

Distributed demand-side contingency-service provisioning while minimizing consumer disutility through local frequency measurements and inter-load communication

Jonathan Brooks and Prabir Barooah

Abstract—We consider the problem of smart and flexible loads providing contingency reserves to the electric grid and provide a Distributed Gradient Projection (DGP) algorithm to minimize loads’ disutility while providing contingency services. Each load uses locally obtained grid-frequency measurements and inter-load communication to coordinate their actions, and the privacy of each load is preserved: only gradient information is exchanged—not disutility or consumption information. We provide a proof of convergence of the proposed DGP algorithm, and we compare its performance through simulations to that of a “dual algorithm” previously proposed in the literature that solved the dual optimization problem. The DGP algorithm solves the primal problem. Its main advantage over the dual algorithm is that it is applicable to convex—but not necessarily strictly convex—consumer disutility functions, such as a model of consumer behavior that is insensitive to small changes in consumption, while the dual algorithm is not. Simulations show the DGP algorithm aids in arresting grid-frequency deviations in response to contingency events and performs better or similarly to the dual algorithm in cases where the two can be compared.

I. INTRODUCTION

For stable and reliable operation of the power grid, generation must match consumption at all timescales [9]. Traditionally, this is achieved through controllable generators that adjust their outputs to track consumption. These services provided by generators are called ancillary services. With the increasing penetration of volatile renewable energies into the power grid, more resources are required to provide these ancillary services. One form of ancillary services is contingency reserves, which are used to restore the consumption-generation balance after a sudden change (e.g., a large generator going off-line). Conventional fossil-fuel generators are often operated at part-load in order to provide spinning reserves (fast-acting contingency reserves). However, rapidly ramping and operating at part-load can result in increased emission rates due to inefficient operation [12]. Part-loading requires additional generators to supply the needs of the grid as well. Building additional fossil-fuel generators to mitigate renewable volatility will reduce the environmental benefits of the renewable energies.

It has been recognized in recent years that an attractive alternative exists: loads can be used to provide spinning reserves by changing their consumption without increasing emissions [16],

[23]. Due to the size of the grid, centralized solutions to the load control problem are not practical. A distributed solution is more attractive and is possible by utilizing the cyber-physical nature of the electric grid whereby “information can be transmitted through actuation and sensing” [1]. In particular, loads can provide primary control by using local frequency measurements [20], [15]. The value of information contained in frequency measurements has been recognized much earlier [21]. Recent work in this vein includes [17], [22] in which loads are turned on or off based on frequency measurements. Information from frequency measurements allows solutions not generally considered in the literature of distributed optimization (e.g., [19], [28], [29]).

Any changes in consumption to help the grid, however, may incur some cost or disutility for the consumer—such as deviation of the indoor temperature from a comfortable range. Thus, there is a need to balance the two—service to the grid and cost to the consumer. In this paper, we consider the problem of designing decision-making algorithms that provide spinning reserves through control of loads while striking this balance.

In this paper, we adopt the problem formulation from [27]: minimize total consumer disutility while returning the consumption-generation mismatch in the grid to zero after a sudden change in generation. The consumption-generation mismatch is estimated by each load from noisy local frequency measurements using a state estimator. The algorithm in [27] is based on solving the dual optimization problem, and we refer to it as the “dual algorithm”. The dual variable, which is constant across the grid, is iteratively estimated using consensus averaging through inter-node communication.

The dual algorithm proposed in [27] requires the consumers’ disutilities to be strictly convex functions of changes in consumption. Quantifying consumers’ disutility in response to consumption changes is challenging, and work in this area is limited. In [7], an exponential function is used to model disutility, while [26] proposes a dynamic disutility model. A study of an industrial aluminum-smelting plant suggests that there may be no disutility for several hours when changing consumption within some threshold of a nominal value, but there *is* significant disutility if consumption is varied too much or for too long [25]. Likewise, [13] showed that consumption in commercial air-conditioning loads can be varied to provide ancillary services without any disutility (adverse effect on

J. Brooks and P. Barooah are with the Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL, 32611 USA. Corresponding author J. Brooks. e-mail: JonathanBrooksUF@gmail.com.

indoor climate) as long as the changes in consumption are small in amplitude and bandwidth-limited. Based on these studies, we hypothesize that an appropriate model of disutility for many consumers is like the function, f_1 , shown in Figure 1. The disutility is zero for small changes in consumption but non-zero disutility for larger changes. Such a consumer’s disutility is modeled by a convex—not strictly convex—function of consumption change.

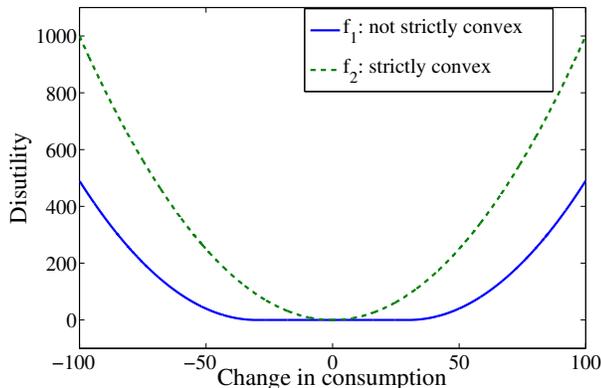


Fig. 1. Two distinct models of consumer disutility.

This work is an extension of our prior work [3], where we proposed a method to solve the *primal* problem in a distributed manner while preserving privacy, which we called the Distributed Gradient Projection (DGP) algorithm. The main advantage of the DGP algorithm over the dual algorithm of [27] is that the DGP algorithm is applicable to disutility functions that are convex but not necessarily strictly convex. The dual algorithm in [27] required a strictly convex disutility function because the inverse image of the function’s gradient is used in the computation, but the inverse image does not exist if there is a flat region in the disutility function (such as in f_1 in Figure 1).

This paper makes several contributions over our preliminary work [3], in which the convergence proof was limited to the case where there are no upper and lower bounds on how much a load can change its consumption. That scenario is unrealistic and was adopted for tractability of analysis. In this work, we remove this assumption and prove convergence in the almost-sure sense as well as global asymptotic stability. The analysis of convergence in a bounded domain is substantially more involved because of the discontinuous nature of the resulting dynamics. The analysis uses the o.d.e. method of stochastic approximation [2], [11]. Connection to the so called Skorokhod problem helps resolve existence and uniqueness issues of the resulting o.d.e, but proving convergence to the desired set—solutions to the optimization problem—required novel analysis. We also present an example of when the algorithm does not converge—when the assumption of strict feasibility of the optimal solutions is violated, showing that the convergence result obtained here is the strongest one possible for the DGP algorithm.

Simulations indicate that the DGP algorithm is effective

in reducing frequency excursions following step changes in generation. Simulation comparisons, in those scenarios where comparison is possible, show that the proposed DGP algorithm performs better than or comparably to the dual algorithm of [27]. On the other hand, our analysis is limited to time-invariant communication, whereas convergence for the dual algorithm was proved to hold even for time-varying communication [27].

This paper is organized as follows. Section II formally defines the problem that we solve. In Section III, we describe the DGP algorithm. We present the convergence result in Section IV and proofs in Section IV-A. Simulations are described in Section V. Finally, Section VI concludes this work and discusses avenues for future extensions.

II. PROBLEM FORMULATION

As in [27], we consider an electric grid with a single frequency throughout the grid, whose nominal value is denoted by ω^* , such as in a microgrid. There are n controllable loads. The deviation of load i ’s consumption from its nominal value is denoted by x_i and incurs a disutility, $f_i(x_i)$. The deviation must lie in $\Omega_i \triangleq [x_i, \bar{x}_i]$, specified a-priori.

