

Frequency violations from random disturbances: an MCMC approach

John Moriarty¹, Jure Vogrinc¹ and Alessandro Zocca²

Abstract—The frequency stability of power systems is increasingly challenged by various types of disturbances. In particular, the increasing penetration of renewable energy sources is increasing the variability of power generation and at the same time reducing system inertia against disturbances. In this paper we are particularly interested in understanding how rate of change of frequency (RoCoF) violations could arise from unusually large power disturbances.

We devise a novel specialization, named *ghost sampling*, of the Metropolis-Hastings Markov Chain Monte Carlo method that is tailored to efficiently sample rare power disturbances leading to nodal frequency violations. Generating a representative random sample addresses important statistical questions such as “which generator is most likely to be disconnected due to a RoCoF violation?” or “what is the probability of having simultaneous RoCoF violations, given that a violation occurs?” Our method can perform conditional sampling from any joint distribution of power disturbances including, for instance, correlated and non-Gaussian disturbances, features which have both been recently shown to be significant in security analyses.

I. INTRODUCTION

Frequency stability is a prime concern of transmission system operators, as frequency instabilities may lead to machines desynchronization and trigger large power outages [ADF⁺05]. Transmission systems are experiencing increased stress and approaching their stability limits due, for example, to the continued connection of power electronics, increased uncertainty, cross-border power flows and the addition of high-voltage direct current (HVDC) links [WEBK15]. In particular, system inertia is decreasing as synchronous machines are replaced by inverter-connected distributed generation. Low inertia levels, together with the variability of renewable generation, can lead to large swings in the power system frequency [UBA14]. While promising mitigations exist including participation from loads [ZTLL14], [VPMB16], distributed energy resources [GZD⁺17] and virtual inertia [PBD17], it is increasingly important to also understand the stability of power system frequency under random disturbances to the network’s power injections. In the Irish transmission system, for example, the rate of change of frequency (RoCoF) has been identified as the key limit to allowing high real-time penetrations of wind generation [CMSR13]. Further it has been shown that frequency fluctuations have a heavy-tailed distribution [SBA⁺18], making strong deviations more likely than would be expected under, for example, a Gaussian model of fluctuations. Beyond questions of system stability it has also been shown that stochastic disturbances can cause significant resistive power losses to be incurred in stabilising the system frequency [TBG15].

In this paper we aim to investigate the extent to which infrequently observed, but large, disturbances to nodal power injections can cause an unexpected rate of change in the power system frequency. In particular, following a disturbance to one or more nodal power injections we model the rate of change of frequency (RoCoF) over the primary control timescale. Motivated by stability considerations for RoCoF relays, we develop a simulation technique to investigate possible causes of RoCoF violations. The frequency considered is either the system average frequency or the set of nodal frequencies, and different characterizations of RoCoF violations are explored.

Estimating the probability of rare events in power systems is computationally challenging. Recent work in this area includes [OM17], [NZZ17a], [NZZ17b] and the present paper is complementary to such studies. Instead we aim to generate a representative sample of disturbances, *conditional* on a RoCoF violation occurring. To this end we present *ghost sampling*, a specialization of the Metropolis-Hastings Markov Chain Monte Carlo (MCMC) method. Relative to current probabilistic power system reliability analyses, a key advantage of MCMC is to allow the power disturbances to have arbitrary joint distribution. Whereas independent disturbances have always been assumed so far in the literature (see [PM17]), the use of MCMC allows the disturbance magnitudes to have both correlation and arbitrary marginal probability distributions. Simultaneous nodal power disturbances with a common cause, due for example to line failures or large weather fronts, may thus be modelled in this framework.

The low probability of RoCoF violations is, however, a challenge to the standard Metropolis-Hastings MCMC algorithm. In common with many other sampling techniques, the latter approach may result in a large part of computational effort being expended in proposing non-violating states. Further, when a violating state is sampled, standard MCMC chains risk becoming ‘stuck’ in its vicinity (see Section 1.11.2 of [BGJM11]). The ‘ghost proposal’ described below mitigates these issues by proposing only violating states.

In the nodal analysis we derive expressions for the full set of RoCoFs at time $t \geq 0$. We also obtain expressions for the set of disturbances whose maximum RoCoF exceeds a given acceptable threshold, which may vary per node. From the representative sample generated we are able to estimate quantities of interest conditional on a violation, in contrast with the average-case analyses common in the literature, see e.g., [PM17], [TBG15].

