

Non-Intrusive Load Management Under Forecast Uncertainty in Energy Constrained Microgrids

Jonathan T. Lee^{*†}, Sean Anderson[†], Claudio Vergara[‡], Duncan S. Callaway^{*}

^{*}Energy and Resources Group, University of California, Berkeley, CA, USA

[†]New Sun Road, P.B.C., Richmond, CA, USA

[‡]ZOLA Electric, San Francisco, CA, USA

Abstract—This paper addresses the problem of managing load under energy scarcity in islanded microgrids with multiple customers and distributed solar generation and battery storage. We explore an understudied, practical approach of scheduling customer-specific load limits that does not require direct control of appliances or a market environment. We frame this as a stochastic, model-predictive control problem with forecasts of solar resource and electricity demand, and develop alternative solutions with two-stage stochastic programming and approximate dynamic programming. We test the efficacy of the alternative solutions against heuristic and deterministic controllers in an environment simulating the customers’ responses to load limits. We show that using forecasts to schedule limits can significantly improve power availability and the customers’ benefits from consumption, even without the controller having a full model of the customers’ responses.

Index Terms—Load management; microgrids; demand-side management; predictive control; optimal control

I. INTRODUCTION

Without measures for microgrid operators to manage load or communicate scarcity, customers in energy-constrained microgrids will experience suboptimal interruptions. For example, in an islanded microgrid with multiple customers sharing limited photovoltaic generation and battery storage capacity, high daytime loads on cloudy days might lead to evening interruptions of low-power / high-value loads such as lighting. This problem could exacerbate inequity across customers, for example, if some are only able to consume electricity in evening hours when interruptions are more prevalent.

We seek to improve the allocation of energy services in time by establishing dynamic *load limits* based on forecasts that allow customers to consume energy over a time window in quantities up to, but not in excess of, the limit. This control problem is related to other flavors of microgrid Energy Management Systems (EMS) and connected methodologically to recent work on Stochastic Unit Commitment (SUC). The classic unit commitment problem schedules generators to minimize startup, shutdown, and variable fuel costs while meeting

an estimate of inflexible load. The stochastic extension typically minimizes a measure of the expectation of costs over a set of uncertain *scenarios* while satisfying constraints [1]–[4].

Solutions to the stochastic microgrid EMS problem in the literature typically employ the same scenario approach as its SUC counterpart, but in different contexts with varying models of physical systems, points of control, and objectives. Generally, the microgrid has local intermittent renewable generation and energy storage, can be either grid-connected or islanded, may contain dispatchable generation, and may have controllable loads. If the microgrid is grid-connected, the main grid is treated as an unconstrained resource, but with a time-varying price entering the optimization problem [5]–[7]. In islanded or off-grid microgrids, dispatchable generation or flexible load [8] is used to balance supply and demand.

Our system of interest can be classified as an islanded EMS where supply-demand balance is met by flexible demand and storage dispatch, and lost load is assigned a cost in the EMS optimization problem. Prior related studies assume load is directly controllable [9], [10], or that customers respond to a pricing signal [7], [11]. Direct load control and time-varying prices are promising pathways; however, they have some limitations. Direct load control requires ubiquitous remotely controllable appliances and is intrusive to customers, particularly if very large demand shifts are required during periods of scarcity. Time-varying pricing requires carefully designed price formation rules and sufficiently responsive load.

In contrast, load limits require only broadcasting a limit to customers and the ability to disconnect load at the meter if the limit is exceeded. Although this approach is more blunt than direct load control or pricing, it is simple and inexpensive to implement. In the simplest case, the load limit can be sent to the customer directly via a mobile interface, in which case they would manually adjust their consumption. More sophisticated smart appliances could automate the adjustment for the customer, but are not required. In either case, using the total load limit preserves privacy and a degree of customer autonomy without distributed automation or a structured market. Load limits are advocated in [12], although that study works within the context of market-based solutions.

In the framework we present in this paper, an operator is held accountable implicitly for unreliable service and chooses load limits that maximize a simplified value metric of each customer’s energy consumption. We show that this choice

This material is based upon work supported by the U.S. National Science Foundation under Grant Nos. 1539585, CPS-1646612, and the Graduate Research Fellowship Program. This collaboration between authors was made possible by the California Renewable and Adaptive Energy group (CAL-RAE) and the Berkeley Hub for Energy Access Research (B-HEAR).

can be formulated mathematically as a *sequential decision problem*, which is well-known to have significant computational complexity [13]. This complexity is intensified by non-convexities in the model of demand subjected to load limits.

We develop two approximations and reformulations to reduce the complexity of the problem. The first uses two-stage stochastic programming with assumptions about the forecast to cast the problem as a less complex mixed integer quadratic program (MIQP). The second uses approximate dynamic programming in conjunction with two-stage stochastic programming to reduce the problem to a sequence of smaller MIQPs. We compare the performance of the approaches in simulation against a baseline model with no control, a heuristic, and a predictive controller that uses only the mean forecast.

The paper contributes a framework for developing stochastic, predictive models for controlling load through consumption limits under forecast uncertainty. The framework is novel in separating the decisions of the customer to respond to load limits from those of the operator to set them, providing a mechanism for evaluating controller performance in the face of model mismatch. We show how stochastic forecasts can be combined with approximate models of the customer response into an optimal decision model that can be solved with out-of-the-box numerical solvers. Our computational experiment results show significant benefits from using forecasts in a receding-horizon control framework, but more modest and variable benefits from using stochastic formulations in place of deterministic forecasts, with the conclusion that model mismatch limits the additional benefit from stochastic approaches. The paper provides a mathematical and computational foundation for exploring different formulations of value and mechanisms to allocate scarce electricity supply.

