

Analysis and Evaluation of an Distributed Optimal Load Coordination Algorithm for Frequency Control

Jonathan Brooks^{*}, Rodrigo D. Trevizan[†], Prabir Barooah^{*}, and Arturo S. Bretas[†]

^{*}Department of Mechanical and Aerospace Engineering
University of Florida, Gainesville, Florida

[†]Department of Electrical and Computer Engineering
University of Florida, Gainesville, Florida

Abstract—The Distributed Gradient Projection (DGP) algorithm was proposed in prior work to allow loads to provide contingency service to the grid using local noisy frequency measurements by varying their demand. Convergence of DGP was established in prior work for a decaying step size. In this paper we modify the algorithm to using a constant step size—constant step size being much more useful for practical implementation. We provide a convergence analysis of the modified algorithm, which we call DGP-C (DGP with Constant step size) and perform extensive simulations in the IEEE 39-bus test system. These studies (i) demonstrate that the DGP-C algorithm is robust to several assumptions made in the analysis and (ii) reveal which factors among the many tested (measurement noise, loads’ response time, etc.) have significant effect on the algorithm’s performance.

NOMENCLATURE

DGP	Distributed Gradient Projection
DGP-C	Distributed Gradient Projection with Constant step-size
HVAC	Heating, Ventilation and Air-Conditioning

I. INTRODUCTION

Many types of ancillary services are required for reliable operation of a power grid [1]. Primary control (PC) is especially important to maintain system stability [2]. Sometimes also referred to as primary frequency control, frequency response etc., primary control refers to actions taken to correct demand-supply imbalance in the first few seconds after a contingency event such as a generator or transmission-line trip. PC therefore must respond in the time scale of seconds. Typically, generators are tasked with providing this service. Frequency responsive loads such as induction machines always helped with primary control in a passive manner. In the smart grid paradigm, a complementary resource that can provide primary control is active loads equipped with sensing, communication and control algorithms.

Because of the need for speedy response, decentralized control of such smart loads, with each load using a locally obtained frequency measurement to determine its control, is attractive. There is a long history of this line of research, going back at least to the early 80’s and picking up interest more recently [3]–[8]. These works show that primary control can be achieved by varying demand as a function of grid frequency. Many of these works implement a droop-type control at the

loads, which varies the load in response to frequency but does not consider the resulting loss to the consumers’ utility, such as [6], [8]. Since consumers consume power to obtain some service or utility, and any change to their nominal power consumption to help the grid may cause a *disutility*. Some works use heuristics to prevent harm to the equipment, such as limiting the minimum time between turning a load on or off [6], [7]. Similarly, in works related to control of electric vehicle batteries for grid support, maintaining the state of charge within certain bounds is usually enough to maintain consumers’ utility [9], [10]. These works do not provide a general framework for distributing the effort among heterogeneous loads in a manner that minimizes the disutility to consumers.

The problem considered in this paper concerns heterogeneous consumers with distinct disutilities. The disutility experienced by a consumer is a function of the demand deviation from the nominal demand. The references [11]–[15] are highly relevant to this topic. The problem formulation in these papers share a common feature: they design distributed algorithms for loads to vary their demand to restore the grid-level demand-supply balance, while apportioning the demand changes so as to minimize the total consumer disutility. In [13] an additional layer is involved to coordinate the actions of the distributed generators. In both [11] and [14], a load uses local frequency measurement to estimate the demand supply imbalance and use that information to solve the optimization problem in a distributed manner. The consumer cost is required to be *strictly* convex in [11]–[13], [15] for their algorithms to be implementable.

A convex – but not necessarily strictly convex – cost is a more realistic model of how most consumers experience disutility, which is described in detail in Section II-B. The ref. [14] developed a distributed algorithm – called *Distributed Gradient Projection* (DGP) algorithm – that was also applicable to convex but not necessarily strictly convex consumer disutility. It was proven in [14] that the distributed decisions computed by the DGP algorithm converge to the central optima almost surely (i.e., with probability 1). The convergence analysis required the use of a decaying step-size.

In practice, the use of decaying step size suffers from severe limitations. As time increases, step size decays, and control effort decays as well. If a disturbance occurs at some large

time interval after the algorithm is turned on, the controllable loads will not react. While in principle this slowing down of the response can be ameliorated by “resetting” the step size, that can lead to spikes in the control effort—and thus in the grid frequency—at the resetting instants. This is unacceptable to grid operators. A constant step size is therefore far more desirable from a practical standpoint.

In this paper we propose a modification of the DGP algorithm: relaxing the step size to be a user-specified constant. The modified DGP algorithm - with a constant step size - is referred to as the *DGP algorithm with constant step-size* (DGP-C) in the rest of this paper. Just like the DGP algorithm, the DGP-C algorithm too can handle non-strictly convex consumer costs. We show that with a sufficiently small step size, the algorithm behaves in a stable manner, and establish a bound on the step size for such stable behavior. We also analyze the mean and variance for strictly convex consumer cost functions. Numerical tests with the IEEE 39 bus systems shows that predictions are still accurate even when consumer costs are non-convex.

A second contribution of this paper is extensive performance evaluation of the DGP-C algorithm in the IEEE 39-bus test system. Robustness to various simplifying assumptions made in the analysis is tested through these studies. These studies identify the design parameters and features of sensors and actuators that are needed for the algorithm to be successfully used in a practical setting. For instance, we find that the algorithm is robust to model mismatch and time delays. However, its performance suffers if the dynamic response of the loads is slow. For the factors that do have a significant effect, such as measurement noise, the simulations provide estimate of the ranges of the corresponding parameters needed to achieve good performance.

