mean square error estimate is equal to its mean [16]. The actual mean used for the estimation of an appliance d_i can be either the historical mean, or the mean in the recent time window W_S, depending on whether a change is detected for appliance d_i or not. As soon as the LSE alerts the smart home management system of a new power allowance, WRAP verifies whether a short-term change has occurred for each appliance d_i. If no change is detected for d_i for time slot τ_j, the algorithm uses the historical mean, i.e., $\lambda_{i,j} = \mu_{i,j}^{H}$. If otherwise a change is detected, then the algorithm uses the distribution of the most recent time window W_S, i.e., $\lambda_{i,j} = \mu_{i}^{W_S}$. Given the values of $\lambda_{i,j}$ for each appliance, (4) is used to calculate the importance factors $\gamma_{i,j}$, which are then used to calculate the emergency schedule described in Section IV-A.

3) Long-Term Change Detection: In order to detect long-term changes in the user behavior, WRAP uses a similar approach based on Welch's t-test. In particular, we compare the historical distribution $< \mu_{i,j}^{H}, \sigma_{i,j}^{H} >$ with the recent time window distribution $< \mu_{i}^{W_L}, \sigma_{i}^{W_L} >$. W_L can be set to several weeks. This test is performed periodically, e.g., daily or weekly. If a change is detected, the historical distribution no longer represents the current usage pattern of an appliance, and hence the new recent set of samples in W_L constitutes the historical distribution.

V. The Problem of a Load Serving Entity

This section describes how LSEs calculate the power allowance M_u for each user u, given the power allowance that resulted from the load curtailment at Bus l requested by the system operator $\Delta P_D(l)$. The following description focuses on a single LSE and, without loss of generality, assumes one LSE is responsible for all users at one bus, to drop the LSE index l. Let $M_{\max}$ be the maximum power allowance resulting from the curtailment for the considered LSE.

The goal of the LSE is to calculate the individual power allowance M_u for each user $u = 1, \ldots, N$, such that $\sum_u M_u = M_{\max}$, and the *aggregate user utility* is maximized. Aggregate user utility is defined as the sum of the utilities of all users served by the LSE. To this purpose, after the LSE receives a load curtailment request, the LSE inquires its users' smart home management systems and receives from them the importance factors for the current time slot predicted by WRAP. The LSE then needs to solve the following optimization problem for time slot τ:

$$\underset{x_i, M_u}{\text{maximize}} \quad \sum_{u=1}^{N} \sum_{i=1}^{n_u} \gamma_{i,\tau}^{u} x_i^{u} \tag{9}$$

$$\text{subject to} \quad \sum_{i=1}^{n_u} x_i^u e_i^u \leq M_u \ \forall\ u; \sum_{u=1}^{N} M_u = M_{\max} \quad (10)$$

where the user u has n_u appliances, e_i^u is the maximum power rating of appliance d_i^u, and $x_i^u = 1$ when u is allowed to use d_i^u, and 0 otherwise. Here, unlike the user's problem in (5), the power allowances M_u are variables to be determined, and the importance factors $\gamma_{i,\tau}^u$ are given by the smart home management systems.

Although the problem above is similar to the emergency schedule optimization in (5), it can suffer severe scalability issues because now the dimension of the problem is multiplied by a large number of users. For this reason, the framework exploits the following *regression based heuristic*, which relaxes the problem into a convex optimization.

A. Regression-Based Heuristic

Consider a specific user u, and let $M_{\max}^u$ be the maximum consumption that u can generate if he utilizes all his appliances at the same time, i.e., $M_{\max}^u = \sum_{i=1}^{n_u} e_i$. Setting $M_u = M_{\max}^u$ would obviously maximize the utility of user u. Since the LSE is aware of the importance factors $\gamma_{i,j}^u$ for each appliance i, it is also able to solve the optimization problem in (5). In fact, according to the heuristic, the LSE solves K instances of the problem for user u using different power allowance levels. At the k-th instance, it sets the power allowance to $\alpha_k M_{\max}^u$, where $\alpha_k \in [0, 1]$ and it is increased at each instance. Note that solving these problems, although NP-hard, is feasible thanks to the limited number of appliances per user.