II. DECISION PROBLEM

We develop a method to set customer-specific load limits in a microgrid where multiple customers share distributed solar and battery storage with limited capacity. The load limit sets a maximum amount of energy that a customer can consume over a window of time. We assume a receding-horizon control (RHC) framework where a microgrid controller acts on behalf of the operator to compute both load limits and power injection setpoints at a fixed time interval, and then transmits these to customers, metering devices, and the distributed energy resources (DERs), as depicted in Figure 1.

The essential states are the state of charge of each battery belonging to each customer and the status of the loads and activities that each customer requires electricity for. The evolution of these states are affected by both the decisions of the microgrid controller and the customer. We assume the microgrid controller cannot control individual loads directly, but that customers can be sent a load limit that is enforced at their meter. We also assume the controller has no knowledge of the customer's decision model, activities, or individual loads, so its decision is to set an upper bound on uncertain consumption. However, we assume the controller is given an exogenously determined forecast of solar power potential

and electricity demand in the absence of consumption. We treat these forecasted variables as stochastic, which given the dynamic nature of the system, presents the controller with a sequential decision problem under uncertainty.

In the following subsections, we present first a relatively simple model of the customer's decision to adjust consumption given a load limit. The purpose of this model is both to capture model uncertainty from the controller's limited information and to define performance metrics for evaluating the control strategy from the perspective of the customer. We then formulate the controller's decision model to set load limits and propose specific approximate methods to make the problem tractable. This is the core contribution of the paper. Lastly, we briefly define a simple feedback controller to compute power injection setpoints to balance state-of-charge between batteries. The purpose of this component is mainly to facilitate simulating power-sharing among DERs and to provide a placeholder for future work to integrate the load-limiting control with optimal power flow models.

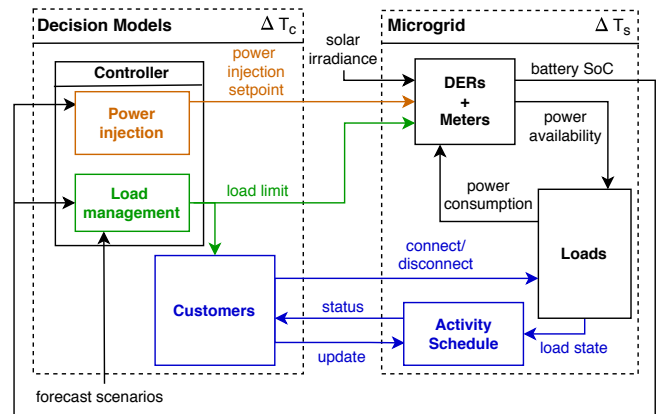


Fig. 1. Receding horizon control system

A. Customer decision and consumption model

We assume customers use their loads to conduct a set of activities that they schedule stochastically around a daily pattern. For example, lights are more likely to be used at night for several hours at a time, and microwaves around meal times for a few minutes. Based on assumptions about appliance ownership and usage patterns which are qualitatively consistent with our field experience, we randomly generate a schedule of activities for each customer that they would carry out if not subjected to limits. Customers derive a value when activities are completed without interruption, but incur an interruption cost otherwise. A customer can cancel an activity before it begins with zero cost but also zero gain.

When a customer is sent a load limit, we assume they cancel or interrupt activities and disconnect the associated loads to maximize their value of completing activities minus any interruption costs from activities already in progress. We introduce this model to emulating behavior to the first order and capture model error in the controller when evaluating performance.

We considered models of thermostatically controlled loads but determined this complexity did not provide additional insight, and recommend future work to comprehensively examine the effects of different types of shiftable and state-dependent loads.

Formally, we assume an activity a has a start time T_a^s , time to complete T_a^c , completion value v_a , interruption cost c_a , and a power consumption P_a when its associated load is on. The activity has two states: its remaining time to completion t_a^r and its status σ_a . The status evolves as a finite state machine with states: $\{0 = \text{queued}, 1 = \text{in progress}, 2 = \text{completed}, 3 = \text{interrupted while in progress}, 4 = \text{cancelled before commencing}\}$. We omit the formal transition rules as they are intuitive. Activities are initialized to $t_a^r = T_a^c$ and $\sigma_a = 0$. When the start time is reached, $\sigma_a \rightarrow 1$ and t_a^r decrements as time passes. Unless the activity is interrupted by either the customer or loss of power in the microgrid, $\sigma_a \rightarrow 2$ when t_a^r reaches zero. Statuses 2, 3, and 4 are terminal and the customer receives v_a for $\sigma_a = 2$ and pays c_a for $\sigma_a = 3$.

At a time t , when the customer is faced with a load limit of average power l over ΔT in the future, the sets of relevant activities are those that are already in progress $\mathcal{A}_1 := \{a \mid \sigma_a = 1\}$, and those that are queued but will start within the time window $\mathcal{A}_0 := \{a \mid \sigma_a = 0 \wedge T_a^s < t + \Delta T\}$. For each $a \in \mathcal{A}_0 \cup \mathcal{A}_1$, the customer chooses either $u_a = 0$ to cancel (for $a \in \mathcal{A}_0$) or interrupt (for $a \in \mathcal{A}_1$) the activity, or $u_a = 1$ to proceed as planned. The energy consumed by each activity over the time window is $P_a \min(t_a^r, \Delta T)$ for $a \in \mathcal{A}_1$ and $P_a \min(t_a^r, \Delta T - \max(T_a^s - t, 0))$ for $a \in \mathcal{A}_0$. For activities that will not be completed within the window, we assume the customer expects no load limit in the next window and effectively receives the completed value for activities still in progress. This allows us to represent their decision $u = \{u_a\}$ as an integer linear program to maximize their utility:

$$\begin{aligned} \max_u \quad & \sum_{a \in \mathcal{A}_0 \cup \mathcal{A}_1} u_a v_a + \sum_{a \in \mathcal{A}_1} u_a c_a \quad (1) \\ \text{s.t.} \quad & \sum_{a \in \mathcal{A}_0} P_a \min(t_a^r, \Delta T - \max(T_a^s - t, 0)) \\ & + \sum_{a \in \mathcal{A}_1} P_a \min(t_a^r, \Delta T) \leq l \Delta T \quad (2) \end{aligned}$$