A preliminary version of this paper was presented in the conference paper [16]. Unlike this paper, no theoretical performance guarantees were provided there. Additionally, the numerical investigations carried out here are far more comprehensive than those in [16]. In particular, the simulations in [16] used three simplifying assumptions that are not likely to hold in practice, which are removed in this study: i) loads/aggregators measured the grid frequency and computed new demand values continuously in time (i.e., sampling period was neglected); ii) loads/aggregators were able to vary demand continuously within their minimum and maximum values (i.e., there was no quantization of changes in demand); and iii) loads/aggregators were able to attain their computed demand values instantaneously (i.e., dynamic response of loads was neglected).

This paper is organized as follows. In Section II, we summarize the DGP algorithm. We present analytical and numerical results in Sections III and IV, respectively. Finally, Section V presents conclusions of this work.

II. ALGORITHM DESCRIPTION

A. Problem Formulation

There are n agents in the power grid with some flexible demand. The change in demand for agent i , from its nominal demand, is denoted by x_i . The demand variation is limited to a range $[\underline{x}_i, \bar{x}_i]$, which is denoted by Ω_i :

$$x_i \in \Omega_i = [\underline{x}_i, \bar{x}_i].$$

For each agent i , there is a disutility function, $f_i(x_i)$, associated with agent i 's change in demand. There is a disturbance to the grid in the form of a sudden, uncontrollable change in generation or load, denoted as g , and the global difference between generation and demand (i.e., the demand-supply imbalance, ignoring losses) is denoted as $u = g - \mathbf{1}^T \mathbf{x}$, where $\mathbf{1}$ is a vector of all 1's and $\mathbf{x}$ is the vector of x_i . It is the objective of the agents to change their demand from their nominal values to obtain $u = 0$ while minimizing their total collective disutility. Formally, the agents are to solve the following optimization problem:

$$\min_{x_i, i=1, \dots, n} \sum_{i=1}^n f_i(x_i), \text{ s.t. } \sum_{i=1}^n x_i = g, x_i \in \Omega_i. \quad (1)$$

Each agent changes demand every T_s seconds, and at time kT_s , each agent i obtains a noisy measurement $\Delta\omega[k]$ of the grid frequency deviation $\omega[k] - \omega^*$ locally, where ω^* is the nominal frequency (usually 60 Hz). Using a state estimator, the agent infers the global power imbalance from the frequency measurements. The power imbalance estimated by agent i at time kT_s is denoted $u_i[k]$.

Additionally, there is a communication network among the agents, whose graph is denoted $G = (V, E)$, where the node set, $V = \{1, 2, \dots, n\}$, is the set of agents and the edge set, $E \subseteq V \times V$, denotes the pairs of agents that can exchange information. The set of agents with whom agent i can exchange information is denoted $N_i = \{j | (i, j) \in E\}$.

a) *Individual load vs. aggregator*: In principle, each agent can be a consumer load. However, in this paper we assume that the agents are load aggregators, which are more suitable for providing contingency services in deregulated electricity markets since individual consumers are not well suited to take part in such markets [17]. Additionally, we assume the aggregators act on the transmission level directly; see Figure 1. The balancing authority has a model of the transmission grid, which is supplied to the aggregators, especially in the state estimator used to infer grid-wide demand-supply imbalance from local frequency measurement.

Once the aggregator i computes its desired demand variation, $u_i[k]$, control actions for the individual loads within its territory need to be decided so that their combined effect produces $u_i[k]$. We do not consider the problem of computing actions for individual consumers within the aggregator in this paper; methods such as those proposed in [18] can be used for making such decisions.

It is important to note that the proposed algorithm is designed for continuous-state loads, but most individual loads are

discrete-state, e.g., on/off. Therefore, even in aggregation, the aggregate loads will still be discrete-state. We have imposed this feature in this work by quantizing the state space of the aggregators. That is, for the results presented in Section IV-D, we round each agent's power consumption to the closest bin, or discrete state, and the effects of this quantization are discussed more thoroughly in Section IV-H.




Fig. 1. The architecture for DGP algorithm's application. The dashed lines indicate communication links, solid lines indicate transmission lines in a power grid. Each agent that applies the DGP algorithm is a load aggregator that operates at the transmission-level.

B. Convex vs. strictly convex consumer cost

Consumers consume power in order to get some service, and a nominal demand profile provides the service they are used to. To help the grid, the demand has to vary from this nominal value, and any such variation can potentially cause consumers to experience a disutility. The disutility is usually modeled as a strictly convex function of the deviation. However, a convex but not necessarily strictly convex is a more realistic consumer model. For example, small variations in power consumption of a commercial heating ventilation and air-conditioning (HVAC) system lead to no perceptible change in indoor climate [19], but larger changes in demand can lead to a noticeable deviation from the set point. A strictly convex disutility cannot capture this phenomenon because every demand variation—no matter

how small—will have a nonzero cost. This difference is illustrated in Figure 2. The same is true also for Aluminum smelters; a small deviation from the nominal demand may not affect the smelting process at all, but beyond a particular threshold the cost is non negligible [20]. In fact one can argue that this is generally true. For any load whose demand can be varied in a continuous manner, variations below a certain threshold cause no perceptible change to the QoS experienced by the consumers. A grid operator can then sign long term contracts with such consumers: in return for the service provided by the loads, the consumer will be paid a fixed monthly amount, simplifying the contract structure and encouraging consumer participation. For variations beyond that threshold, which might be occasionally needed to handle large spikes in demand-supply imbalance, there will be a non-zero disutility to the consumer, and in that case the payment will be proportional to the disutility.




Fig. 2. Two distinct models of consumer disutility, f_2 which is strictly convex and f_1 which is convex but not strictly so. The cost function f_1 models a consumer who does not experience any cost if the demand deviation from the nominal is smaller than a threshold (in this case ± 25).

C. The DGP Algorithm

The update law of the DGP algorithm consists of three main parts: i) a generation-matching step, ii) a gradient-descent step, and iii) a projection step. The generation-matching step is used to drive the demand-supply imbalance, u , to 0. The gradient-descent step utilizes agent-to-agent communication to equalize the agents' gradients. Finally, the projection step limits each load aggregator i 's demand change to its feasible range, Ω_i .