Let $\delta_1, \ldots, \delta_K$ be the optimal solutions of such instances, where δ_k is the solution for α_k. The pairs (α_k, δ_k), $k = 1, \ldots, K$ are used to infer a continuous function $H_u : \mathbb{R}^+ \to [0, 1]$, which relates the power allowance M_u to the utility achieved by user u. A *regression* technique is adopted to approximate this function. Since $H_u(\cdot)$ is monotonically increasing, power law regressions can be used, that is $H(\cdot)$ can be approximated as $H(M) = \alpha M^\beta$, where $\alpha, \beta \in \mathbb{R}$ [17], and $\beta \leq 1$ to ensure that $H_u(\cdot)$ is concave.

The LSE calculates the functions $H_u(\cdot)$ for each of the N users, and solves the following optimization problem:

$$\underset{M_u}{\text{maximize}} \quad \sum_{u=1}^{N} H_u(M_u) \quad (11)$$

$$\text{subject to} \quad \sum_{u=1}^{N} M_u = M_{\max}; M_u \leq M_{\max}^u \ \forall u \quad (12)$$

The problem is a relaxation in the continuous domain of problem (9), and it returns an assignment of the maximum loads M_u to the users. Note that the problem does not solve for the decision variables x_u^i explicitly, which is the key for the reduction in complexity. Since the objective function is concave, the problem can be solved using standard convex optimization techniques [18].

VI. Simulation Results

The framework is tested on the IEEE 39-bus system modeled after the ISO New England system with 10 generators, 46 lines, and 21 nonzero load buses [19], [20]. The one-line diagram of this system is shown in Fig. 2. All nonzero load buses are assumed to have capability of load curtailment with the Internet of Things technologies, but not all these buses are necessarily subjected to a load curtailment in contingency cases. We use synthetic (randomly generated) and real traces (dataset) to model the user appliance usage. Real traces are taken from the data repository Tracebase [21], which collects the power consumption of various electrical appliances, with a resolution of several samples per second. Some of the considered appliances and their maximum power ratings are listed in TABLE II.

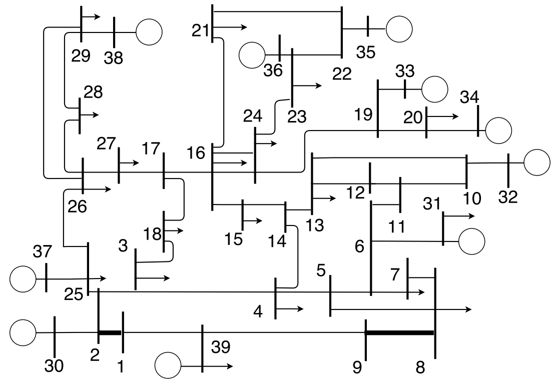

Fig. 2. The IEEE 39-bus system used in the simulations.

TABLE I
LOAD CURTAILMENT BY BUS

Bus ID	Curtailment (MW)	Bus ID	Curtailment (MW)
1	15.9	8	0.00492
4	0.0246	15	0.0355
7	0.000189	18	0.0315

A. The System Operator's Problem

The following simulations concern contingencies where a single line has failed. First, the one-component failure cases were identified by running the DC optimal power flow (OPF) problems with each line taken out. As long as the line failure did not isolate a generator bus with the rest of the system, the solution existed, and only these cases were studied in this work. When a generator is isolated as a result of a line failure, it changes the topology of the system, which changes the network matrix and the PTDF matrix H of the system.

This work focuses on the cases where the power flow solution exists even after one line failed. The case where Line 16 that connects Buses 8 and 9 (marked as a thick line in Fig. 2) has failed is presented. The resulting active power flows in the lines as a result of the DC OPF for this case are shown in Fig. 3(a).