B. Operator load-limit decision model with forecasts

We assume the microgrid uses RHC with fixed time-step ΔT_c over a horizon T . The controller decides on an action u_t to take on behalf of the operator at time t , based on the current state x_t and a probabilistic forecast \mathcal{W}_t of exogenous disturbances w_t . In our problem, x_t is a vector of the stored energy $E_{n,t}^{\text{stor}}$ in each customer n 's battery, u_t is the vector of load limits $l_{n,t}$, and w_t is the solar generation potential $P_{n,t}^g$ and electricity demand $P_{n,t}^l$ for each customer. We assume \mathcal{W}_t is a finite set of S scenarios consisting of distributed generation and demand values at each time over the horizon for each customer. Each scenario has a probability of occurrence p_s , which we assume to be uniformly $\frac{1}{S}$, but could be given explicitly by the forecast algorithm or tuned to hedge against particular outcomes. We assume the scenarios can be derived

from historical measurements, but do not present algorithms for doing so in this paper. The dynamics f are given by:

$$E_{n,t+1}^{\text{stor}} = E_{n,t}^{\text{stor}} + P_{n,t}^c \Delta T_c, \quad 0 \leq E_{n,t}^{\text{stor}} \leq E_n^{\text{max}} \quad (3)$$

where $P_{n,t}^c$ is the average net charge power into each customer's battery. $P_{n,t}^c$ is determined implicitly by the controller's action, the state variables and disturbances across all customers, and constraints defined subsequently, such as the capacity of each battery E_n^{max} and conservation of energy.

A critical detail in RHC is that the operator makes the next decision *after* observing a realization of the forecast w_t , the new state x_{t+1} , and given a new forecast \mathcal{W}_{t+1} ; however, to make the optimal decision u_t at time t , they have to compute what decision they *would* make at the next time-step given all possible outcomes, and so on over the horizon. This requires assuming how the forecast will be updated as realizations are observed, which we denote with the function g . The proper definition of g is ambiguous without additional information about the forecasting process, but has implications for the decision model; we discuss this in detail after stating the decision model in its general form.

The objective is to maximize the expected benefit of using electricity in the current time period plus the expected future benefit in subsequent time periods. This multi-stage decision problem can be represented mathematically in general with (4)-(8), where u'_t denotes hypothetical actions to take at the present time t and $\tau \in [t, t+T-1]$ denotes time-steps over the horizon. Note that the variables defined for $\tau > t$ are predicted future trajectories. Similarly f and g are models and do not necessarily match the physical or simulated system dynamics exactly. Q_t determines the expected benefit over the forecast horizon for any state and action, and is defined recursively as a sum of the expected present benefits b and the future benefits V_{t+1} given the new state and new forecast. V_τ is the maximum value from time τ assuming the operator acts optimally given state x_τ and forecast \mathcal{W}_τ .

$$u_t = \operatorname{argmax}_{u'_t} Q_t(x_t, u'_t, \mathcal{W}_t) \quad (4)$$

$$Q_\tau(x_\tau, u_\tau, \mathcal{W}_\tau) = \mathbb{E}_{\mathcal{W}_\tau} [b(x_\tau, u_\tau, w_\tau) + V_{\tau+1}(x_{\tau+1}, \mathcal{W}_{\tau+1})] \quad (5)$$

$$V_\tau(x_\tau, \mathcal{W}_\tau) = \max_{u_\tau} Q_\tau(x_\tau, u_\tau, \mathcal{W}_\tau) \quad (6)$$

$$x_{\tau+1} = f(x_\tau, u_\tau, w_\tau) \quad (7)$$

$$\mathcal{W}_{\tau+1} = g(\mathcal{W}_\tau, x_\tau, u_\tau, w_\tau) \quad (8)$$

$$V_{t+T}(x, \mathcal{W}) \equiv 0 \quad (9)$$

We assume for simplicity with (9) that the future benefit at the end of the horizon is zero regardless of the final battery state, but this can be replaced with any linear or quadratic function. We define the benefit b as a quadratic function of the *actually used* load power P^u averaged over a time-step. P^u is not directly controllable, but is a stochastic variable influenced nonlinearly by the load limit l , whose realization depends on the customer decision and information not available to the controller. To formulate the controller's decision, we model it

as (11), which is an overestimate of consumption because the customer is unlikely to be able to adjust *exactly* to the limit.

$$b(x_t, u_t, w_t) = \frac{1}{N} \sum_n \left(P_{n,t}^u - \frac{1}{2P_{n,t}^{l,\max}} P_{n,t}^u \right)^2 \quad (10)$$

$$P_{n,t}^u = \min(l_{n,t}, P_{n,t}^l) \quad (11)$$

The appropriate choice of b in different contexts is an important topic that requires careful study beyond the scope of this paper. We select the quadratic form for the common case where there is diminishing marginal value of consumption. In contrast, a linear function would value all consumption equally and effectively not steer the operator to take any actions to “keep the lights on” by reducing the usage of a few high power loads, which is our qualitative objective. We show in the results that using this form yields desired behavior despite b not representing any direct value. Eq. (10) can be modified in several ways while preserving the same structure: it can be shaped for different rates of diminishing marginal value, and weighted differently for particular customers over times of day. These parameters can be functions of past load limits or consumption. Note also that b is increasing up to the maximum possible load, $P_{n,t}^{l,\max}$, which is the power rating of their meter.