The DGP algorithm is summarized below [21].

DGP Algorithm (at load aggregator i , time k):

- 1) Obtain an estimate $u_i[k]$ of the demand-supply imbalance u from measurement $\Delta\omega[k]$ using a state estimator. The *generation-matching step* is then $y[k]u_i[k]$, where $y[k]$ is a step size.
- 2) Compute gradient $\frac{d}{dx_i}f_i(x_i[k])$, transmit gradient value to neighbors, and receive neighbors' gradient values. Compute the *gradient-descent step* $\Delta x_i[k]$ as

$$\Delta x_i[k] = \sum_{j \in N_i} \nabla f_j(x_j[k]) - \nabla f_i(x_i[k]) .$$

3) Compute

$$x_i[k+1] = P_{\Omega_i} x_i[k] + c\gamma[k]\Delta x_i[k] + \gamma[k]\tilde{u}_i[k], \quad (2)$$

where $P_{\Omega_i}[\cdot]$ denotes the standard projection operator, $\gamma[k]$ is a step size, and c is a positive constant.

Note that any estimator may be used in step 1 that estimates demand-supply imbalance from grid frequency. In the simulations reported in this paper, we shall use the estimator in [16], which is the one used in [11]. In the sequel, boldface small letters will denote vectors obtained by stacking corresponding scalar quantities. For instance, $\mathbf{x}[k]$ will refer to the vector of $x_i[k]$'s.

III. PERFORMANCE ANALYSIS OF THE DGP-C ALGORITHM

The DGP algorithm with a constant step-size γ is called the DGP-C algorithm.

A. Convergence with a constant step size

In previous work [14], it was shown that with a decaying step size ($\gamma[k] \rightarrow 0$ as $k \rightarrow \infty$), the DGP algorithm converges almost surely to X^* , the set of optimal solutions of problem (1). In this paper, our focus is on the constant-step-size case, i.e., $\gamma[k] \equiv \gamma$ for some constant γ .

We make the following assumptions for our analysis, as in [14].

Assumption 1.

- 1) Each f_i is convex for each i with a (not necessarily unique) minimum at $x_i = 0$.
- 2) f_i is continuously differentiable for each i .
- 3) ∇f_i is Lipschitz for each i .
- 4) The communication graph, $\mathcal{G}$, is connected.
- 5) The solution set to (1) is contained within the strict interior of Ω .
- 6) The estimation error $\tilde{u}_i[k] := \tilde{u}_i[k] - u_i[k]$ is a martingale difference sequence for each i .

Since f is a design variable, the assumptions on f are not restrictive. The assumption on the connectivity of the communication graph means there is a path for information flow from any agent to any other. It has been established previously that $\tilde{u}_i[k]$ is a martingale difference sequence [11], [14] if the estimator described in [22] is used; however, the results in this section hold for any estimator that yields an error sequence that is a martingale difference sequence.

We now present the convergence result.

Theorem 1. *Let Assumption 1 hold. If $\gamma[k] \equiv \gamma > 0$, then for any $\delta > 0$, the fraction of time that $\mathbf{x}[k]$ (vector of $x_i[k]$'s in (2)) spends in the δ -neighborhood of X^* on $[0, T]$ goes to 1 in probability as $\gamma \rightarrow 0$ and $T \rightarrow \infty$.*

Proof. The proof invokes results from what is known as the o.d.e. method of stochastic approximation [23], which involves

analyzing a deterministic, continuous-time, o.d.e. analogue of the discrete-time iterations (2). The o.d.e. analogue of (2) is

$$\dot{\mathbf{x}}(t) = \Gamma_{\Omega} \mathbf{x}(t) [-cL\nabla f(\mathbf{x}(t))]^T + u(t)\mathbf{1}, \quad (3)$$

where L is the Laplacian matrix of the communication graph $\mathcal{G}$ [24], $\Gamma_{\Omega} \mathbf{y}(t)$ denotes the continuous-time projection operator of $\mathbf{y}(t) \in \mathbb{R}^n$ onto Ω . That is, the i th component of $\Gamma_{\Omega} \mathbf{y}[z]$ is

$$\Gamma_{\Omega} \mathbf{y}[z]_i = \begin{cases} 0; & y_i = \min \Omega_i, z_i < 0 \\ 0; & y_i = \max \Omega_i, z_i > 0 \\ z_i; & \text{o.w.} \end{cases}$$

Theorem 2.1 (Chapter 8) in [25] states that the fraction of time that $\mathbf{x}[k]$ spends in the δ -neighborhood of the limit set of (3) goes to 1 in probability as $\gamma \rightarrow 0$. It was proved in [14] that the solution to the ode (3) converges to the optimal set: $\mathbf{x}(t) \rightarrow X^*$, i.e., the limit set of the ode is X^* . Therefore, the theorem follows from the two results from [25] and [14].

Although convergence of the DGP-C algorithm established above is not as strong as the almost-sure convergence of the DGP algorithm with decaying step size that was established in [14], Theorem 1 indicates that the DGP-C algorithm can still “work well” provided the step size γ is small. The numerical results presented later, and in [16], are consistent with this prediction. Note that Theorem 1 merely states that good performance can be expected if γ is sufficiently small, but does not provide an estimate of how small γ should be. In contrast, the next result, Proposition 1, does provide such an estimate.

B. Asymptotic Behavior for Quadratic Cost

The following proposition characterizes the asymptotic behavior of the DGP-C algorithm. For the convenience of analysis, we limit ourselves to a quadratic model of consumer disutility. Although this is a rather strong assumption, the numerical results in the next section show that the results still predict the behavior of the DGP algorithm even when the disutility is not quadratic.