Let Δg be the generation deviation from the nominal value. The problem is for the loads to decide how much to change their own consumption so that the consumption-generation mismatch is diminished while the resulting total disutility of the loads is minimized:

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

Load i can obtain a noisy measurement of the grid frequency and can use it to make a decision on x_i . In addition, the computation of the decision variables, x_i , must be distributed in the following sense. There is a communication graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the node set, $\mathcal{V} = \{1, 2, \dots, n\}$, is simply the loads and the edge set, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$, specified a-priori, determines which pairs of loads can exchange information. The set of neighbors of load i , with which it can exchange information, is defined by $\mathcal{N}_i \triangleq \{j \mid (i, j) \in \mathcal{E}\}$, and we denote $n_i \triangleq |\mathcal{N}_i|$ as the number of neighbors of load i . The frequency measurements are essential since every load can use them to estimate the equality-constraint violation, $u \triangleq \Delta g - \sum_{i=1}^n x_i$. How this is done is described in Section III-A.

Although Problem (1) does not include time, the proposed algorithm is an iterative approach, so time does play a role. The reason for this is twofold: i) due to the limited information available to the loads, they are unable to find the solution in just one step, so they move “toward” the solution and then re-evaluate at the new location ii) due to generator dynamics, the frequency, ω , is dependent on the loads, so as the loads change in time, so does the frequency, which the loads can then use for feedback control. Time is measured by a discrete iteration counter: $k = 0, 1, \dots$. The generation at time k is denoted by $g[k]$ so that the generation change from nominal is $\Delta g[k] \triangleq g[k] - g^*$, where g^* is the nominal generation. We assume that, at $k = 0$, total load and total generation are

equal, and we limit ourselves to step changes in generation. That is, $\Delta g[k] = 0$ for $k < K$ for some K , and $\Delta g[k] = \bar{g}$ for $k \geq K$, where \bar{g} is the size of the step change.

III. DISTRIBUTED GRADIENT PROJECTION (DGP) ALGORITHM

To describe the algorithm, we define the consumption-generation mismatch at iteration k :

$$u[k] \triangleq \Delta g[k] - \sum_{i=1}^n x_i[k] = \Delta g[k] - \mathbf{1}^T \mathbf{x}[k], \quad (2)$$

where $\mathbf{x}[k] \triangleq [x_1[k], \dots, x_n[k]]^T$ and $\mathbf{1} \in \mathbb{R}^n$ is a vector of all ones. The frequency deviation from its nominal value is denoted by $\Delta \omega[k] \triangleq \omega[k] - \omega^*$. Neither $\sum_{i=1}^n x_i[k]$ nor $\Delta g[k]$ is known to any of the loads. However, load i can obtain a noisy measurement, $\Delta \tilde{\omega}_i[k]$, of the frequency deviation. It uses this measurement to estimate the mismatch, which is denoted by $\hat{u}_i[k]$,

The update law of the DGP algorithm comprises of 3 main operations: (i) a generation-matching step, (ii) a gradient descent step, and (iii) a projection step. The first step uses the estimated mismatch, $\hat{u}_i[k]$, to compute a consumption change that will reduce the mismatch. For the second step, pure gradient descent, though possible, will violate the equality constraint (consumption-generation matching). Therefore the gradient descent step is designed to be orthogonal to the generation-matching step; i.e., it does not change the total consumption. The updates computed by the first two steps are added, and their sum is then projected onto Ω_i to respect the upper and lower bounds on consumption change.

The update law of the DGP algorithm at load i at time k is summarized below:

DGP Algorithm:

- 1) Obtain an estimate of $u_i[k]$ (call it $\hat{u}_i[k]$) from the measurement, $\Delta \tilde{\omega}_i[k]$, using a state estimator (described in Section III-A).
- 2) Compute gradient $\frac{d}{dx_i} f_i(x_i[k])$, transmit gradient value to neighbors, receive neighbors' gradient values, and compute

$$\Delta x_i[k] := -n_i \nabla f_i(x_i[k]) + \sum_{j \in \mathcal{N}_i} \nabla f_j(x_j[k]). \quad (3)$$

- 3) Update load change as $x_i[k+1] = P_{\Omega_i}[x_i[k] + \alpha[k] \Delta x_i[k] + \gamma[k] \hat{u}_i[k]]$, where $P_{\Omega_i}[\cdot]$ denotes the standard projection operator and $\alpha[k], \gamma[k]$'s are step sizes.

Note that loads only exchange gradient information, so each load's disutility and consumption information remain private and are not shared with other loads.

A. Estimation of consumption-generation mismatch using frequency measurements

We borrow the estimation method proposed in [27] for estimating u , although it is possible to use any estimator in the DGP algorithm. The power grid is modeled as a discrete-time LTI system with consumption-generation mismatch, $u[k]$, as

the input and frequency deviation from nominal, $\Delta \omega[k]$, as the output. At each time k , load i obtains the noisy measurement, $\Delta \tilde{\omega}_i[k]$, to estimate the state of the plant by using the estimator in [10], which was developed for estimating the state of a system with an unknown input. Once the state estimate is obtained, each load estimates the unknown input by effectively assuming that the most recent output is error-free and then solving for the previous input from the state equations.

We denote the estimation error at time k by

$$\epsilon[k] \triangleq \hat{u}[k] - u[k] \mathbf{1}, \quad (4)$$

where $\hat{u}[k]$ is the column vector of $\hat{u}_i[k]$'s. Define the σ -algebra, $\mathcal{F}[k-1] := \sigma(\epsilon_i[\ell-1] \mid i \in \mathcal{V}; 1 \leq \ell \leq k)$. It was shown in [10] that

$$\mathbb{E}[\epsilon_i[k] | \mathcal{F}[k-1]] = 0. \quad (5)$$

The following proposition is reproduced from [27].

Proposition 1 ([27]). *Let A , B , and C denote the process, input, and output matrices of the state-space model of the power grid used in the state estimator described in [10], and let I_A be an identity matrix the size of A . If every eigenvalue of $(I_A - B(CB)^{-1}C)A$ lies within the unit circle, then $\lim_{k \rightarrow \infty} \mathbb{E}[(\epsilon_i[k])^2 | \mathcal{F}[k-1]]$ exists.*

The following corollary is a straightforward consequence of (5) and Proposition 1 based on the definition of a martingale-difference sequence.

Corollary 1. *If the condition for Proposition 1 holds, then (5) and Proposition 1 imply that the estimation error sequence, $\epsilon[k]$, is a martingale-difference sequence.*

IV. CONVERGENCE ANALYSIS

We make the following assumptions for our analysis.

Assumption 1. (Assumptions on disutility).

- 1) $f_i(x_i)$ is convex for each i with a (not necessarily unique) minimum at $x_i = 0$.
- 2) $f_i(x_i)$ is coercive for each i ; i.e., $\{x_i | f_i(x_i) \leq F\}$ is compact for every $F \geq 0$ for each i .
- 3) $f_i(x_i)$ is continuously differentiable for each i .
- 4) $\nabla f_i(x_i)$ is Lipschitz for each i .

Assumption 2. (Geometric assumptions).

- 1) The domain, Ω , is compact.
- 2) The communication graph, \mathcal{G} , is connected.
- 3) The disturbance is a constant: $\Delta g[k] \equiv \bar{g}$ for all $k \geq 0$.

Assumption 3. (Technical assumptions).

- 1) $\alpha[k] = c\gamma[k]$ for some positive constant c .
- 2) The function, $\gamma[k] \rightarrow 0$, satisfies $\sum_{k=0}^{\infty} \gamma[k] = \infty$ and $\sum_{k=0}^{\infty} (\gamma[k])^2 < \infty$.
- 3) The estimation error sequence, $\epsilon[k]$, is a martingale-difference sequence.

Assumption 1 is readily satisfied because the disutility functions are a modeling choice. Assumption 2(1) is always met in practice since a load cannot change its demand outside its maximum rated power and 0. Assumption 2(3) means that

we only consider a step-change in generation, which is a good approximation of a contingency event. Assumptions 3(1-2) are satisfied by choosing $\alpha[k]$ and $\gamma[k]$ appropriately. Assumptions 3(2-3) are standard in the field of stochastic approximation [2]. Note that Assumption 3(3) holds in this work due to Proposition 1.

The main convergence result is the following.