To specify g , one must assume whether each scenario represents a single trajectory, or a Markov process where the possible values at each moment in time are independent of prior values. The former implies up to S possible trajectories and final states, while the latter implies S^T , effectively leading to two different *scenario trees* after time $t + 1$, illustrated in Figure 2. Assuming for the illustration that the scenarios are unique over the first time-step, the two interpretations respectively imply that the operator assumes either, after observing w_t , that 1) they will know with certainty what trajectory they are on and then act optimally with perfect information, or 2) they will again face an uncertain forecast with no gained information.

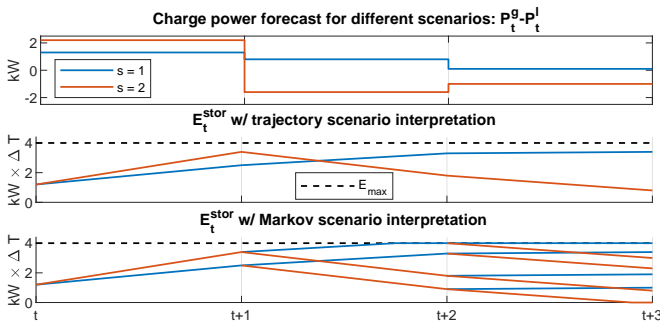


Fig. 2. Alternative interpretations of a 2-scenario forecast with a single battery over a horizon of three time-periods.

Both interpretations are approximations of the optimal decision because the forecast itself is an approximation of reality via a finite number of scenarios.¹ Here, we focus not on which

¹We refer the reader to [13, Ch. 6] for additional discussion showing how the trajectory interpretation can in fact be cast as an approximate solution to the Markov interpretation.

is correct – it depends entirely on the details of the forecasting algorithm – but develop solutions for both and compare their performance in simulation. We show that due to having fewer trajectories, the trajectory interpretation can be computed with two-stage stochastic programming, while the Markov interpretation requires additional approximate dynamic programming techniques to solve.

1) *Two-stage stochastic programming solution with trajectory forecast*: The key insight and distinction of our model from others is that the operator cannot directly control load, but can only indirectly influence it via a non-convex, piecewise-linear constraint as in (11). Otherwise, the problem employs the standard two-stage stochastic model by assuming each scenario is a distinct trajectory [5], [7], [9]:

$$Q_t(x_t, u_t, W_t) = \sum_s p_s \left(b(x_t, u_t, w_{t,s}) + \sum_{\tau=t+1}^{T-1} b(x_{\tau,s}, u_{\tau,s}, w_{\tau,s}) \right) \quad (12)$$

where the single recourse decision $u_{\tau,s}$ for each scenario s is a trajectory with a corresponding state $x_{\tau,s}$ from time $t + 1$. The optimization problem includes the constraints eqs. (13-26) with variables specified by customer n , scenario s , and over time τ as in (12). $E_{n,t,s}^{\text{stor}}$ is fixed at the initial condition $E_{n,t}^{\text{stor}}$ for each scenario. P^w is wasted solar (i.e. curtailed when batteries are full), and P is net flow into the network. $P_n^{c,\max}$ is the maximum charge power of a battery, assumed for simplicity to be the same as discharge power.

$$P_{n,\tau,s}^c = P_{n,\tau,s}^g - P_{n,\tau,s}^w - P_{n,\tau,s}^u - P_{n,\tau,s} \quad (13)$$

$$E_{n,\tau+1,s}^{\text{stor}} = E_{n,\tau,s}^{\text{stor}} + P_{n,\tau,s}^c \Delta T_c \quad (14)$$

$$0 = \sum_n P_{n,\tau,s} \quad (15)$$

$$0 \leq P_{n,\tau,s}^w \leq P_{n,\tau,s}^g \quad (16)$$

$$0 \leq P_{n,\tau,s}^u \leq P_{n,\tau,s}^l \quad (17)$$

$$-P_n^{c,\max} \leq P_{n,\tau,s}^c \leq P_n^{c,\max} \quad (18)$$

$$-P_n^{\max} \leq P_{n,\tau,s} \leq P_n^{\max} \quad (19)$$

$$\forall n \in [1, N], \forall \tau \in [t, t + T - 1], \forall s \in [1, S]$$

$$0 \leq E_{n,\tau,s}^{\text{stor}} \leq E_n^{\max} \quad (20)$$

$$\forall n \in [1, N], \forall \tau \in [t + 1, t + T], \forall s \in [1, S]$$

To cast the problem in a generic form for standard numerical optimization solvers, we replace (11) with the equivalent set (21)-(26) using binary variables $q_{n,s}$ and the constant $M_n := \max_s P_{n,t,s}^l$ [14]. These constraints, along with (17), give two disjoint cases for whether or not the load limit is binding in scenario s : $q_{n,s} = 1 \implies P_{n,t,s}^u = l_{n,t}$, and $q_{n,s} = 0 \implies P_{n,t,s}^u = P_{n,t,s}^l$. Note that the constraints only include the decision $l_{n,t}$ at the first time-step, and that only one decision is made for all scenarios, reflecting that the action must be taken before a scenario is realized. In contrast, the operator assumes they will be taking actions with certainty for $\tau \geq t + 1$, meaning they can set a load limit exactly to the desired consumption in that scenario. In the case where the optimal

load limit is the maximum over the forecast, i.e. $l_{n,t} = M_n$, then any $l_{n,t} \geq M_n$ is optimal, so the controller selects no load limit with $l_{n,t} = \infty$.

$$q_{n,s} \in \{0, 1\} \quad (21)$$

$$P_{n,t,s}^u \leq l_{n,t} \quad (22)$$

$$l_{n,t} \leq P_{n,t,s}^u + (1 - q_{n,s})M_n \quad (23)$$

$$l_{n,t} \leq P_{n,t,s}^l + (1 - q_{n,s})M_n \quad (24)$$

$$P_{n,t,s}^l \leq P_{n,t,s}^u + q_{n,s}M_n \quad (25)$$

$$P_{n,t,s}^l \leq l_{n,t} + q_{n,s}M_n \quad (26)$$

This is a mixed integer quadratic program (MIQP) with NS binary variables, and $\mathcal{O}(NST)$ continuous variables and constraints. This scaling in dimension is not to be confused with the complexity of solving the MIQP, which itself scales nonlinearly with the number of variables and constraints.