Proposition 1. *Let $\Omega_i = \mathbb{R}$ and $f_i(x_i) = q_i x_i^2$ for all $i \in \mathcal{V}$, where $q_i > 0$ for each i . Let λ be the eigenvalue of $A = I - \gamma(cLQ + \mathbf{1}\mathbf{1}^T)$ with maximum real part, where I is the identity matrix of size n , L is the graph Laplacian of the communication graph $\mathcal{G}$ [24], and $Q = \text{diag}(q_1, q_2, \dots, q_n)$. Then the dynamics of the loads' changes in demand are stable if and only if $\gamma < 2/|\lambda|$. Furthermore, the mean, $\mu[k]$, and covariance, $\Sigma[k]$, of the loads' changes in demand are given by*

$$\mu[k+1] = A\mu[k] + Bg\mathbf{1},$$

$$\Sigma[k+1] = A\Sigma[k]A^T + BW[k]B^T,$$

respectively, where $B = \gamma I$ and $W[k] = E[\tilde{u}_i[k]\tilde{u}_i[k]^T]$.

The proposition assumes $\Omega_i = \mathbb{R}$ for each i , i.e., that there is no projection. The nonlinearity in the dynamics (2) due to the

projection operator makes analysis challenging [26]. However, we can argue that for asymptotic analysis, the projection operator is inactive “most of the time” as follows. Theorem 1 shows that the iterates spend almost all of their time close to the optimal solution of problem (1) as long as the step size is small. Since the solution set of problem (1) lies within the strict interior of Ω by Assumption 1, for a small V , the iterates of the DGP algorithm spend most of their time within the strict interior of Ω . Hence, asymptotically the projection is inactive most of the time. Removing the projection altogether is a strong assumption, but as we will see in Section IV-B, Proposition 1 is still useful since it accurately predicts the behavior of the DGP algorithm *even when this assumption is not satisfied*.

Proof. For $\gamma[k] = \gamma$, we may rewrite the DGP update law (2):

$$\begin{aligned} \mathbf{x}[k+1] &= \mathbf{x}[k] + \gamma(-\mathbf{c}LQ - \mathbf{1}\mathbf{1}^T)\mathbf{x}[k] \\ &\quad + \gamma(g\mathbf{1} + [k]) \\ &= A\mathbf{x}[k] + B(g\mathbf{1} + [k]), \end{aligned} \quad (4)$$

where A and B are as defined above. Hence, the power consumption is a linear dynamical system under the DGP algorithm. Therefore, stability is attained if all eigenvalues of A lie strictly inside the unit circle. We begin by showing that the eigenvalues of $\mathbf{c}LQ + \mathbf{1}\mathbf{1}^T$ are real and positive. Because L is a Laplacian matrix, each column of L sums to 0 [24]. Since Q is diagonal, post-multiplication by Q scales the columns by diagonal entries of Q . It follows that the columns of $\mathbf{c}LQ$ also sum to 0, and $\mathbf{c}LQ$ has positive diagonal entries and nonpositive off-diagonal entries, just like the graph Laplacian. By applying Gershgorin’s circle theorem [27] to $(\mathbf{c}LQ)^T$ and utilizing the fact that a matrix and its transpose have the same eigenvalues, we conclude that $\mathbf{c}LQ$ has eigenvalues with nonnegative real parts. Because L and Q are real, symmetric matrices and Q is positive definite, $\mathbf{c}LQ$ has real eigenvalues¹. Now suppose $(\lambda, \mathbf{v})$ is an eigenpair of $\mathbf{c}LQ + \mathbf{1}\mathbf{1}^T$. That is,

$$(\mathbf{c}LQ + \mathbf{1}\mathbf{1}^T)\mathbf{v} = \lambda\mathbf{v}. \quad (5)$$

We consider two possibilities: $\mathbf{1}^T\mathbf{v} \neq 0$ and $\mathbf{1}^T\mathbf{v} = 0$. If $\mathbf{1}^T\mathbf{v} \neq 0$, then pre-multiplying (5) by $\mathbf{1}^T$ yields $n\mathbf{1}^T\mathbf{v} = \lambda\mathbf{1}^T\mathbf{v} \Rightarrow \lambda = n > 0$. If $\mathbf{1}^T\mathbf{v} = 0$, then (5) reduces to $\mathbf{c}LQ\mathbf{v} = \lambda\mathbf{v}$, so $(\lambda, \mathbf{v})$ is also an eigenpair of $\mathbf{c}LQ$. Thus λ is real and nonnegative. For $\lambda = 0$, we find a contradiction. Q has a trivial null space, which implies that $Q\mathbf{v}$ is an eigenvector of L associated with a zero eigenvalue. Because G is connected, there is only one such eigenvalue, and any corresponding eigenvector is parallel to the $\mathbf{1}$ vector, but each nonzero entry of Q is positive, so each element of $Q^{-1}\mathbf{1}$ is also positive. Therefore $\mathbf{1}^T\mathbf{v} \neq 0$. Thus we have a contradiction, so $\lambda \neq 0$. Therefore $\mathbf{1}\mathbf{1}^T + \mathbf{c}LQ$ has strictly positive eigenvalues.

¹This fact must be well known, but we were unable to find a reference. It can be proven as follows. Let $(\lambda, \mathbf{v})$ be an eigenpair of $\mathbf{c}LQ$, so $\mathbf{c}LQ\mathbf{v} = \lambda\mathbf{v}$. Then $\mathbf{c}V^*QLQ\mathbf{v} = \lambda\mathbf{v}^*Q\mathbf{v}$, where V^* denotes the conjugate transpose of V . Because L and Q are real, symmetric matrices, taking the conjugate transpose yields $\mathbf{c}V^*QLQ\mathbf{v} = \lambda^*\mathbf{v}^*Q\mathbf{v}$. From this and the previous equation, it follows that $0 = (\lambda - \lambda^*)\mathbf{v}^*Q\mathbf{v}$. Since Q is positive definite, $\mathbf{v}^*Q\mathbf{v} \neq 0$, which implies $\lambda = \lambda^*$. Therefore λ is real.