As can be seen from the figure, Line 1 (between Buses 1 and 2, marked as a thick line in Fig. 2) resulted in a power flow very close to the limit, and the system operator may decide to reduce this line flow. The system operator can set the criterion in advance for which s/he decides to take actions and apply load curtailment. Once the criterion has been violated and load

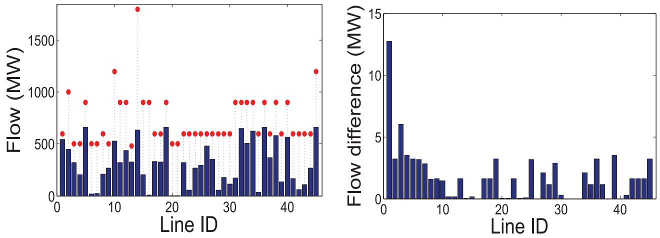

(a) (b)

Fig. 3. (a) Active power flows in the lines after Line 16 is out with the red markers denoting the line flow limits; (b) Active power flow difference in all lines before and after load curtailments.

curtailment is deemed appropriate, then ΔP_D is calculated for all the adjustable load buses, as described in Section III. In this work 10% margin of the line MVA rating is used as the reliability criterion. The difference of the absolute power flows before and after the load curtailment is depicted in Fig. 3(b).

The total amount of system load to be curtailed as a result is 16.04 MW, in order to reduce 6.37 MW of power flow in Line 1. Some generators were able to reduce their output due to this curtailment, and the output reduction from each of Generators 1, 3, 6, 9, and 10 was 3.25 MW, with the other generators' output unchanged. The amounts of load curtailment by bus is shown in TABLE I. All the other nonzero load buses that are not shown in the table did not have any curtailment. The other single-line failure cases yielded comparable results.

B. The Users' Problem

This subsection first studies the prediction accuracy of the WRAP algorithm. Subsequently, it analyzes the user utility achieved by the emergency schedule calculated with the prediction provided by WRAP.

1) Accuracy of WRAP: Synthetic traces are first used for simulations since changes in the data pattern can be manipulated at specific time instants. This shows the benefits of WRAP's change detection mechanisms in a controlled setting. Then the performance in real settings is analyzed using the real traces from the data repository Tracebase [21].

a) Synthetic traces: To generate synthetic traces, the length of each time slot is set to 1 hour. For each appliance d_i and time slot j, the utilization $\lambda_{i,j} \in [0, 1]$ is randomly generated. This represents the fraction of time that d_i is utilized on average during time slot j. Then, to simulate the variability of the user behavior, for each day the actual utilization at τ_j is generated using a Gaussian's distribution with mean $\lambda_{i,j}$ and variance σ.

The traces consider 200 days of observations, simulating a change in the user behavior by selecting a new value of $\lambda_{i,j}$, for each appliance d_i, every 50 days. Additionally, $\sigma = 0.2$, which according to the experiments, well approximates the realistic variability in user habits.

Since WRAP considers each appliance independently, the following experiment focuses on a single device. In particular, WRAP predicts the utilization $\lambda_{i,j}$ at the time slot j of the $(k+1)$-th day, using the values of $\lambda_{i,j}$ generated for the same time slot of the previous k days. The results show the accuracy of the predicted values in terms of the mean square error with respect to the actual values in the traces.

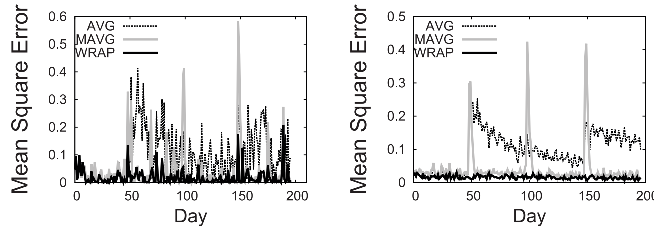

(a) (b)

Fig. 4. Prediction error: single time slot (a), average of all time slots (b).