2) *Approximate dynamic programming solution with Markov forecast*: If the forecast is considered Markov, Eqs. (4)-(8) can be solved with backwards recursion, which in practice requires computing and storing values of $V_\tau(x)$ for each possible x . Computing this if each of N batteries is approximated with X discrete state-of-charge regions requires $(X + 1)^N$ samples, which is intractable. We address this by employing state-space aggregation, approximating the state by the sum of energy stored in all batteries \hat{x}_τ and sampling it uniformly at $X + 1$ points indexed by i . We denote samples of the aggregated state and value function $\tilde{x}(i)$ and $\tilde{V}_\tau(i)$. The continuous and sampled forms are related by piecewise linear interpolation in (31)-(32), with weights r_i satisfying SOS2 constraints defined for each scenario in (36)-(42).²

$$\hat{x}_\tau := \sum_n E_{n,\tau}^{\text{stor}} \quad (27)$$

$$\hat{x}^{\max} := \sum_n E_n^{\max} \quad (28)$$

$$\hat{x}_{\tau+1} = \hat{x}_\tau + \sum_n P_{n,\tau}^c \quad (29)$$

$$\tilde{x}(i) = \frac{i}{X} \hat{x}^{\max} \quad \forall i \in \{0, 1, \dots, X\} \quad (30)$$

$$\hat{x}_\tau = \sum_i r_i \tilde{x}(i) \quad (31)$$

$$\hat{V}_\tau(\hat{x}_\tau) := \sum_i r_i \tilde{V}_\tau(i) \quad (32)$$

Given the above, we can now define the optimization problem with objective (33) for computing the value function $\tilde{V}_\tau(i)$ at a sample of the state space i at time τ , given values of the next step value function at all samples of the state space $\tilde{V}_{\tau+1}(j) \forall j$, and forecast scenarios $w_{\tau,s}$:

$$\tilde{V}_\tau(i) = \max_{u'} \sum_s p_s(b(\tilde{x}(i), u', w_\tau, s) + \hat{V}_{\tau+1,s}) \quad (33)$$

The constraints are the same as the previous two-stage stochastic formulation $\forall s \in [1, S]$ and $\forall n \in [1, N]$, except that only

²SOS2 refers to ‘‘special ordered sets of type 2’’ constraints [14], which have a structure that can be exploited for better performance by some solvers.

one time-step τ is considered (the load limit constraints (22)-(26) are defined for time τ), the individual state-of-charge dynamics (14) are replaced with the aggregate dynamics (34) and likewise for battery capacity (35), and the SOS2 constraints are included:

$$\hat{x}_{\tau+1,s} = \tilde{x}(i) + \Delta T_c \sum_n P_{n,s}^c \quad (34)$$

$$0 \leq \hat{x}_{\tau+1,s} \leq \hat{x}^{\max} \quad (35)$$

$$\hat{x}_{\tau+1,s} = \sum_j r_{s,j} \tilde{x}(j) \quad (36)$$

$$\hat{V}_{\tau+1,s} = \sum_j r_{s,j} \tilde{V}_{\tau+1,s}(j) \quad (37)$$

$$\sum_j r_{s,j} = 1 \quad (38)$$

$$\sum_j y_{s,j} \leq 2 \quad (39)$$

$$y_{s,j} \in \{0, 1\} \quad (40)$$

$$0 \leq r_{s,j} \leq y_{s,j} \quad \forall j \in [0, X] \quad (41)$$

$$y_{s,j} + y_{s,k} \leq 1 \quad \forall j \in [0, X - 2], \forall k \in [j + 2, X] \quad (42)$$

In general, b should be redefined on the aggregated state space, but our form in (10) does not directly depend on state, so we use the same b . Note that (33)-(42) define an optimization problem only over one time-step. The solution process consists of starting at time $\tau = t + T - 1$ with $\tilde{V}_{\tau+1}(j) = \tilde{V}_{t+T}(j) \equiv 0$, solving the above problem to determine $\tilde{V}_\tau(i)$ for each $i \in \{0, X\}$, repeating for $\tau = \tau - 1$, and stopping after solving for $\tau = t + 1$. This entails solving $(X + 1)(T - 1)$ MIQPs, each with a dimension on the order of NS . Once \tilde{V}_{t+1} has been determined, we solve the problem again, but only given the initial state x_t to determine the optimal action u_t to take at time t using \tilde{V}_{t+1} as an approximation of V_{t+1} .