Theorem 1. *If Assumptions 1, 2, and 3 hold and all solutions to Problem (1) are strictly feasible, $\mathbf{x}[k]$ converges to a solution to Problem (1), almost surely (a.s.).*

The proof of this result relies on the so called o.d.e. method of stochastic approximation, which establishes a rigorous connection between noisy discrete iterations and a continuous-time o.d.e. [2], [11], stated next.

Proposition 2 (Theorem 2.1 (Chapter 5) in [11]). *Consider the sequence $\{\mathbf{y}[k]\}$ generated by the iteration*

$$\mathbf{y}[k+1] = P_\Omega \left[\mathbf{y}[k] + \gamma[k] (h(\mathbf{y}[k]) + \epsilon[k]) \right],$$

where P_Ω is the projection operator onto Ω , $h(\mathbf{y}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz, and $\{\epsilon[k]\}$ is a martingale-difference sequence. If $\gamma[k]$ satisfies Assumption 3(2) then $\mathbf{y}[k]$ converges a.s. to some limit set of the o.d.e.,

$$\dot{\mathbf{y}}(t) = \Gamma_{\Omega, \mathbf{y}(t)} [h(\mathbf{y}(t))],$$

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

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

By using (4) and Assumption 3, the update law of the DGP algorithm can be written as

$$\begin{aligned} x_i[k+1] &= P_{\Omega_i} \left[x_i[k] \right. \\ &\quad \left. + \gamma[k] \left(c \sum_{j \in \mathcal{N}_i} (\nabla f_j(x_j[k]) - \nabla f_i(x_i[k])) \right) \right. \\ &\quad \left. + u[k] + \epsilon_i[k] \right], \end{aligned} \quad (6)$$

or compactly,

$$\begin{aligned} \mathbf{x}[k+1] &= P_\Omega \left[\mathbf{x}[k] \right. \\ &\quad \left. + \gamma[k] (-cL\nabla f(\mathbf{x}[k])^T + u[k]\mathbf{1} + \epsilon[k]) \right], \end{aligned} \quad (7)$$

where $\nabla f(\mathbf{x})$ is the gradient of $f(\mathbf{x}) := \sum_{i=1}^n f_i(x_i)$ and $L \in \mathbb{R}^{n \times n}$ is the graph Laplacian of \mathcal{G} [6]: $L_{ii} = n_i$ (number of neighbors of node i), and for $j \neq i$, $L_{ij} = -1$ if $j \in \mathcal{N}_i$, $L_{ij} = 0$ if $j \notin \mathcal{N}_i$. By Proposition 2, the iterates, $\mathbf{x}[k]$, converge a.s. to a limit set of the o.d.e.,

$$\dot{\mathbf{x}}(t) = \Gamma_{\Omega, \mathbf{x}(t)} [-cL\nabla f(\mathbf{x}(t))^T + u(t)\mathbf{1}], \quad (8)$$

$$u(t) := \bar{g} - \mathbf{1}^T \mathbf{x}(t) \quad (9)$$

The task is to prove that the trajectories of o.d.e. (8) converge to the set of solutions to Problem (1), which we denote by X^* . Because of the convexity, $\mathbf{x}^* \in X^*$ if and only if it satisfies the first-order necessary conditions of optimality [14]. We denote the boundary of Ω by $\partial\Omega$, and we denote the interior of Ω by Ω° . Because of the hypothesis about strict feasibility in Theorem 1, only the equality constraint is active, which is also regular, and thus it is straightforward to verify that

$$X^* = \{ \mathbf{x} \in \Omega^\circ \mid \mathbf{1}^T \mathbf{x} = \bar{g}, \nabla f_i(x_i) = \nabla f_j(x_j), \forall i, j \in \mathcal{V} \}. \quad (10)$$

Because each f_i is convex, each gradient $\nabla f_i(x_i)$ is nondecreasing, and the optimal set, X^* , is connected. It follows that the gradient, $\nabla f(\mathbf{x})$, at any optimal point $\mathbf{x}^* \in X^*$ is unique, which can be seen via contradiction. Suppose $\mathbf{x}_1^*, \mathbf{x}_2^* \in X^*$ with $\|\nabla f(\mathbf{x}_1^*)\| < \|\nabla f(\mathbf{x}_2^*)\|$. Then $f(\mathbf{x}_1^*) < f(\mathbf{x}_2^*)$ by convexity, so \mathbf{x}_2^* is not optimal—a contradiction. Therefore, $\nabla f(\mathbf{x}^*)$ is unique for all $\mathbf{x}^* \in X^*$. We denote the corresponding scalar by ∇f^* :

$$\nabla f^* \triangleq \nabla f_i(x_i), \text{ where } \mathbf{x} \in X^* \text{ and } i \in \{1, \dots, n\}. \quad (11)$$

The constant, ∇f^* , is called the *optimal gradient* in the sequel.

Since all limit sets of o.d.e (8) are contained within its ω -limit set, which we call ω_L , it suffices to show that the optimal set, X^* , contains the ω -limit set: $\omega_L \subseteq X^*$. Theorem 1 then follows immediately by Proposition 2.

In many applications of the o.d.e. method, the main hurdle in analyzing convergence is to establish boundedness of the iterates, $\mathbf{x}[k]$ [2]. In our case, presence of the projection step guarantees boundedness trivially. However, this projection makes the limit sets of the o.d.e. difficult to characterize due to the discontinuous nature of projected dynamical systems [18]. In fact, existence and uniqueness of a solution to (8) needs to be established first. Because the right-hand side of (8) is discontinuous, existence of classical solutions to the o.d.e. (i.e., solutions that are continuously differentiable) is not guaranteed. For discontinuous dynamics, the notions of Caratheodory and Filippov solutions are used. Caratheodory solutions are essentially a generalization of classical solutions that are absolutely continuous and differentiable almost everywhere, while Filippov solutions are absolutely continuous maps that satisfy a differential inclusion almost everywhere [5]. When dealing with Filippov solutions, the way in which the differential inclusion is defined plays a crucial role in both existence and uniqueness. This can further complicate the analysis.

The o.d.e. (8) is closely related to the so called Skorokhod problem, which deals with stochastic differential equations with boundaries [24]. The connection with the Skorokhod problem and ours is useful; known results on the Skorokhod problem provide guarantees of not only existence but also uniqueness of Caratheodory solutions to o.d.e. (8) [18]. For completeness, the relevant result from [18] is presented below.

Proposition 3 (Theorem 2.5 in [18]). *Let $h(\cdot)$ be Lipschitz continuous and*

$$\dot{\mathbf{x}}(t) = \Gamma_{\Omega, \mathbf{x}(t)} [h(\mathbf{x}(t))], \quad (12)$$

where Γ is the projection defined in Proposition 2 and Ω is compact. For any $\mathbf{x}(0) \in \Omega$, there exists a unique

Caratheodory solution, $\mathbf{x}(t)$, to (12) starting from $\mathbf{x}(0)$. Furthermore, $\mathbf{x}(t)$ is continuous with respect to the initial condition, $\mathbf{x}(0)$.

The second challenge is analyzing limiting behavior of the o.d.e. The trajectory may “evolve[] along a ‘section’ of $[\partial\Omega]$... At a later time the solution may re-enter $[\Omega^0]$, or it may enter a lower[-]dimensional part of $[\partial\Omega]$.” [18]. The trajectory may go on doing so without ever converging to a limit set with tractable structure. This is a severe hurdle in analyzing dynamical systems evolving in a bounded region, where the boundedness is enforced through a projection operator. In [18], this hurdle was sidestepped by *assuming* that the ω -limit set is contained in the set of fixed points. That is equivalent to assuming that complicated limit sets such as limit cycles do not arise. The reference [2] also mentions that, in presence of the projection operator onto domains with non-smooth boundaries (such as in our case), existence of an o.d.e. limit of the discrete-time algorithm may be a non-trivial issue. The classic reference [11] on stochastic approximation and o.d.e. methods mentions that basic results on convergence to a limit set (such as Proposition 2) is rarely useful since the limit set can be the whole of the bounded domain in which the o.d.e. evolves. This was also the main argument in [18] that is mentioned above. Only in the special case when the right-hand side of the o.d.e. is a gradient-descent system can convergence to a set of fixed points be established [11]. However, (8) is not a gradient-descent system.