Now, let λ_i be an eigenvalue of $\mathbf{c}LQ + \mathbf{1}\mathbf{1}^T$. Then the eigenvalues of A are $1 - \gamma\lambda_i$. Thus, if $\gamma < 2/\bar{\lambda}$, where $\bar{\lambda}$ is the maximum eigenvalue of $\mathbf{c}LQ + \mathbf{1}\mathbf{1}^T$, then $|1 - \gamma\lambda_i| < 1$ for all λ_i . Therefore, A is stable.

That $\boldsymbol{\mu}[k+1] = A\boldsymbol{\mu}[k] + Bg\mathbf{1}$ follows from (4) and the fact that the estimator in [11], [22] is unbiased. To determine the steady-state covariance, let $\tilde{\mathbf{x}}[k] = \mathbf{x}[k] - \boldsymbol{\mu}[k]$, and note that the expression for $W[k] = E[\tilde{\mathbf{x}}[k]\tilde{\mathbf{x}}[k]^T]$ is given in [11]. The covariance at time $k+1$ is

$$\begin{aligned} \Sigma[k+1] &= E[(\mathbf{x}[k+1] - \boldsymbol{\mu}[k+1])(\mathbf{x}[k+1] - \boldsymbol{\mu}[k+1])^T] \\ &= E[(A\tilde{\mathbf{x}}[k] + B[k])(A\tilde{\mathbf{x}}[k] + B[k])^T] \\ &= A\Sigma[k]A^T + B[k]B[k]^T, \end{aligned}$$

where we have used that $\mathbf{x}[k]$ and $[k]$ are uncorrelated.

IV. NUMERICAL RESULTS

A. Simulation setup

The DGP-C algorithm was tested in the IEEE 39-bus test system, implemented in SimPowerSystems [28]. This system has 10 synchronous machines, each with governor control and PSS, and 19 load buses. We assume that at each load bus there is an aggregator, which can modulate a small percentage (5%) of the load’s consumption.

The gains used in the numerical tests are $c = 2$ and $\gamma[k] = 0.06/(\max\{\beta_i\})$ for all k . The disutility at load bus i is

$$f_i(x_i) = \begin{cases} 0, & |x_i| < a_i \\ \frac{1}{\beta_i}(x_i - a_i)^2, & x_i \geq a_i \\ \frac{1}{\beta_i}(x_i + a_i)^2, & x_i \leq -a_i \end{cases} \quad (6)$$

where $a_i \geq 0$ and $\beta_i > 0$. Note that if $a_i > 0$, the function f_i is convex but not strictly convex, and if $a_i = 0$ then f_i is a quadratic—and therefore strictly convex—function. To estimate the power imbalance using the frequency measurements, the load aggregators use a discrete-time, LTI model of the grid identified from the 39-bus system [16] (see Figure 3).




Fig. 3. Step response of IEEE 39-bus system and of the identified LTI model of the system that loads use in the DGP algorithm [16].

A note about the simulation results that will be presented in the next sections. In Section IV-B, we relax some practical constraints to show that the algorithm performs according to the theoretical predictions presented earlier. In Section IV-D, we evaluate the effects of a number of practical limitations present in real-life application.

B. Comparison with theory

Firstly, the DGP-C algorithm is tested in an ideal scenario, i.e., assuming loads can continuously vary demand with no constraints in terms of saturation, speed of actuation, or delay in communications. The only non-ideal aspect considered is the noise in frequency measurements, modeled as zero-mean Gaussian noise with a standard deviation (σ_f) of 0.01% of the value of the synchronous frequency (60 Hz). In this simulation, β_i is chosen equal to 0.2 for all load aggregators.

The goal of this test is to observe if the system will behave as predicted by theory when a disturbance in the form of a 150 MW increase in the nominal load in bus 27 is applied at $t = 10$ seconds. Two tests were run: i) $a_i = 0$ MW and ii) for $a_i = 2$ MW. The value 2 MW is roughly one quarter of 7.89 MW, which is the expected change in demand of all 19 load aggregators if they equally share the control effort. The resulting grid frequency deviation, $\Delta\omega$, is shown in Fig. 4(a), and values of x_i for the cases of $a_i = 0$ MW and $a_i = 2$ MW are depicted in Fig. 4(b) and 4(c), respectively. The speed deviation, $\Delta\omega$, shown is the mean deviation of all 10 generators from the nominal frequency (60 Hz).




Fig. 4. Simulation on the IEEE-39 bus test system, with both strictly and not-strictly convex disutilities, for a contingency event at $t = 0$. The control actions (demand variation) of all the load buses are shown (bottom two plots). The mean and std. deviation $\mu[k]$, $\sigma[k]$ are computed from Proposition 1.

In Fig. 4(a), we notice that the DGP-C algorithm quickly reduces the frequency deviation of the system to zero in both cases, in spite of the highly noisy frequency measurements. Fig. 4(b) and 4(c) show that the demand changes remain within three standard deviations ($\pm 3\sigma[k]$) of the mean ($\mu[k]$),

where $\mu[k]$, $\sigma[k]$ are computed from Proposition 1. These results numerically confirm the accuracy of the predictions by Proposition 1. This accuracy is achieved despite the strong assumptions made in the proposition, that the projection operator is not active. In fact, even the assumption that the optimal set, X^* , lies in the strict interior of Ω is violated for this scenario.

It is important to stress that even when a non-strictly convex and non-quadratic disutility function was used ($a_i = 2 \neq 0$), the numerical results are consistent with the prediction of Proposition 1 that assumed $a_i = 0$. We believe this is due to the fact that the not-strictly convex disutility functions used in the simulations were asymptotically quadratic. In the remainder of the paper we use the non-strictly convex cost function defined by (6) with $a_i > 0$.

C. Constant vs. decaying step size

We now present numerical results to corroborate the motivation for using a constant step size in DGP-C instead of a decaying step size. Figure 5 shows a simulation comparison between DGP-C (with constant step size) and DGP with decaying step size, with and without periodic reset.