3) *Alternative deterministic solutions*: The two controllers of primary interest are described above, but we also define three alternative controllers for use in the computational experiments. The first trivially sets no load limit, the second sets limits according to the piecewise-linear feedback rule (43), using only the aggregated state of charge (27)-(28) and no forecasts, and the third uses a single forecast, computed as the mean over all scenarios, without considering uncertainty. The single forecast formulation is actually equivalent to the stochastic trajectory forecast with $S = 1$, making the binary variables extraneous and reducing the problem to a QP.

$$l_{n,t} = \begin{cases} 0.01P_n^{l,\max} & 0 \leq \hat{x}_t < 0.1\hat{x}^{\max} \\ 0.05P_n^{l,\max} & 0.1\hat{x}^{\max} \leq \hat{x}_t < 0.2\hat{x}^{\max} \\ 0.1P_n^{l,\max} & 0.2\hat{x}^{\max} \leq \hat{x}_t < 0.3\hat{x}^{\max} \\ \infty & 0.3\hat{x}^{\max} \leq \hat{x}_t \end{cases} \quad (43)$$

C. Power dispatch model

In a microgrid with DERs, a dispatch mechanism is required to maintain power balance and coordinate the charge power of each individual batteries. We model solar generation and batteries interfaced with grid-forming converters, where each

group n tracks a setpoint $P_{n,t}^{\text{inj},*}$ of power to inject into the network and the total imbalance is shared by an automatic generation control described below. The primary control objective here, given assumptions to ignore network constraints, is to keep states of charge equally balanced to each other to prevent losing instantaneous power capacity if some were to become drained before others. This is an open research area, but we achieve sufficient balancing with a simple, centralized, proportional feedback controller with gain $K = 2$ and ΔT_c the time-step between control action:

$$P_{n,t}^{\text{inj},*} = \frac{1}{K\Delta T_c} \left(E_{n,t}^{\text{stor}} - \frac{1}{N} \sum_n E_{n,t}^{\text{stor}} \right) \quad (44)$$

Integrating more sophisticated predictive power dispatch models with load-limiting to account for network constraints and losses is an important area for future work that becomes increasingly relevant in larger microgrids.

III. MICROGRID SIMULATION MODEL

To evaluate controller performance, we develop a simulation model of a distributed microgrid to capture interruption events and the evolution of battery states. We use a quasi-static simulation of the steady-state behavior of the the primary and secondary controls of the DER power converters, which govern power sharing and the availability of supply. We introduced grid frequency Δf_t as a state variable in the simulation model to maintain instantaneous power balance. The DERs act as synchronous interconnected areas that maintain power balance using classic droop and automatic generation control subject to constraints on the solar availability and battery charge [15]. We assume the charge and discharge capacity is constrained by the battery inverter rating $P_n^{c,\text{max}}$, the free capacity of the battery, and a linear power derating when the battery state-of-charge is within 10% of its limits. These dynamic constraints are captured respectively by the three terms in the min functions defining the maximum charge $P_{n,t}^{c,+}$ and discharge $P_{n,t}^{c,-}$:

$$\begin{aligned} P_{n,t}^{c,+} &= \min \left(P_n^{c,\text{max}}, \frac{E_n^{\text{max}} - E_{n,t}^{\text{stor}}}{\Delta T_s}, \frac{P_n^{c,\text{max}}(E_n^{\text{max}} - E_{n,t}^{\text{stor}})}{0.9E_{n,t}^{\text{stor}}} \right) \\ P_{n,t}^{c,-} &= \min \left(P_n^{c,\text{max}}, \frac{E_{n,t}^{\text{stor}}}{\Delta T_s}, \frac{P_n^{c,\text{max}} \cdot E_{n,t}^{\text{stor}}}{0.1} \right) \end{aligned} \quad (45)$$

The net injection $P_{n,t}^{\text{inj}}$ of each “area” n of DERs tracks the setpoint $P_{n,t}^{\text{inj},*}$ with a frequency response stiffness β_n subject to the charge and solar generation capacity constraints as well as conservation of energy given the loads $P_{n,t}^u$:

$$\begin{aligned} P_{n,t}^{\text{inj}} &= \min(P_{n,t}^g + P_{n,t}^{c,-}, \max(-P_{n,t}^{c,+}, P_{n,t}^{\text{inj},*} - \beta_n \Delta f_t)) \\ 0 &= \sum_n P_{n,t}^{\text{inj}} - P_{n,t}^u \end{aligned} \quad (46)$$

We set the stiffness of area n as proportional to the total inverter capacity: $\beta_n = \beta(P_n^{c,\text{max}} + P_n^{g,\text{max}})$ where $P_n^{g,\text{max}}$ is the PV inverter capacity, and we choose $\beta = 4$. The above system has either a unique solution for Δf_t or no solution; in the latter case, a blackout is implied. In the event

of a blackout, meters disconnect all load (thus interrupting customer activities) until the aggregate state of charge reaches 10%, and the DERs come back online automatically. When there is no blackout, the solar generation, curtailment, and battery charge are recovered from Δf_t and $P_{n,t}^{\text{inj}}$ by minimizing curtailment, and the battery stored energy is updated incremented by $P_{n,t}^c \Delta T_s$.

As shown in Fig. 1, the control system sets limits for each customer and a power injection setpoint for each DER every $\Delta T_c = 4$ hours. Within that window, the DERs, loads, meters, and customer activity states are simulated on a $\Delta T_s = 2$ minute time-step. We assume the customer updates their activity schedule whenever they receive a new limit and that individual meters enforce load limits by disconnecting load if the limit is exceeded.

IV. COMPUTATIONAL EXPERIMENTS

We conducted two computational experiments with multiple trials to assess the efficacy and computational tractability of the proposed algorithms using the experimental methodology and terminology proposed in [16]. All modelling code and data are available on GitHub, including the complete implementation of the models above, all experimental parameters, and additional data visualization.³ We ran the experiments on a personal computer using an Intel i7-7600U CPU Dual Core, 2.80 GHz CPU with 16 GB of memory. We used CVX version 2.1, build 1127, with MATLAB 2018a to develop the optimization problems with Gurobi 9.0.1 as the solver [17], [18]. In the simulation and timing results, we used MATLAB compiled binaries and the Gurobi API directly instead of CVX to improve performance.