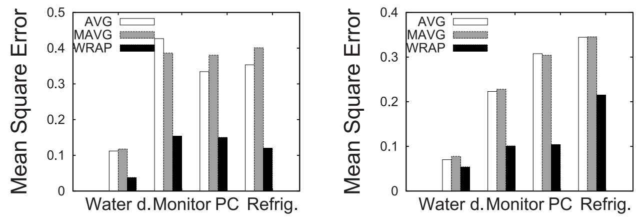

(a) (b)

Fig. 5. Prediction error: single day (a), average over 30 days (b).

WRAP is compared with two other standard prediction techniques. *Average* (AVG): this approach predicts the next $(k+1)$-th value of $\lambda_{i,j}$ as the average of all the previous k values. *Moving Average* (MAVG): this scheme predicts the next $(k+1)$-th value of $\lambda_{i,j}$ as the average of the last observed w values of $\lambda_{i,j}$, where w is the size of the time window. To ensure the reactivity of the approach, we set w equal to 5 days.

Fig. 4 (a) shows the results for the three approaches for a single time slot, while Fig. 4 (b) shows the average across all the time slots of a day. Before the first change occurs, all the three approaches perform similarly, since in the synthetic traces, the appliance utilization is drawn from the same distribution. However, after 50 days, a change in the usage pattern occurs. The error of AVG suddenly increases, since it keeps using all the previous dataset. MAVG, instead, is more reactive, but it still incurs high errors for a few days after the change. Additionally, it often overreacts to the fluctuation in user appliance utilization during stationary periods.

WRAP is able to promptly react to the changes and achieves significantly lower error than the other approaches, thanks to the short- and long-term change detection mechanisms. In fact, as soon as a change occurs, the short-term mechanism detects the change and predicts using only the most recent observed values. When the change in the user habits persists, the long-term mechanism eventually discards the former knowledge and only considers the observed values after the change.

b) Real traces: These experiments consider as appliances PC, refrigerator, monitor, and water dispenser, from the data repository Tracebase [21]. Note that these are the only appliances for which the repository provides at least 30 days of data. Similar experiments are performed as with the synthetic traces, in which k days of observation are used to predict the $(k+1)^{\text{th}}$.

Fig. 5(a) focuses on the 17$^{\text{th}}$ day and it shows the performance of the approaches averaging the mean square errors across all the time slots of that day. The results show that for

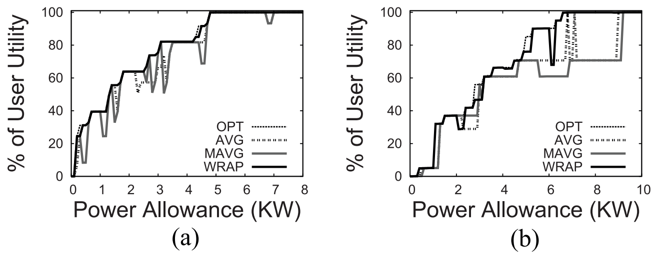

Fig. 6. Utility of a single user for time slot 11 a.m. (a) and 2 p.m. (b).

some appliances, such as the water dispenser, all approaches incur in a low estimation error due to the regular usage pattern. Differently, for other appliances, the results for AVG and MAVG strictly depend on the considered appliance. In particular, AVG outperforms MAVG for appliances, such as refrigerator and PC, which have a stationary utilization pattern, with minor short term variations. On the contrary, for appliances with non-stationary pattern, such as monitor, AVG is worse than MAVG, since the most recent days are more representatives of the future utilization. WRAP, always achieves the lowest error with respect to the other approaches, thanks to its adaptability to short and long term variations. Fig. 5(b) shows the average error over 30 days of predictions. The results confirm that our approach achieves the best performance. Note that, the difference between AVG and MAVG is smoothed by averaging over several days.