Our main convergence result, Theorem 1, is obtained without making the strong assumptions that are typical, such as assuming that no limit cycles exist. The structure of the Laplacian matrix plays an important role in the proofs of the main theorem and the technical results needed for the theorem proof. A key technical result is that the equality constraint is eventually satisfied, which is stated in the next lemma.

Lemma 1. *If Assumptions 1, 2, and 3 hold and all solutions to Problem (1) are strictly feasible, $u(t) \rightarrow 0$, where $u(t)$ is defined in (9) and $\mathbf{x}(t)$ —in the definition of $u(t)$ —is governed by o.d.e. (8).*

A similar result was obtained in [4]. However, the analysis in that work uses a penalty method in place of projection. Consequently, the result of Lemma 1 is obtained immediately. In our case, where we retain the projection operation, the proof of Lemma 1 is not as straightforward. We provide our proof of the lemma in Section IV-A. We are now ready to prove Theorem 1.

Proof of Theorem 1: Recall the discussion after Proposition 2: we must prove that the trajectories of o.d.e. (8) converge to the set, X^* . By Proposition 3, the solution to o.d.e. (8) is continuous with respect to the initial condition, so ω_L is an invariant set [8]. Our goal is to use this invariance property to characterize ω_L and show that $\omega_L \subseteq X^*$. For convenience, we write the i^{th} element of (8) below:

$$\dot{x}_i(t) = \Gamma_{\Omega_i, x_i(t)} \left[c \sum_{j \in \mathcal{N}_i} \left(\nabla f_j(x_j(t)) - \nabla f_i(x_i(t)) \right) + u(t) \right]. \quad (13)$$

By Lemma 1, $\mathbf{x}(t) \rightarrow \{\mathbf{x} \in \Omega \mid \mathbf{1}^T \mathbf{x} = \bar{g}\} =: U$; this implies that all limit sets are contained within U . Because the ω -limit set is the union of all limit points, we have $\omega_L \subseteq U$. Let

$$\underline{\nabla} f(t) \triangleq \min_i \nabla f_i(x_i(t)). \quad (14)$$

Imagine a trajectory starting in ω_L : $\mathbf{x}(0) \in \omega_L \subseteq U$, which implies $\mathbf{x}(t) \in U$ for all t by invariance of ω_L . Since $u(t) \equiv 0$ for this trajectory, o.d.e. (8) reduces to $\dot{\mathbf{x}}(t) = \Gamma_{\Omega, \mathbf{x}(t)} [-cL \nabla f(\mathbf{x}(t))^T]$. That is, for every $i \in \mathcal{V}$, $\dot{x}_i(t)$ is the average of its neighbors’ gradients at t , which shows that $\dot{x}_{r(t)} \geq 0$, where $r(t) := \arg \min_i \nabla f_i(x_i(t))$. Therefore, $\underline{\nabla} f(t) = \nabla f_{r(t)}(x_{r(t)}(t))$ is nondecreasing by convexity. (Note that this does not require uniqueness of $r(t)$, so when two gradients are both minimal, both are nondecreasing.) Because $\underline{\nabla} f(t)$ is bounded as well, which comes from boundedness of the domain, Ω , $\underline{\nabla} f(t)$ converges. Denote the limit by F so that trajectories starting in ω_L converge to $\mathcal{F} := \{\mathbf{x} \in \Omega \mid \min_i \nabla f_i(x_i) = F; \mathbf{1}^T \mathbf{x} = \bar{g}\} \subseteq \omega_L \subseteq U$. Therefore, for any $\mathbf{x}(0) \in \omega_L$, $\mathbf{x}(t) \rightarrow \mathcal{F}$.

We now show that $\mathbf{x}(t) \rightarrow \mathcal{F}$ for any $\mathbf{x}(0) \in \Omega$. Suppose we have some trajectory, $\mathbf{x}(t)$ with $\mathbf{x}(0) = \mathbf{x}_o \in \Omega$. From the definition of the ω -limit set, there exists a point, $\mathbf{y} \in \omega_L$, and a sequence, t_0, t_1, \dots with $t_n \rightarrow \infty$ as $n \rightarrow \infty$, such that $\mathbf{x}(t_n) \rightarrow \mathbf{y}$ as $n \rightarrow \infty$. Let $\mathbf{x}_n(t) := \mathbf{x}(t_n + t)$ for every t so that $\mathbf{x}_n(0) = \mathbf{x}(t_n)$. Let \mathbf{x}_{ω_L} be a trajectory that starts from \mathbf{y} : $\mathbf{x}_{\omega_L}(0) = \mathbf{y} \in \omega_L$. By construction $\lim_{n \rightarrow \infty} \mathbf{x}_n(0) = \mathbf{y} = \mathbf{x}_{\omega_L}(0)$. Therefore, for any $T \in [0, \infty)$, $\mathbf{x}_n(T) \rightarrow \mathbf{x}_{\omega_L}(T)$ as $n \rightarrow \infty$ by continuity of $\mathbf{x}(t)$ with respect to $\mathbf{x}(0)$ (see Proposition 3). Now, we have $\lim_{n \rightarrow \infty} \mathbf{x}(t_n + T) = \lim_{n \rightarrow \infty} \mathbf{x}_n(T) = \mathbf{x}_{\omega_L}(T)$. This is true for any $T \geq 0$, so we may take the limit as $T \rightarrow \infty$ to obtain $\mathbf{x}(t) \rightarrow \mathcal{F}$ because $\mathbf{x}_{\omega_L}(T) \rightarrow \mathcal{F}$ as $T \rightarrow \infty$. Therefore, $\mathbf{x}(t) \rightarrow \mathcal{F}$ for any $\mathbf{x}(0) \in \Omega$, so $\omega_L \subseteq \mathcal{F}$. Since $\mathcal{F} \subseteq \omega_L$ also, we have $\mathcal{F} = \omega_L$. Therefore, \mathcal{F} is invariant by invariance of ω_L .

Next, we show that \mathcal{F} is invariant only if $F = \nabla f^*$. Consider again a trajectory starting in \mathcal{F} : $\mathbf{x}(t_0) \in \mathcal{F}$ for some t_0 . If $\nabla f_i(x_i(t_0)) = F$ for every $i \in \mathcal{V}$, then the first-order necessary conditions for optimality are satisfied, so $F = \nabla f^*$, and $\mathcal{F} = X^*$, and the proof is complete. Suppose there exists an $\ell \in \mathcal{V}$ such that $\nabla f_\ell(x_\ell(t_0)) > F$. Let $q := \arg \min_i \nabla f_i(x_i(t_0))$; then $\dot{x}_q(t_0) \geq 0$ with equality if and only if $\nabla f_j(x_j(t_0)) = F$ for all $j \in \mathcal{N}_q$, in which case q is not unique. Denote the different choices for q by q_i , $i \in \mathcal{V}$. Because $\nabla f_\ell(x_\ell(t_0)) > F$ and the graph \mathcal{G} is connected, $\dot{x}_{q_j}(t_0) > \epsilon$ for some q_j and some $\epsilon > 0$. It follows that there exists some $t_1 > t_0$ such that $\nabla f_{q_j}(x_{q_j}(t_1)) > F$. Note that $u(t) = 0$ implies $\nabla f_{q_j}(x_{q_j}(t)) > F$ for all $t \geq t_1$; this is because $\dot{x}_{q_j}(t) \rightarrow \mathbb{R}_{\geq 0}$ as $\nabla f_{q_j}(x_{q_j}(t)) \rightarrow F$ (see (13)). Hence, $\dot{x}_{q_i}(t_1) > \delta$ for some $q_i \in \mathcal{N}_{q_j}$ and some $\delta > 0$. By continuing this argument, we see that eventually there exists some $t_2 > t_0$ such that $\underline{\nabla} f(t_2) > F$; that is, when the gradient of one q_j increases, that causes the gradient of all $q_i \in \mathcal{N}_{q_j}$ to eventually increase, and so on. However, $\underline{\nabla} f(t_2) > F$ implies $\mathbf{x}(t_2) \notin \mathcal{F}$, so \mathcal{F} is not invariant—a contradiction. Therefore, if $\mathbf{x}(t_0) \in \mathcal{F}$, there exists no $\ell \in \mathcal{V}$ such that $\nabla f_\ell(x_\ell(t_0)) > F$. Hence, $\nabla f_i(x_i(t_0)) = F$ for all $i \in \mathcal{V}$,

which satisfies the optimality conditions, so $F = \nabla f^*$, and $\omega_L = \mathcal{F} = X^*$, which completes the proof. ■

The following lemma shows that the optimal solution set, X^* , is not only globally attractive, but also globally asymptotically stable. The notion of stability of an invariant set used here is the one from [8].