In each trial, we simulate a microgrid of N customers by randomly distributing 300 W PV units and 2 kWh battery units with 1.2 kW charge power. We set the total solar capacity to produce the average unconstrained demand of 330 W, which was computed by simulating users’ activities, and 3 kWh of total battery capacity per kW of PV. This results in and average 1.5 kWp PV and 4.5 kWh of storage per user, but variable distributed, and ensures energy scarcity. Each customer is assumed to have a maximum possible load of $P_n^{l,\text{max}} = 10$ kW. Customers are assumed for simplicity to have the same activities and loads with parameters given in Table I, but multiple types are supported in the simulation. The tables dictating the probability of a customer scheduling an activity to start in each hour of the day are not shown for space reasons, but are available in the repository.

We used satellite-measured solar irradiance from a location on Lake Victoria, Uganda, spanning 2004 to 2019 at one minute resolution, to generate irradiance forecasts and realizations [19]. This region has active development of energy-scarce, isolated microgrids and exhibits daily variation in irradiance. In each experimental trial, we randomly select one year to use as realization, and draw S times with replacement from the remaining fifteen years for forecasts. We created

³Code: <https://github.com/Energy-MAC/pscc2020-load-limiting>.

TABLE I
ACTIVITY PARAMETERS (TIME IN MINUTES)

Activity	Watts	Min Time	Max Time	Compl. Val.	Int. Cost
Electronics 1	50	5	15	0.5	1
Electronics 2	75	30	180	4	2
TV	50	30	240	1	5
Lighting 1	300	5	260	2	10
Lighting 2	450	5	30	2	6
Microwave	650	2	10	2	5
Hair dryer	1800	2	17	2	5
Clothes Washer	500	30	60	3	5
Clothes Dryer	2500	45	60	3	5
Dishwasher	1200	60	90	3	5

sample load forecasts by simulating the customer load model with random activity schedules S times.

A. Controller Efficacy

In this experiment, we use $N = 7$ customers and $S = 15$ forecast scenarios with 48 hour horizons and simulate the RHC for 28 days. These, and the parameters defined in previous sections, comprise the experiment parameters. For each trial, we draw a random start day, random DER configuration, random customer activity schedule, and random forecasts and realizations as confounding variables. For these confounding variables, we compare each of the five controllers as independent variables: no load limit, proportional feedback, deterministic forecast, the two-stage model, and the approximate dynamic programming model. For each of these, we simulate the RHC and define three key performance metrics on the outputs: the value of the quadratic objective function (10) applied to realized consumption averaged over the 4 hour decision interval, the net customer utility derived from their successful completion and interruption of the loads, and the per unit average service availability index (ASAI), which is the fraction of time power was available averaged across customers (ASAI [20]). The objective values and customer utility are average per user per 4 hour time-step. These results are shown in Fig. 3, where the bar height is the median and the range shows the 5th and 95th percentile values across trials. We conducted 150 trials, observing that the coefficients of variation across trials for the performance metrics stabilize by 100 trials.

The predictive controllers significantly improve customer utility and power availability, but they do not improve the quadratic objective measure they explicitly maximize. We expect this is due to model mismatch where the controller assumes customers adjust load *exactly* to the limits, but they in fact reduce load *below* the limits. This is consistent with Fig. 4, which shows that the predictive controllers overestimate the objective even when accounting for forecast uncertainty. As expected, the no control case has the highest mean load because there is no curtailment. The objective values correspond closely to the mean load and are only slightly lower because the quadratic term is small, especially at normally low load, which leads us to conclude that the greater consumption drives the higher objective value.




Fig. 3. Key performance metrics across trials.

The key result is that despite the model mismatch, optimizing for the simple quadratic value of consumption produces an outcome that allows customers to respond to scarcity with lower interruption costs and greater utility. This may not be the outcome in some cases, for example if customers have very high-value and high-power, daytime loads, but if this is known to the microgrid operator, this can be addressed with weights in the benefit function. Further, gains on the feedback controller could be tuned to give better performance in particular cases, although it would likely be challenging to set gains that are effective across a wide variety of cases.

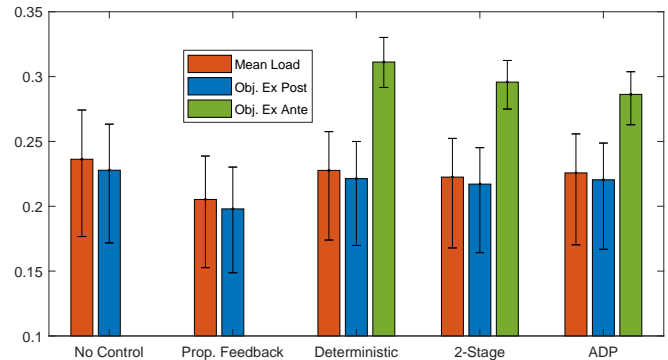


Fig. 4. Objective *ex post* and *ex ante* values with mean load.

Among the predictive controllers, the two stochastic approaches yield similar results to each other; however, Fig. 3 shows they tend towards slightly higher utility and minimally higher ASAI than the deterministic. The deterministic overestimates the objective relative to the 2-stage, which also overestimates relative to the ADP formulation if the forecast values are independent in time. This can be shown theoretically and is supported empirically in Fig. 4. In supplemental figures provided in the GitHub repository, we show that the stochastic approaches impose load limits more of the time but at higher and less restrictive levels when they are imposed, resulting in reductions in interruption costs. Essentially, they perform some effective hedging, but the benefits are small and the deterministic approach provides satisfactory performance.