For the decaying-step-size case, we use $\gamma[k] =$




Fig. 5. Comparison of performances of DGP with constant step size, decaying step size with timer reset, and decaying step size without timer reset.

$0.06/(19 \max\{\beta_i\}) \cdot 1/(k \cdot T_s)^{0.8}$, for $T_s = 2/60$ s. For the “decaying step size with reset” case, the timer is reset every 20 seconds: $\tau_{rst} = 20$ s, when k is reset to 0. Two contingency events were simulated: a sudden load increase of 150 MW at bus 27 at $t = 10$ s and a load decrease at bus 27 of the same 150 MW at $t = 50$ s. As we can see from Figure 5, DGP’s performance is poorest without timer reset. This is expected since beyond a certain value of the time index k , the step size $\gamma[k]$ is too small for the loads to respond. With decaying step size but with timer reset, performance is improved, but still poorer than that of DGP-C, especially if the reset time and moment when the contingency occurs are out of synchrony, which is what we would expect in practice.

D. Evaluation of DGP-C algorithm under practical limitations

Real-life implementations of the DGP-C algorithm would face a number of limitations whose model is not included in the formulation of the proposed solution, such as: i) model

TABLE I
STANDARD VALUES FOR PARAMETRIC STUDIES

Parameter	Nominal
τ_d	Delay in communication between aggregators
T_s	Sampling period of controller
$ \Omega_l $	Number of bins (discrete states) of loads
σ_f	Standard deviation of frequency measurements
σ_l	Standard deviation of noise in load demand
f_c	Cutoff frequency of load dynamics
	100 ms
	2/60 s
	51
	0.01%
	0.01%
	0.32 Hz

mismatch, ii) communication delay between load aggregators, iii) communication topology, iv) frequency measurement noise, v) presence of uncertain, uncontrollable renewables and stochastic loads, vi) sampling period of discrete controller and measurements, vii) quantization of controller output, and viii) bandwidth of the actuator. We did not evaluate analytically these known effects in Section III, but in Sections IV-D to IV-H we present case studies that assess the effect of those practical limitations.

The parameters related to the practical considerations, along with their standard values used in the simulations, are presented in Table I, and the definition of all practical considerations are provided along with the presentation of the corresponding parametric study. When one parameter is varied to study its effect, others are held constant at their nominal values. The parameters of the cost function (6) are $\alpha_i = 0.05\bar{x}_i$ and β_i is chosen from a uniform distribution on $[0.1, 0.3]$ for each i . The nominal edge set for the communication graph is $E = \{(i, j) | |i - j| = 1\}$.

We have studied two types of disturbance. The first is a disturbance in the form of a 150-MW increase in load at bus 27 applied to the system between 10 seconds and 50 seconds. The second type of disturbance is the loss of generator 5 (508 MW) at $t = 5$ seconds, which also represents the introduction of a model error because the model used by the state estimator (Fig. 3) that considers the original 10 generators does not match the model of the 9-generator system under disturbance. All parametric tests were performed for both disturbance types. Because most of them showed analogous results, we will show the results for the 150-MW load disturbance, except when model mismatch is concerned.

E. Effect of Time Delay and Model Mismatch

To evaluate model mismatch and the effect of communication time delay, we disconnect generator 5 (508 MW) from the system at 5 seconds; generator 5 accounts for nearly 10% of the total generation in the system. There is a communication delay (τ_d) for each load aggregator, i.e., each aggregator only has access to past values of its neighbors' gradients. To assess the effects of communication delay, we use three different values for each scenario: $\tau_d = 0$ ms (no delay), $\tau_d = 100$ ms, and $\tau_d = 1$ second of communication latency between adjacent nodes. Fig. 6 shows the results of the simulation where generator 5 is disconnected from the grid. Results for nominal operation (without smart loads) are shown for comparison. The

frequency deviation caused by the generator's disconnection is halved compared to the scenario without smart loads. The DGP-C algorithm achieves this while using the original model, which is no longer accurate due to changes caused by the generator disconnection.




Fig. 6. Effects of change in topology and communication delay: loss of generator 5 (508 MW).

For this type of disturbance, the performance of the method is only marginally affected by the time delay in communications, as shown in Fig. 6. For the 150-MW load-increase disturbance (omitted in this paper), a larger time delay in communications has a more noticeable degradation in the response time of the DGP-C algorithm. In particular, the longer delay results in a longer time for the disutility to reach steady state. This happens because the delay affects the gradient-descent step in the algorithm, which corresponds to minimizing disutility. This difference between the response of both types of disturbance happens because the loss of the generator creates a very large mismatch between load and generation, thus resulting in a very high value for $u_i[k]$ in (2). Therefore, in the case of a large power imbalance, such as the case where a generator is disconnected, the response of the DGP-C algorithm is dominated by the generation-matching term ($y[k]u_i[k]$), which is not affected by communication delay because it only depends on local frequency measurements, while the gradient-descent term ($cy[k]\Delta x_i[k]$) has smaller values. The same effect is not reproduced in the case of load increase, where the generation-matching term is not as dominant.

F. Effect of Measurement Noise

From this section on, we will only consider the case where a 150-MW increase in load occurs at bus 27 between 10 and 50 seconds. To evaluate the effect of frequency measurement noise, we have set the standard deviation of the frequency noise measurements to $\sigma_f = 0.01\%$, 0.1% , and 1% . Even though the frequency measurements in this implementation are supposed to be taken on the transmission side, where measurements tend to be accurate, we test the performance of the method with larger noise since the method is flexible enough to be used in microgrids and distribution systems, where robustness to noisy frequency measurements is more important.

The results shown in Fig. 7 demonstrate that a very noisy measurement can severely harm the performance of the DGP-C algorithm in terms of reducing frequency deviation. When compared to the performance of the system without smart loads, the results are superior for the cases where the standard deviation of noise is 0.01% and 0.1%, and they are similar when it is equal to 1%. *We conclude that the quality of frequency measurements is critical for this application.* It is important to note that the control systems of the synchronous machines that keep the system stable do not rely on noisy measurements.