Lemma 2. *Let the conditions of Theorem 1 hold. Then the optimal set, X^* , is globally asymptotically stable.*

The proof of this lemma is provided in Section IV-A.

Comment on strict feasibility: The assumption of strict feasibility of all solutions to Problem 1 is necessary for the convergence of the DGP algorithm to hold. If X^* is not strictly feasible, then the DGP algorithm is not guaranteed to converge to X^* ; in fact it may converge to non-optimal points. This may be seen through the following 2D counterexample. Let $\mathbf{x} \in \mathbb{R}^2$, $\bar{g} = 1$, $f_1(x_1) = (x_1)^2$, $f_2(x_2) = (x_2)^2$, $\Omega = [0, 1/4] \times [0, 1]$, and $c = 1$. \mathcal{G} consists of the two nodes and one edge connecting them. The constraint, $u = 0$, is satisfied on the line $x_1 + x_2 = 1$. It is straightforward to verify that $\mathbf{x}^* = [1/4, 3/4]^T$ is the unique solution to Problem (1), but \mathbf{x}^* is not strictly feasible since it lies on the boundary of Ω . Now, from (8),

$$\begin{aligned} \dot{x}_2(t)|_{\mathbf{x}^*} &= \Gamma_{\Omega_2, x_2(t)}[-2x_2(t) + 2x_1(t) + u(t)]|_{\mathbf{x}^*} \\ &= \Gamma_{\Omega_2, 0.75}[-2 \cdot 0.75 + 2 \cdot 0.25 + 0] \\ &= -1, \end{aligned}$$

which shows that $[1/4, 3/4]^T$ is not an equilibrium point of o.d.e. (8). Therefore, the DGP algorithm is not guaranteed to converge to X^* . It can also be shown that for this example $[1/4, 5/12]^T$ is an equilibrium point of o.d.e. (8) and that it is attractive, but $[1/4, 5/12]^T$ is not a solution to Problem (1). Therefore, from Proposition 2, the iterates will converge almost surely to a non-optimal point.

A. Proofs

Proof of Lemma 1: First, we introduce a few definitions. Define the *projection-less* derivative at time t for load i , $\dot{p}_i(t)$, such that $\dot{x}_i(t) = \Gamma_{\Omega_i, x_i(t)}[\dot{p}_i(t)]$. That is,

$$\dot{p}_i(t) \triangleq c \sum_{j \in \mathcal{N}_i} (\nabla f_j(x_j(t)) - \nabla f_i(x_i(t))) + u(t). \quad (15)$$

Let

$$X_i^* := \{m \in \mathbb{R} \mid \exists \mathbf{x} \in X^* \text{ s. t. } \mathbf{x}_i = m\}. \quad (16)$$

Figure 2 shows an illustration for a 2-D case. Note that $x_i \in X_i^*$ for all i does not imply $\mathbf{x} \in X^*$ because it may not satisfy the equality constraint, $\mathbf{1}^T \mathbf{x} = \bar{g}$. However, the sets, X_i^* , satisfy a property that is useful in the proofs: if $\mathbf{x} \in X_1^* \times X_2^* \times \dots \times X_n^*$, then $\nabla f_i(x_i) = \nabla f^*$ for every i ; where ∇f^* is the optimal gradient; see (11). It should be noted, however, that the converse is true only if f_i is strictly convex.

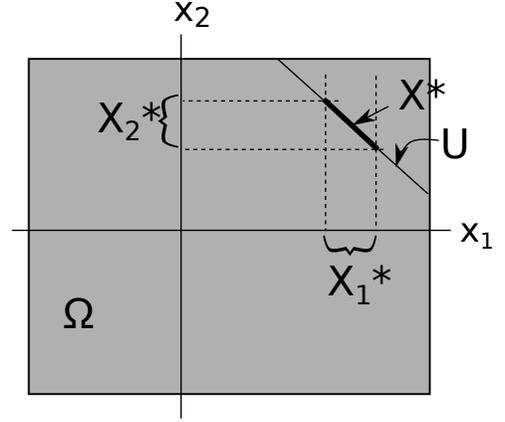


Fig. 2. An illustration of the relevant sets for a 2-D case. The equality constraint is satisfied on the line segment, which is the set U . The thick sub-segment is the optimal solution set, X^* .

For this reason, we shall call X_i^* the *critical gradient set* of load i . Let us define the collections of loads,

$$\begin{aligned} \mathcal{A}(t) &:= \{i \in \mathcal{V} \mid x_i(t) > X_i^*\}, \\ \mathcal{B}(t) &:= \{i \in \mathcal{V} \mid x_i(t) < X_i^*\}, \\ \mathcal{M}(t) &:= \{i \in \mathcal{V} \mid x_i(t) \in X_i^*\}, \end{aligned} \quad (17)$$

At t , $\mathcal{A}(t)$ is the set of loads above their critical gradient sets, $\mathcal{B}(t)$ is the set of loads below their critical gradient sets, and $\mathcal{M}(t)$ is the set of loads at their critical gradient sets. By definition, $\mathcal{A}(t)$, $\mathcal{B}(t)$, and $\mathcal{M}(t)$ are mutually disjoint. Define also $a(t) := |\mathcal{A}(t)|$ and $b(t) := |\mathcal{B}(t)|$; note that $a(t) + b(t) \leq n$ for all t . Additionally, let

$$\begin{aligned} \mathcal{S}(t) &:= \{i \in \mathcal{V} \mid \dot{p}_i(t) \neq \dot{x}_i(t)\}, \\ \mathcal{S}_A(t) &:= \{i \in \mathcal{A} \mid \dot{p}_i(t) > 0 = \dot{x}_i(t)\} \subset \mathcal{S}(t), \\ \mathcal{S}_B(t) &:= \{i \in \mathcal{B} \mid \dot{p}_i(t) < 0 = \dot{x}_i(t)\} \subset \mathcal{S}(t). \end{aligned} \quad (18)$$

That is, $\mathcal{S}(t)$ is the set of loads for which the projection operation is currently active, $\mathcal{S}_A(t)$ is the set of loads in $\mathcal{A}(t)$ currently constrained at their upper bounds due to projection, and $\mathcal{S}_B(t)$ is the set of loads in $\mathcal{B}(t)$ constrained at their lower bounds due to projection. Recall that Ω° is the interior of Ω , and note that $X_i^* \subset \Omega_i^\circ$ because $X^* \subset \Omega^\circ$; therefore, projection does not affect loads in $\mathcal{M}(t)$. It follows that $\mathcal{S}(t) = \mathcal{S}_A(t) \cup \mathcal{S}_B(t)$. Likewise, we have $\mathcal{S}_A(t) \subset \mathcal{A}(t)$ and $\mathcal{S}_B(t) \subset \mathcal{B}(t)$. Now, we may observe

$$\begin{aligned} \dot{u}(t) &= -\mathbf{1}^T \dot{\mathbf{x}}(t) \\ &= -\mathbf{1}^T \dot{\mathbf{p}}(t) + \sum_{i \in \mathcal{S}(t)} \dot{p}_i(t) \\ &= -nu(t) + \sum_{i \in \mathcal{S}_A(t)} \dot{p}_i(t) + \sum_{i \in \mathcal{S}_B(t)} \dot{p}_i(t), \end{aligned} \quad (19)$$

where the final equality follows from $\mathcal{S}(t) = \mathcal{S}_A(t) \cup \mathcal{S}_B(t)$.