2) User Utility of Emergency Schedule: This section compares WRAP, AVG, and MAVG in terms of the individual user utility achieved by the corresponding emergency schedule. The emergency schedule defines the appliances that are allowed to be ON, and it is calculated by solving the optimization problem in (5) with the predicted importance factors and power allowance given by the LSE. The optimal schedule (OPT) is also shown for comparison, which is calculated by solving the optimization problem with the actual importance factors, assuming perfect knowledge of the user future behavior. The user utility of a schedule is calculated as the sum of the actual importance factors of the appliances allowed by the schedule.

The experiments consider real traces for 12 appliances. A contingency occurs on Day 16, and the previous 15 days are used for the prediction. Figs. 6(a) and (b) show the utility of a single user for time slots 11 a.m. and 2 p.m., respectively, under different power allowances given by the LSE (x-axis). The accuracy of WRAP allows to perform very close to the optimal, unlike the other methods. Note that only OPT yields a monotonically increasing utility with respect to the power allowance. This is because of the inaccuracy in predicting the importance factors by the other methods.

C. The LSE's Problem

As described in Section V, an LSE receives a load curtailment from the system operator, and it calculates the power allowance for its users to maximize the aggregate user utility. The LSE executes the regression-based heuristic to efficiently approximate the optimal solution of (9). This section studies the accuracy of the approximation provided by the heuristic. The experiments consider 8,000 users with 12 appliances each, using the real traces. Since the traces do not provide data for multiple users, the available data are replicated to represent the number of users in the simulations.

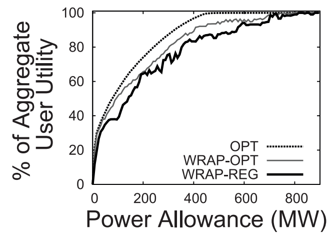

Fig. 7. Aggregate utility achieved by all methods with 8000 users.

Fig. 7 shows the aggregate user utility (i.e., the sum of individual users' utilities), expressed as the percentage of maximum utility achievable with no curtailment. The experiments show the performance of the regression heuristic (WRAP-Reg) given varying power allowances to the LSE (x-axis). The results are compared with the other two solutions: the optimal solution (OPT) assuming perfect knowledge of the future user behavior, which is the solution to the NP-hard optimization problem defined in (9); and an approximated solution (WRAP-OPT) where the LSE solves the same NP-hard problem but uses WRAP to predict the importance factors.

WRAP-OPT well approximates the optimal solution, again validating the accuracy of the prediction algorithm. However, this approach still requires to solve an NP-hard problem, and is thus not applicable in scenario where the LSE has a large number of users. Instead, WRAP-Reg achieves similar results close to the optimal, with significantly lower complexity. It should be noted that different prediction techniques would only impact the aggregate user utility, and not the load curtailment quantities at the system level.

Now the specific case of the IEEE 39-bus system presented in Section VI-A is considered. Recall that in this case scenario, Line 16 fails. The system operator, to prevent a cascading failure, notifies the LSE at Bus 1 that a curtailment of 15.9 MW is needed. It is assumed that the LSE's users' demand was 40 MW before the curtailment, therefore the power allowance at the LSE after curtailment is 24.1 MW.

The experiments compare OPT, WRAP-Reg and MAVG in this scenario to distribute the 24.1 MW to the users. The results of AVG are omitted since it performs similar to MAVG. To calculate the power allowances M_u's, OPT optimally solves (9), while MAVG evenly distributes the power among users, i.e., $M_u = M_{\max}/N \ \forall u$. To calculate the emergency schedule, OPT uses the actual importance factors, while MAVG uses the factors predicted with this strategy. WRAP-Reg adopts WRAP to predict the importance factors and the regression heuristic to distribute the load.