Fig. 7. Effect of noise in frequency measurement: 150-MW load disturbance at bus 27.

If we compare this result to Fig. 4, we notice that the practical limitations introduced from Section IV-D onwards have greatly degraded the performance of the DGP-C algorithm, which is not capable of returning the frequency of the system to its nominal value.

G. Speed of Actuation

We consider that there is a phase lag between the control action and the change in demand due to the dynamics of loads and delay in communications (between aggregators and their subordinate loads). Because loads cannot respond instantaneously, their dynamics were modeled as a low-pass filter. We examine the performance with various cutoff frequencies (f_c), including 0.032 Hz and 0.32 Hz, which were inspired by the cutoff frequencies experimentally obtained for commercial-building HVAC fans [19] and variable-speed heat pumps [29], respectively. We also tested for cutoff frequencies of 0.01 Hz and 3.2 Hz to emulate loads with slower and faster responses, respectively, as shown in Fig. 8.

The results show that good performance can be obtained for all cases, when compared to the case without smart loads. Tests for f_c of 0.01 Hz and below, however, have shown that the performance of smart loads whose response is very slow (can only respond in tens of seconds) is less effective or even ineffective for the time range studied in this simulation.




Fig. 8. Effect of cutoff frequency of loads: 150-MW load disturbance at bus 27.

In practice, the control system and measurements would be digital (therefore calculated in discrete time). Unless otherwise noted, the values for sampling period (T_s) of controller and measurements is 2/60 seconds. Tests were performed for sampling periods of $T_s = 2/60$, 5/60, 10/60, and 0.5 seconds. The results in Fig. 9 show that the performance of the DGP-C algorithm is harmed by increasing the sampling period beyond 2 cycles, and for $T_s > 10/60$ seconds, large excursions of all signals—especially total disutility—appear. For $T_s = 0.5$ seconds, the method is no longer more efficient than the baseline case without smart loads.




Fig. 9. Effect of different sampling periods on the performance of DGP-C under a 150-MW load increase at bus 27.

H. Other effects

Other parametric tests were performed and the results are summarized below.

- **Graph topology.** The scenario of disturbance by load increase was re-examined using a different communication-graph topology, having a connected graph with smaller

graph diameter. The results are similar to those to the graph with the original topology. We believe the reason is the use of local frequency measurements that provide global information. As a result, change in the communication topology do not change the speed of convergence much.

- **Number of bins.** Suppose each agent can only change its demand by discrete increments; that is, changes in demand are placed into “bins.” Because the DGP-C algorithm was designed for continuous changes in demand rather than discrete, for implementation with such discrete loads, each load aggregator rounds its calculated demand change to its closest bin. The effect of changing the number of bins was studied by considering $|\Omega_d| = 5, 9, \text{ and } 51$. The results demonstrate that the DGP-C algorithm can function adequately for all three cases. However, the frequency deviation was larger for a smaller number of bins.
- **Stochastic loads.** The net load, which is nominal demand minus renewable generation, is a stochastic process due to the randomness in demand and renewable energy generation. The effect of intermittent and uncertain renewable energy generation was modeled as an increase in the fluctuation of the demand from loads. Those variations can be decomposed into small- and large-scale deviations. Large variations can be modeled as a step in consumption of a given load, such as the 150 MW increase disturbance applied to bus 27. The random small-scale fluctuations of both renewable generation and load are modeled as Gaussian additive noise in demand of each load with zero-mean and standard deviation (σ_l) equal to 0.01% of the nominal demand of the load. We have tested standard deviation of each load from $\sigma_l = 0.01\%$ to 1%. The results show that the DGP-C algorithm can compensate for the 10-fold increase in the uncertainty of load, even when $\sigma_l = 1\%$. In comparison to the case without smart loads, frequency deviations are reduced by the DGP-C algorithm.

V. CONCLUSION

We provided analytical and numerical results for the DGP-C algorithm with measurement noise and other practical issues. A bound on the step size was established for stability, and the mean and variance of the demand changes were analyzed. Despite the strong assumptions used in the analysis of mean and variance, the theoretical results still accurately predict the behavior of the DGP-C algorithm seen in the simulations, in which those assumptions are not satisfied.

The DGP-C algorithm was shown to successfully arrest frequency deviations from the nominal value in most scenarios. In particular, through simulations we examined the effects of model mismatch, communication delay, communication topology, frequency-measurement noise, presence of uncertain, uncontrollable, stochastic renewable generation, sampling period for discrete-time control and measurements, quantization of controller output, and speed of actuator response.

The simulations revealed that the DGP-C algorithm is robust to most of these factors, but it is sensitive to noise in frequency measurements and sampling rate of implementation. Lower cutoff frequencies for actuator bandwidths significantly degraded performance compared to higher cutoff frequencies. This implies that only sufficiently fast loads can effectively provide frequency control using the DGP-C algorithm, as expected. A cutoff frequency corresponding to the bandwidth of a commercial HVAC fan’s power deviation still significantly improved performance compared to the scenario without smart loads. This indicates that commercial HVAC systems are fast enough to provide meaningful frequency control service to the grid.

In this work network constraints such as line flow limits were not taken into account in setting up the optimization problem [15]. This is an area of future work. The analysis in this paper assumes that feedback interconnection between the smart loads and the rest of the power grid does not cause instability. Another important direction of future work is the analysis of closed-loop stability.

ACKNOWLEDGMENT

The research reported here was partially supported by the National Science Foundation under grant CPS-1646229, by the DOE GMLC program under grant “Virtual Batteries”, and by Nanjing SilverMicro Electronics Ltd. under grant 19050100.