Let the function, $y_i : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, $i = 1, \dots, n$, and the function, $y : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$, be defined as

$$\begin{aligned} y_i(m) &\triangleq \inf_r \{|m - r| : r \in X_i^*\}, \quad i \in \mathcal{V}, \\ y(\mathbf{x}) &\triangleq \sum_{i \in \mathcal{V}} y_i(x_i). \end{aligned} \quad (20)$$

At time t , $y_i(x_i(t))$ is the Euclidean distance between the component, $x_i(t)$, and the critical gradient set, X_i^* , and $y(t)$ is the total distance of all loads from their respective critical gradient sets. For the sake of compactness, we will write $y_i(t)$ and $y(t)$ in place of $y_i(x_i(t))$ and $y(x(t))$ in the sequel. Next, we more closely examine the term, $-L\nabla f(x(t))$, in (8). Without loss of generality, let $\nabla f_1(x_1(t)) \leq \nabla f_2(x_2(t)) \leq \dots \leq \nabla f_n(x_n(t))$.

$$\begin{aligned} C_B(t) &\triangleq c \sum_{i \in \mathcal{B}(t)} \sum_{j \in \mathcal{N}_i} (\nabla f_j(x_j(t)) - \nabla f_i(x_i(t))) \\ &= c \sum_{i=1}^{b(t)} \sum_{j \in \mathcal{N}_i} (\nabla f_j(x_j(t)) - \nabla f_i(x_i(t))). \end{aligned} \quad (21)$$

For $i = 1$, every term in (21) is nonnegative. Suppose $1 \notin \mathcal{N}_2$; then for $i = 2$, every term in (21) is nonnegative. Now suppose $1 \in \mathcal{N}_2$; then $2 \in \mathcal{N}_1$, and the respective terms in (21) cancel for $i = 1, j = 2$ and $i = 2, j = 1$. Therefore, the terms of (21) for $i = 1, 2$ are nonnegative regardless of whether 1 and 2 are neighbors. This argument may be continued for all $i \in \mathcal{B}(t)$. It follows that $C_B(t) \geq 0$.

We may define a similar sum over $\mathcal{A}(t)$:

$$C_A(t) \triangleq c \sum_{i \in \mathcal{A}(t)} \sum_{j \in \mathcal{N}_i} (\nabla f_j(x_j(t)) - \nabla f_i(x_i(t))). \quad (22)$$

By the same argument as above, we may show that $C_A(t) \leq 0$.

Combining (15) and (21), we can see that the total change in consumption for loads in $\mathcal{B}(t)$ without projection is

$$\sum_{i \in \mathcal{B}(t)} \dot{p}_i(t) = C_B(t) + b(t)u(t).$$

Similarly from (15) and (22), the total change in consumption for loads in $\mathcal{A}(t)$ without projection is

$$\sum_{i \in \mathcal{A}(t)} \dot{p}_i(t) = C_A(t) + a(t)u(t).$$

Note that

$$\begin{aligned} \sum_{i \in \mathcal{B}(t)} \dot{x}(t) &= \sum_{i \in \mathcal{B}(t)} (\dot{p}_i(t)) - \sum_{i \in \mathcal{S}_B(t)} (\dot{p}_i(t)) \\ &= C_B(t) + b(t)u(t) - \sum_{i \in \mathcal{S}_B(t)} \dot{p}_i(t) \end{aligned}$$

and

$$\begin{aligned} \sum_{i \in \mathcal{A}(t)} \dot{x}(t) &= \sum_{i \in \mathcal{A}(t)} (\dot{p}_i(t)) - \sum_{i \in \mathcal{S}_A(t)} (\dot{p}_i(t)) \\ &= C_A(t) + a(t)u(t) - \sum_{i \in \mathcal{S}_A(t)} \dot{p}_i(t). \end{aligned}$$

Because $x(t)$ is a Caratheodory solution (see Proposition 3), it is absolutely continuous and differentiable almost everywhere. It follows that $y(t)$ is absolutely continuous and differentiable

almost everywhere. Then we have

$$\begin{aligned} \dot{y}(t) &= \sum_{i \in \mathcal{A}(t)} (\dot{x}_i(t)) - \sum_{i \in \mathcal{B}(t)} (\dot{x}_i(t)) \\ &= (C_A(t) + a(t)u(t) - \sum_{i \in \mathcal{S}_A(t)} \dot{p}_i(t)) \\ &\quad - (C_B(t) + b(t)u(t) - \sum_{i \in \mathcal{S}_B(t)} \dot{p}_i(t)) \\ &= C_A(t) - C_B(t) + (a(t) - b(t))u(t) \\ &\quad - \sum_{i \in \mathcal{S}_A(t)} (\dot{p}_i(t)) + \sum_{i \in \mathcal{S}_B(t)} (\dot{p}_i(t)). \end{aligned} \quad (23)$$

Observe that, for $u(t) \geq 0$, we have

$$\begin{aligned} \dot{y}(t) + \dot{u}(t) &= C_A(t) - C_B(t) + (a(t) - b(t) - n)u(t) \\ &\quad + 2 \sum_{i \in \mathcal{S}_B(t)} \dot{p}_i(t) \\ &\leq -u(t) = -|u(t)| \leq 0, \end{aligned} \quad (24)$$

where we have used $a(t) \leq n-1$ if $u(t) \geq 0$ (i.e., if $u(t) \geq 0$, at least one load i must be at or below its critical gradient set, X_i^* , by definition of $u(t)$). Similarly, for $u(t) \leq 0$, we have

$$\begin{aligned} \dot{y}(t) - \dot{u}(t) &= C_A(t) - C_B(t) + (a(t) - b(t) + n)u(t) \\ &\quad - 2 \sum_{i \in \mathcal{S}_A(t)} \dot{p}_i(t) \\ &\leq u(t) = -|u(t)| \leq 0. \end{aligned} \quad (25)$$

Consider the function

$$\begin{aligned} z : \mathbb{R}^n &\rightarrow \mathbb{R}_{\geq 0} : z(x) = y(x) + |\bar{g} - \mathbf{1}^T x| \\ &\Rightarrow z(t) = y(t) + |u(t)|. \end{aligned} \quad (26)$$

Just as $y(t)$ is absolutely continuous and differentiable almost everywhere, so is $|u(t)|$, and therefore, $z(t)$. Then, we may combine (24) and (25) to obtain

$$\dot{z}(t) \leq -|u(t)| \leq 0. \quad (27)$$

Hence, $z(t)$ is non-increasing almost everywhere.

Now, suppose $u(t)$ does not converge to 0. For some $t_1 < t_2$, let $|u(t)| > \mu$ for $t_1 < t < t_2$ and some $\mu > 0$. From (27) and the first fundamental theorem of calculus, it then follows that

$$z(t_2) - z(t_1) \leq - \int_{t_1}^{t_2} \mu dt.$$

Because $u(t)$ is continuous and $u(t) \not\rightarrow 0$, there exists a sequence, t_1, t_2, \dots , with $t_k \rightarrow \infty$ as $k \rightarrow \infty$ such that $|u(t)| > \mu$ for all $t_1 < t < t_2, t_3 < t < t_4, \dots$. During all other intervals (e.g., $[t_2, t_3]$), $z(t)$ does not increase since $\dot{z}(t) \leq 0$ for all t . This implies $z(t)$ decreases without bound, which is impossible since the boundedness of the domain, Ω , implies $z(t)$ must be bounded. Therefore, $u(t) \rightarrow 0$. \square

Proof of Lemma 2: Global asymptotic stability of an invariant set requires stability and global attractiveness (see Definition 8.1 in section 8.4 in [8]). Theorem 1 establishes global

attractiveness of X^* , so only stability remains to be proven. For stability, it is sufficient to show that each neighborhood of X^* is positively invariant [8]. In the proof of Lemma 1, it was shown that $z(t)$ is non-increasing (27), where $z(t)$ is defined in (26). Let $Z_m \triangleq \{x \in \Omega | z < m\}$. It follows from the definition that any neighborhood of X^* takes the form of Z_m for some m [8]. Since $z(t)$ is non-increasing, it follows that Z_m is a positively invariant neighborhood of X^* for each $m \geq 0$. Therefore, X^* is stable. This concludes the proof. \square

V. SIMULATION RESULTS

A. Simulation setup

Figure 3 shows the system architecture used for design and simulation. The generator dynamics block shown in Figure 3 also includes local controls that are usually present in generators. The loss of generation is modeled as an exogenous disturbance, \bar{g} , in the figure. The estimator in the figure is the one described in Section III-A to estimate the consumption-generation mismatch from local, noisy frequency measurements. The process disturbance, ζ , and measurement noise, ξ_i , at each load are modeled as wide-sense stationary white noise.