To test computational performance, we varied the number of customers $N \in \{5, 15\}$, the number of scenarios $S \in \{5, 15\}$, and the time-steps in the forecast $T \in \{12, 24, 36\}$, and recorded the time for each formulation of the decision algorithm to converge to a solution. Table II shows the median time over 20 trials with random forecasts and initial states. We observed that for larger products of NS approaching the range of 300, the solver does not reliably converge within an hour, so we do not show results for problems of this size. We observed that for these problems that do not converge, approximate solutions are reached relatively quickly, but that thousands of successive iterations in the branch-and-bound algorithm continue with minimal improvements.

TABLE II
TIMING RESULTS (SECONDS)

N	S	T	Time: Det.	Time: 2 Stage	Time: DP
5	5	12	0.01	0.01	0.97
5	5	24	0.01	0.02	1.9
5	5	36	0.01	0.03	2.6
5	15	12	0.01	0.03	5
5	15	24	0.01	0.06	9.3
5	15	36	0.01	0.09	13
15	5	12	0.01	0.04	1.7
15	5	24	0.01	0.07	3.5
15	5	36	0.02	0.11	4.7
15	15	12	0.01	0.11	30
15	15	24	0.01	0.24	47
15	15	36	0.02	0.39	61

The results show that the approximate dynamic program generally takes longer to solve, but that the two-stage solution exhibits poor scaling with the forecast horizon. Both formulations are tractable for a real-time control scheme for products of NS up to around 100 with a forecast horizon of 24-36 hours. The tractability for larger products of NS requires more research into solver customization, appropriate solution tolerance, and convex relaxations.

V. CONCLUSIONS

This paper develops a mathematical framework for managing electricity consumption in energy-constrained microgrids by scheduling load limits to improve the availability and value of electricity service. We propose two techniques for incorporating stochastic forecasts into the decision to schedule load limits, and show how these can be modelled as mixed-integer programs. We find that both improve metrics of the value of electricity service and are tractable with an out-of-the-box MIQP solver for microgrids on the order of 15 customers, but that a deterministic approach, using only a single forecast, yields comparable performance improvements in our particular test case but with much lower computational complexity. Our modelling approach and simulation environment contribute a foundation for exploring different formulations of value and mechanisms to allocate scarce electricity supply.

- [1] S. Takriti, J. Birge, and E. Long, "A stochastic model for the unit commitment problem," *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1497–1508, 1996.
- [2] L. Wu, M. Shahidehpour, and T. Li, "Stochastic Security-Constrained Unit Commitment," *IEEE Transactions on Power Systems*, vol. 22, no. 2, pp. 800–811, May 2007.
- [3] P. Ruiz, C. Philbrick, E. Zak, K. Cheung, and P. Sauer, "Uncertainty Management in the Unit Commitment Problem," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 642–651, May 2009.
- [4] A. Papavasiliou, S. S. Oren, and R. P. O'Neill, "Reserve Requirements for Wind Power Integration: A Scenario-Based Stochastic Programming Framework," *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2197–2206, Nov 2011.
- [5] A. Parisio and L. Glielmo, "Stochastic Model Predictive Control for economic/environmental operation management of microgrids," in *2013 European Control Conference (ECC)*. IEEE, Jul 2013, pp. 2014–2019.
- [6] W. Su, J. Wang, and J. Roh, "Stochastic Energy Scheduling in Microgrids With Intermittent Renewable Energy Resources," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1876–1883, Jul 2014.
- [7] M. Rahimiyan and L. Baringo, "Real-time energy management of a smart virtual power plant," *IET Generation, Transmission & Distribution*, vol. 13, no. 11, pp. 2015–2023, Jun 2019.
- [8] G. Strbac, "Demand side management: Benefits and challenges," *Energy Policy*, vol. 36, no. 12, pp. 4419–4426, Dec 2008.
- [9] D. E. Olivares, J. D. Lara, C. A. Canizares, and M. Kazerani, "Stochastic-Predictive Energy Management System for Isolated Microgrids," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2681–2693, Nov 2015.
- [10] A. Parisio, E. Rikos, and L. Glielmo, "A Model Predictive Control Approach to Microgrid Operation Optimization," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 5, pp. 1813–1827, Sep 2014.
- [11] M. Vahedipour-Dahraie, H. Rashidizadeh-Kermani, H. R. Najafi, A. Anvari-Moghaddam, and J. M. Guerrero, "Stochastic security and risk-constrained scheduling for an autonomous microgrid with demand response and renewable energy resources," *IET Renewable Power Generation*, vol. 11, no. 14, pp. 1812–1821, Dec 2017.
- [12] K. Margellos and S. Oren, "Capacity Controlled Demand Side Management: A Stochastic Pricing Analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 706–717, Jan 2016.
- [13] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 4th ed. Athena Scientific, 2017, vol. 1.
- [14] E. M. L. Beale and J. J. H. Forrest, "Global optimization using special ordered sets," *Mathematical Programming*, vol. 10, no. 1, pp. 52–69, Dec 1976.
- [15] J. D. Glover, M. S. Sarma, and T. Overbye, *Power system analysis & design, SI version*. Cengage Learning, 2012.
- [16] J. D. Lara, J. T. Lee, D. S. Callaway, and B. Hodge, "Computational experiment design for operations model simulation," *Electric Power Systems Research*, in press.
- [17] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.
- [18] —, "CVX: Matlab Software for Disciplined Convex Programming, version 2.1," <http://cvxr.com/cvx>, Mar 2014.
- [19] "Copernicus Atmospheric Monitoring Radiation Service." [Online]. Available: <http://www.soda-pro.com/web-services/radiation/cams-radiation-service>
- [20] R. Billinton and R. N. Allan, "Reliability of Electric Power Systems: An Overview," in *Handbook of Reliability Engineering*, H. Pham, Ed. London: Springer-Verlag, 2003, ch. 28, pp. 511–528.