REFERENCES

- [1] B. Kirby, “Ancillary services: Technical and commercial insights,” 2007, prepared for Wärtsilä North America Inc.
- [2] Prepared by NERC RS Committee, “Balancing and Frequency Control: A Technical Document Prepared by the NERC Resources Subcommittee,” NERC Technical Report: <http://www.nerc.com/comm/OC/Pages/RS/Resources-Subcommittee.aspx>, pp. 1–53, January 26 2011.
- [3] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox, “Homeostatic utility control,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-99, no. 3, pp. 1151–1163, may 1980.
- [4] D. J. Hammerstrom, J. Brouss, D. P. Chassin, G. R. Horst, R. Kajfasz, P. Michie, T. V. Oliver, T. A. Carlon, C. Eustis, O. M. Jarvegren *et al.*, “Pacific northwest GridWise testbed demonstration projects; part II. Grid Friendly Appliance project,” Pacific Northwest National Laboratory (PNNL), Richland, WA (US), Tech. Rep., 2007.
- [5] J. A. Short, D. G. Infield, and L. L. Freris, “Stabilization of grid frequency through dynamic demand control,” *Power Systems, IEEE Transactions on*, vol. 22, no. 3, pp. 1284–1293, 2007.
- [6] A. Molina-Garcia, F. Bouffard, and D. S. Kirschen, “Decentralized Demand-Side Contribution to Primary Frequency Control,” *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 411–419, 2011.
- [7] P. J. Douglass, R. Garcia-Valle, P. Nyeng, J. Ostergaard, and M. Togeby, “Smart demand for frequency regulation: Experimental results,” *Smart Grid, IEEE Transactions on*, vol. 4, no. 3, pp. 1713–1720, 2013.
- [8] H. W. Qazi and D. Flynn, “Synergetic frequency response from multiple flexible loads,” *Electric Power Systems Research*, vol. 145, pp. 185–196, 2017.
- [9] J. C. Hernández, P. G. Bueno, and F. Sanchez-Sutil, “Enhanced utility-scale photovoltaic units with frequency support functions and dynamic grid support for transmission systems,” *IET Renewable Power Generation*, vol. 11, pp. 361–372, February 2017.
- [10] J. Hernández, F. Sanchez-Sutil, P. Vidal, and C. Rus-Casas, “Primary frequency control and dynamic grid support for vehicle-to-grid in transmission systems,” *International Journal of Electrical Power & Energy Systems*, vol. 100, pp. 152–166, 2018.
- [11] C. Zhao, U. Topcu, and S. H. Low, “Optimal load control via frequency measurement and neighborhood area communication,” *Power Systems, IEEE Transactions on*, pp. 3576–3587, 2013.

- [12] C. Zhao, U. Topcu, N. Li, and S. Low, "Design and stability of load-side primary frequency control in power systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, 2014.
- [13] D. Wu, J. Lian, Y. Sun, T. Yang, and J. Hansen, "Hierarchical control framework for integrated coordination between distributed energy resources and demand response," *Electric Power Systems Research*, vol. 150, pp. 45–54, 2017.
- [14] J. Brooks and P. Barooah, "Distributed demand-side contingency-service provisioning while minimizing consumer disutility through local frequency measurements and inter-load communication," *arXiv preprint arXiv:1704.04586*, 2017.
- [15] E. Mallada, C. Zhao, and S. Low, "Optimal load-side control for frequency regulation in smart grids," *IEEE Transactions on Automatic Control*, vol. 62, no. 12, pp. 6294–6309, Dec 2017.
- [16] J. Brooks, R. D. Trevizan, P. Barooah, and A. S. Bretas, "Performance assessment of an optimal load control algorithm for providing contingency service," in *2017 North American Power Symposium (NAPS)*, Morgantown, WV, Sep 2017, pp. 1–6.
- [17] G. Barbose, C. Goldman, and B. Neenan, "A survey of utility experience with real time pricing," *Lawrence Berkeley National Laboratory*, 2004.
- [18] B. Biegel, L. Hansen, P. Andersen, and J. Stoustrup, "Primary control by ON/OFF demand-side devices," *IEEE Trans. on Smart Grid*, vol. 4, pp. 2061–2071, Dec 2013.
- [19] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental evaluation of frequency regulation from commercial building HVAC systems," *IEEE Transactions on Smart Grid*, vol. 6, pp. 776–783, 2015.
- [20] D. Todd, M. Caufield, B. Helms, A. P. Generating, I. M. Starke, B. Kirby, and J. Kueck, "Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of ALCOA INC.," *ORNL/TM*, vol. 233, 2008.
- [21] J. Brooks and P. Barooah, "Consumer-aware load control to provide contingency reserves using frequency measurements and inter-load communication," in *American Control Conference*, July 2016, pp. 5008–5013.
- [22] P. K. Kitanidis, "Unbiased minimum-variance linear state estimation," *Automatica*, vol. 23, no. 6, pp. 775–778, 1987.
- [23] V. S. Borkar and S. P. Meyn, "The ode method for convergence of stochastic approximation and reinforcement learning," *SIAM Journal on Control and Optimization*, vol. 38, no. 2, pp. 447–469, 2000.
- [24] C. Godsil and G. Royle, *Algebraic Graph Theory*, ser. Graduate Texts in Mathematics. Springer, 2001.
- [25] H. Kushner and G. G. Yin, *Stochastic approximation and recursive algorithms and applications*, 2nd ed., 2003, vol. 35.
- [26] V. S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.
- [27] R. S. Varga, *Geršgorin and his circles*. Springer Science & Business Media, 2010, vol. 36.
- [28] A. Moeini, I. Kamwa, P. Brunelle, and G. Sybille, "Open data iee test systems implemented in SimPowerSystems for education and research in power grid dynamics and control," in *Power Engineering Conference (UPEC), 2015 50th International Universities*. IEEE, 2015, pp. 1–6.
- [29] Y.-J. Kim, E. Fuentes, and L. K. Norford, "Experimental study of grid frequency regulation ancillary service of a variable speed heat pump," *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 3090–3099, 2016.