For ease of comparison between the DGP algorithm and the dual algorithm, we use the same generator dynamics, noise statistics, and communication graph as in [27], and the reader is referred to that work for more detailed information about the simulation model or implementation of the state estimator.

Even without the use of intelligent loads, the local generator control will change the generator set point in response to frequency deviation to match consumption, which will restore the frequency to its nominal value on its own. Intelligent loads are supposed to help the generator in reacting to frequency deviations faster so that large excursions of system frequency are avoided.

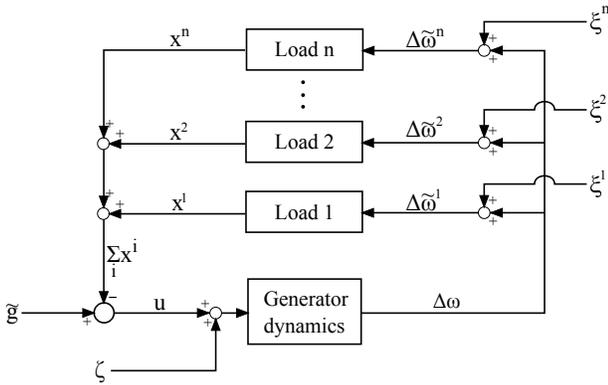


Fig. 3. System architecture for simulations. Inter-load communication is not shown.

For each load i , we choose $\Omega_i = [-\bar{x}_i, \bar{x}_i]$, where \bar{x}_i is chosen from a uniform distribution and then normalized so that $\sum_{i=1}^n \bar{x}_i = 60$ MW (as in [27]).

We test the performance of the DGP algorithm with two distinct disutility functions. The first is a convex but not strictly

convex function:

$$f_i(x_i) = \begin{cases} 0, & |x_i| < a_i \\ q_i(x_i - a_i)^2, & x_i \geq a_i \\ q_i(x_i + a_i)^2, & x_i \leq -a_i \end{cases} \quad (28)$$

where $a_i = 0.1\bar{x}_i$. The consumer does not experience any disutility as long as the load variation is within $\pm a_i$. The second disutility function is strictly convex:

$$f_i(x_i) = q_i(x_i)^2. \quad (29)$$

For both disutility functions, we pick q_i to be an arbitrary positive number such that $1/q_i$ is chosen from a uniform distribution on the interval of $[0.1, 0.3]$. This is chosen for comparison with [27], which makes a similar choice for disutility functions.

The initial conditions are $g[0] = 200$ MW and $u[0] = 0$. Two generation contingencies are modeled as step changes:

$$g[k] = \begin{cases} 200 \text{ MW}, & 0 \leq kT < 20 \text{ s} \\ 190 \text{ MW}, & 20 \leq kT < 50 \text{ s} \\ 170 \text{ MW}, & 50 \leq kT, \end{cases}$$

where $T = 0.1$ seconds is the discretization interval.

Simulations are conducted with the communication network in [27], where load i communicates with loads from $\max\{1, i - n_0\}$ to $\min\{n, i + n_0\}$, where $n_0 \leq n$. We use $n = 1000$ and $n_0 = 1$. The network is shown in Figure 4.

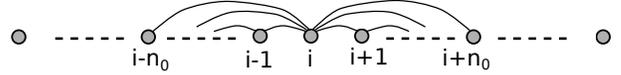


Fig. 4. The communication graph. To avoid clutter, not all edges are shown.

Additionally, we use $c = 5$ and $\gamma[k] = \gamma[0]/(k^{0.8})$ for $k > 0$, with $\gamma[0] = 1.5q/n$, where $q \triangleq \min_i q_i$.

With all of these parameter choices, Assumptions 1, 2, and 3 are satisfied. Note that Assumption 3(3) is satisfied from the discussion in Section III-A.

B. Results with non-strictly convex disutility

Figure 5 shows simulation results for the DGP algorithm with consumer disutility function (28). The dual algorithm (DA) from [27] is not applicable because the inverse of $\nabla f(x)$ must exist in Ω to implement DA, which is not the case when $|x_i| \leq a_i$.

The system frequency without smart loads (i.e., with generator-only control) is shown in red as well. Using DGP, the loads are able to assist the generator in avoiding large frequency deviations from the nominal when each contingency occurs.

C. Comparison with dual algorithm: results with strictly convex disutility

Figure 6 shows results of DGP and DA with quadratic disutilities (29). DGP results in a significantly smaller frequency drop compared to both generator-only control and

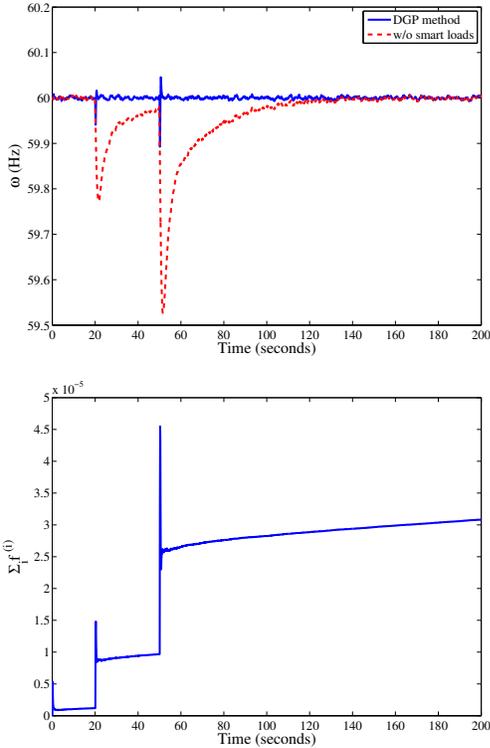


Fig. 5. Performance of the DGP algorithm with consumer disutility that is not strictly convex. Step changes in generation occur at 20 and 50 seconds.

DA. Although DA returns the frequency to the nominal value faster than generator-only control, it does not reduce the initial frequency drop as much as DGP.

However, the consumer disutility is significantly lower for DA than for DGP. This is because DA is responding more slowly than DGP, so the equality constraint is not being satisfied—resulting in a lower cost. The slower response of DA is due to the inversion of the derivative of each load’s disutility function. Because the derivative of each disutility function is rather steep, the inverse is quite flat, so large changes in its argument still result in small changes in its value—leading to small changes in consumption. Conversely, DGP aggressively meets the equality constraint because of the generation-matching step. This results in a lower frequency deviation but more disutility.

Effect of network size and structure: Although not reported here, simulations with varying number of loads ($n = 10, 100$) and varying amount of communication ($n_0 = 10, 100, 1000$) showed similar trends as in the $n = 1000, n_0 = 1$ case. It was observed in [27] that DA showed similar behavior. The reason for this insensitivity to network size and structure is likely the use of local frequency measurements which provides global information, effectively creating a virtual communication link between a fictitious central node and all nodes of the network.