TABLE II shows the solution of the LSE problem, and the emergency schedules, for a specific user affected by the curtailment of the LSE at Bus 1. The LSE gives a power allowance to the user of 3,273 W under OPT, 3,267 W under Wrap-Reg, and 2,750 W under MAVG. The corresponding utility is 99.93%, 99.93% and 33.17%, respectively. Note that the goal of the framework is to maximize the *aggregate* user utility, not necessarily the utility of an individual user. In this

TABLE II
AN EXAMPLE OF THE APPLIANCE SCHEDULE OF A SINGLE USER

Appliance	Power rating (W)	OPT (3273 W)	WRAP (3267 W)	MAVG (2750 W)
Coffee maker	1305	ON	ON	OFF
Monitor	106	ON	ON	OFF
PC	166	OFF	OFF	ON
Refrigerator	1075	ON	ON	ON
Dish washer	2132	OFF	OFF	OFF
Water dispenser	360	ON	ON	OFF
User utility	-	99.64%	99.59%	33.17%

scenario, OPT achieves 83.1% aggregate utility, while WRAP-Reg 73.6% and MAVG 47%. The inaccuracy of MAVG in estimating the importance factors negatively affects both the power allowance distribution and the emergency schedule, and ultimately the aggregate user utility. Conversely, the high accuracy of WRAP-Reg is able to achieve only 10% less utility than OPT, which however assumes perfect knowledge of the future users' behavior and has higher complexity.

VII. CONCLUSION

This paper proposes, for the first time, a framework for contingency management that involves the system operator, the LSEs and the end-users. The framework enables the system operator to prevent subsequent failures by relieving lines possibly overloaded after the contingency. This is achieved through flexible loads at the user level realized with the emerging paradigm of the Internet of Things. The framework provides efficient algorithmic solutions to: 1) determine the curtailment at each bus, 2) calculate the resulting power allowance for each user and, 3) predict the user's near-future behavior to minimize the impact of the curtailment on the user utility. Results on the New England 39-bus test system, using real traces, show that the framework is effective in keeping lines within their capacity margins, with minimal impact on the user utility.

REFERENCES

[1] Executive Office of the President, "Economic benefits of increasing electric grid resilience to weather outages," White House Office Sci. Technol., U.S. Dept. Energy, Washington, DC, USA, Tech. Rep., 2013. [Online]. Available: http://energy.gov/sites/prod/files/2013/08/f2/Grid%20Resiliency%20Report_FINAL.pdf.

[2] U.S.–Canada Power System Outage Task Force, "Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations," Nat. Resour. Canada, U.S. Dept. Energy, Ottawa, ON, Canada, Tech. Rep., Apr. 2004. [Online]. Available: http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/Outage_Task_Force_-_DRAFT_Report_on_Implementation.pdf.

[3] U.S.–Canada Power System Outage Task Force, "Final report on the implementation of the task force recommendations," Nat. Resour. Canada, U.S. Dept. Energy, Ottawa, ON, Canada, Tech. Rep., Sep. 2006. [Online]. Available: http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf.

[4] L. Atzori, A. Iera, and G. Morabito, "The Internet of things: A survey," *Elsevier Comput. Netw.*, vol. 54, no. 15, pp. 2787–2805, 2010.

[5] A. Zanella, N. Bui, A. P. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, Feb. 2014.

[6] L. Goel, V. P. Aparna, and P. Wang, "A framework to implement supply and demand side contingency management in reliability assessment of restructured power systems," *IEEE Trans. Power Syst.*, vol. 22, no. 1, pp. 205–212, Feb. 2007.

[7] L.-R. Chang-Chien, L. N. An, T.-W. Lin, and W.-J. Lee, "Incorporating demand response with spinning reserve to realize an adaptive frequency restoration plan for system contingencies," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1145–1153, Sep. 2012.

[8] M. Humayun, A. Safdarian, M. Z. Degefa, and M. Lehtonen, "Demand response for operational life extension and efficient capacity utilization of power transformers during contingencies," *IEEE Trans. Power Syst.*, vol. 30, no. 4, pp. 2160–2169, Jul. 2015.