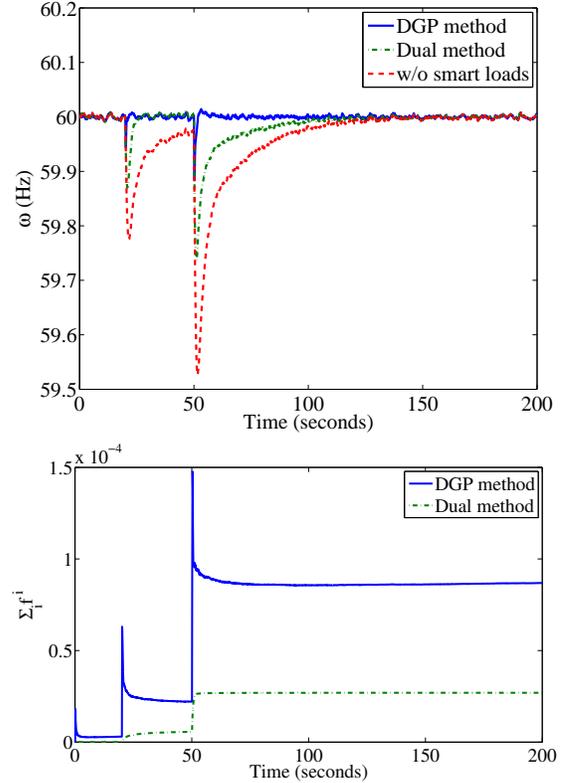


Fig. 6. Performance of the DGP and dual algorithms with quadratic consumer disutility with projection. Step changes in generation occur at 20 and 50 seconds.

VI. CONCLUSION

The DGP algorithm solves a constrained optimization problem in a distributed manner to aid a power grid in maintaining system frequency near its nominal value while minimizing consumers’ disutility. The main advantage of the DGP algorithm over the dual algorithm of [27] is that it is applicable to disutility functions that are convex but not necessarily strictly convex, while the dual algorithm can be used only for strictly convex functions. Disutility functions that are not strictly convex model more realistic consumer behavior that is insensitive to small changes in consumption. Simulations also show that the DGP algorithm performed either better than or similarly to the dual algorithm from [27] in maintaining frequency.

The convergence of the DGP algorithm required that the optimal points lie in the strict interior of the domain. When capacity of the loads is small, the optimal solution is likely to lie on the boundary. An open problem is the design of a distributed primal algorithm, if such an algorithm exists, that is guaranteed to converge to an optimal solution lying on the boundary.

The analysis in this paper assumed a time-invariant communication graph. In contrast, the dual algorithm in [27] was proved to converge even with a time-varying communication graph. Convergence analysis of the DGP algorithm for the time-varying case is left for future work.

One issue that was ignored in our analysis—as well as that

in [27]—is the effect of feedback interconnection between the generator control system and the load control algorithm: simulations reported here and in [27] are conducted with both the control systems in place, and simulation results do not show instability. However, analysis to rule out possible instabilities is lacking. A related issue is actuator dynamics. It is assumed that loads can react as fast as asked by the load control algorithm. The phase lag due to loads' inertia can reduce performance or even cause instability. These are interesting topics worthy of future investigation.

ACKNOWLEDGMENT

The authors thank C. Zhao and S. Low for their assistance in reproducing the results of [27], Sean Meyn for pointing out the connection to the Skorokhod problem, and Vivek Borkar for several useful comments, including pointing us to the work by Nagurney.

REFERENCES

- [1] S. Bolognani and S. Zampieri. A distributed control strategy for reactive power compensation in smart microgrids. *IEEE Trans. on Automatic Control*, 58(11), November 2013.
- [2] Vivek S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.
- [3] Jonathan Brooks and Prabir Barooah. Consumer-aware load control to provide contingency reserves using frequency measurements and inter-load communication. In *American Control Conference*, pages 5008 – 5013, July 2016.
- [4] Ashish Cherukuri and Jorge Cortes. Initialization-free distributed coordination for economic dispatch under varying loads and generator commitment. *Automatica*, 2014. submitted.
- [5] Jorge Cortés. Discontinuous dynamical systems - a tutorial on solutions, nonsmooth analysis, and stability. *IEEE Control Systems Magazine*, 28(3):36–73, 2008.
- [6] Chris Godsil and Gordon Royle. *Algebraic Graph Theory*. Graduate Texts in Mathematics. Springer, 2001.
- [7] Prajwal Khadgi, Lihui Bai, and Gerald Evans. Modeling demand response using utility theory and model predictive control. In *IIE Annual Conference. Proceedings*, page 1262. Institute of Industrial Engineers-Publisher, 2014.
- [8] H.K. Khalil. *Nonlinear Systems 3rd*. Prentice hall Englewood Cliffs, NJ, 2002.
- [9] Brendan Kirby. Ancillary services: Technical and commercial insights. 2007. prepared for Wärtsilä North America Inc.
- [10] Peter K Kitanidis. Unbiased minimum-variance linear state estimation. *Automatica*, 23(6):775–778, 1987.
- [11] Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. 2nd edition, 2003.
- [12] Debra Lew, Greg Brinkman, E Ibanez, BM Hodge, and J King. The western wind and solar integration study phase 2. *National Renewable Energy Laboratory, NREL/TP-5500, 55588*, 2013.
- [13] Yashen Lin, Prabir Barooah, Sean Meyn, and Timothy Middelkoop. Experimental evaluation of frequency regulation from commercial building HVAC systems. *IEEE Transactions on Smart Grid*, 6:776 – 783, 2015.
- [14] David G. Luenberger. *Linear and Nonlinear Programming*. Springer, 2 edition, 2003.
- [15] Y. V. Makarov, Lu. S., J. Ma, and T. B. Nguyen. Assessing the value of regulation resources based on their time response characteristics. Technical Report PNNL-17632, Pacific Northwest National Laboratory, Richland, WA, June 2008.
- [16] Michael Milligan and Brendan Kirby. Utilizing load response for wind and solar integration and power system reliability. In *Wind Power Conference, Dallas, Texas*, 2010.
- [17] Angel Molina-García, François Bouffard, and Daniel S Kirschen. Decentralized demand-side contribution to primary frequency control. *Power Systems, IEEE Transactions on*, 26(1):411–419, 2011.
- [18] Anna Nagurney and Ding Zhang. *Projected dynamical systems and variational inequalities with applications*, volume 2. Springer Science & Business Media, 2012.
- [19] Angelia Nedic and Asuman Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009.
- [20] Prepared by NERC RS Committee. Balancing and Frequency Control: A Technical Document Prepared by the NERC Resources Subcommittee. NERC Technical Report, January 26 2011.
- [21] F.C. Schweppe, R.D. Tabors, J.L. Kirtley, H.R. Outhred, F.H. Pickel, and A.J. Cox. Homeostatic utility control. *PAS-99(3):1151 –1163*, May 1980.
- [22] Joe A Short, David G Infield, and Leon L Freris. Stabilization of grid frequency through dynamic demand control. *Power Systems, IEEE Transactions on*, 22(3):1284–1293, 2007.
- [23] Pierluigi Siano. Demand response and smart grids: a survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, 2014.
- [24] Anatoliy V Skorokhod. Stochastic equations for diffusion processes in a bounded region. *Theory of Probability & Its Applications*, 6(3):264–274, 1961.
- [25] D Todd, M Caufield, B Helms, Alcoa Power Generating, Inc M Starke, B Kirby, and J Kueck. Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of ALCOA INC. *ORN/MTM*, 233, 2008.
- [26] Bowen Zhang, Michael C Caramanis, and John Baillieul. Optimal price-controlled demand response with explicit modeling of consumer preference dynamics. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 2481–2486. IEEE, 2014.
- [27] Changhong Zhao, Ufuk Topcu, and Steven H Low. Optimal load control via frequency measurement and neighborhood area communication. *Power Systems, IEEE Transactions on*, pages 3576–3587, 2013.
- [28] Minghui Zhu and Sonia Martínez. On distributed convex optimization under inequality and equality constraints. *Automatic Control, IEEE Transactions on*, 57(1):151–164, 2012.
- [29] Minghui Zhu and Sonia Martínez. An approximate dual subgradient algorithm for multi-agent non-convex optimization. *Automatic Control, IEEE Transactions on*, 58(6):1534–1539, 2013.