[9] A. Rabiee, A. Soroudi, B. Mohammadi-Ivatloo, and M. Parniani, "Corrective voltage control scheme considering demand response and stochastic wind power," *IEEE Trans. Power Syst.*, vol. 29, no. 6, pp. 2965–2973, Nov. 2014.

[10] K. Basu, V. Debusschere, and S. Bacha, "Residential appliance identification and future usage prediction from smart meter," in *Proc. IEEE IECON*, Vienna, Austria, Nov. 2013, pp. 4994–4999.

[11] A. T. Kaliappan, S. Sathiakumar, and N. Parameswaran, "Flexible power consumption management using Q learning techniques in a smart home," in *Proc. IEEE CEAT*, Lankgkawi, Malaysia, 2013, pp. 342–347.

[12] P. Wei, Y. Ni, and F. F. Wu, "Decentralised approach for congestion management and congestion price discovering," *IEE Proc. Gener. Transm. Distrib.*, vol. 149, no. 6, pp. 645–652, Nov. 2002.

[13] F. C. Schweppe, M. C. Caramanis, R. D. Tabors, and R. E. Bohn, *Spot Pricing of Electricity*. New York, NY, USA: Springer, 2013.

[14] T. H. Cormen, *Introduction to Algorithms*. Cambridge, MA, USA: MIT Press, 2009.

[15] B. L. Welch, "The generalization of students' problem when several different population variances are involved," *JSTOR Biometrika*, vol. 34, nos. 1–2, pp. 28–35, 1947.

[16] A. Gut, *An Intermediate Course in Probability*. New York, NY, USA: Springer, 2009. [Online]. Available: http://link.springer.com/book/10.1007%2F978-1-4419-0162-0.

[17] F. A. Graybill and H. K. Iyer, *Regression Analysis*. Belmont, CA, USA: Duxbury Press, 1994.

[18] O. Güler, *Foundations of Optimization*. New York, NY, USA: Springer, 2010.

[19] T. Athay, R. Podmore, and S. Virmani, "A practical method for the direct analysis of transient stability," *IEEE Trans. Power App. Syst.*, vol. PAS-98, no. 2, pp. 573–584, Mar. 1979.

[20] Illinois Center for a Smarter Electric Grid, Information Trust Institute, University of Illinois at Urbana–Champaign. (2015). *IEEE 39-Bus System*. [Online]. Available: http://publish.illinois.edu/smartergrid/ieee-39-bus-system/, accessed Feb. 10, 2016.

[21] A. Reinhardt *et al.*, "On the accuracy of appliance identification based on distributed load metering data," in *Proc. IFIP SustainIT*, Pisa, Italy, 2012, pp. 1–9.

Stefano Ciavarella received the degree (Hons.) in computer science from the Sapienza University of Rome, Italy, where he is currently pursuing the Ph.D. degree with the Computer Science Department. He is a Visiting Scholar with the Computer Science Department, Missouri University of Science and Technology. His research interests lie in the area of network management, wireless sensor networks, emergency scenarios in cyber-physical systems, and smart grid security.

Jhi-Young Joo (S'07–M'13) received the B.Eng. and M.Eng. degrees from Seoul National University in 2005 and 2007, respectively, and the Ph.D. degree from Carnegie Mellon University in 2013, all in electrical and computer engineering. She is an Assistant Professor with the Department of Electrical and Computer Engineering, Missouri University of Science and Technology. Her research interests are in modeling and optimization of power systems and markets.

Simone Silvestri (M'11) received the degree (Hons.) and the Ph.D. degree in computer science from the Sapienza University of Rome, Italy. He is currently an Assistant Professor with the Computer Science Department, Missouri University of Science and Technology. His research interests lie in the area of network management, hybrid wireless sensor networks, interdependent cyber-physical systems, and smart grid security. He serves as the Program Committee Member of several international conferences.