II. MODEL DESCRIPTION

A power system described by the graph $G = (\mathcal{G}, \mathcal{E})$ is considered, with nodes (buses) $\mathcal{G} = \{1, \dots, n\}$ and m edges (transmission lines) $\mathcal{E} \subseteq \mathcal{G} \times \mathcal{G}$. It is assumed

¹John Moriarty and Jure Vogrinc are with the School of Mathematics, Queen Mary University of London, London E1 4NS, UK. Email: j.moriarty@qmul.ac.uk and j.vogrinc@qmul.ac.uk

²Alessandro Zocca is with the California Institute of Technology, Pasadena CA, 91125, USA. Email: azocca@caltech.edu

that G is a reduced network in which each bus houses a generation unit, since passive loads can be eliminated via Kron reduction [DB13], [DB10].

Writing ω_j for the frequency at node $j \in \mathcal{G}$, the time evolution of nodal frequencies is modelled via linearized dynamics as

$$M_j \dot{\omega}_j + D_j \omega_j = p_j^{\text{in}} - p_j^{\text{out}}, \quad \forall j \in \mathcal{G}, \quad (1)$$

where $M_j > 0$ is the inertia of the generator at node $j \in \mathcal{G}$, $D_j > 0$ is the damping/droop control coefficient, and p_j^{in} and p_j^{out} represent, respectively, the mechanical power injected by the generator at node j and the net electrical power drawn by the network from node j ; see [Kun94] for more details.

Reactive power injections and reactive power flows are neglected and the standard assumptions of lossless lines, time-invariant identical voltage magnitudes across all nodes and small-signal approximations [PMVB05], [WWS14] are made. In view of these assumptions, the so-called DC power flow approximation holds, namely

$$p_j^{\text{out}} = \sum_{i \in \mathcal{G}} f_{i,j} = \sum_{i \in \mathcal{G}} B_{i,j} (\theta_i - \theta_j), \quad (2)$$

where $f_{i,j}$ describes the power flow on line $e = (i, j) \in \mathcal{E}$, $B_e = B_{i,j} \geq 0$ is the (effective) susceptance between nodes i and j and θ_j denotes the phase angle at node $j \in \mathcal{G}$. Note that an arbitrary but fixed orientation has been chosen for the edges in \mathcal{E} , which is captured by the incidence matrix $C \in \{-1, 0, 1\}^{n \times m}$ of G , that is

$$C_{i,e} = \begin{cases} 1 & \text{if } e = (i, j), \\ -1 & \text{if } e = (j, i), \\ 0 & \text{otherwise.} \end{cases}$$

Denoting by $B \in \mathbb{R}^{m \times m}$ the diagonal matrix with the susceptances $\{B_e\}_{e=1, \dots, m}$ as diagonal entries, the relation between line flows and phase angles may be rewritten in matrix form as

$$f = BC^T \theta,$$

where $f \in \mathbb{R}^m$ and $\theta \in \mathbb{R}^n$ are the vectors of line flows and phase angles, respectively.

We are interested in how, starting from an equilibrium point, the network reacts to a vector $u \in \mathbb{R}^n$ of nodal power disturbances. In view of (1) and (2), the *deviations* from their nominal values of the nodal frequencies and line power flows are then described by

$$M_j \dot{\omega}_j = -D_j \omega_j + u_j - \sum_{i: (i,j) \in \mathcal{E}} f_{i,j}, \quad \forall j \in \mathcal{G}, \quad (3a)$$

$$\dot{f}_{i,j} = B_{i,j} (\omega_i - \omega_j), \quad \forall (i, j) \in \mathcal{E}, \quad (3b)$$

where, with a minor abuse of notation, the variables ω and f henceforth denote deviations from the corresponding nominal values at equilibrium. This means, in particular, that at equilibrium all variables in equations (3) are equal to 0.

The entries u_j , $j \in \mathcal{G}$ of the random disturbance vector $u \in \mathbb{R}^n$ are modelled as continuous random variables with joint probability density function π , so that for any measurable subset $A \subseteq \mathbb{R}^n$ we have

$$\mathbb{P}[u \in A] = \int_A \pi(u_1, \dots, u_n) du_1 \dots du_n. \quad (4)$$

The correlation in renewable generation, alongside correlation in other factors such as loads, has been shown to have a significant effect on power system risk assessment [LZW⁺15]. One advantage of our approach is that the random disturbances u_j are not required to be independent. This is because we aim to simulate typical disturbances u causing frequency violations rather than, for example, to derive closed form expressions for synchronization performance as in [PM17] or [TBG15].

Further, the errors in renewable power forecasts have been shown to have significantly non-Gaussian distributions. For example, fat tails have been demonstrated in wind power forecast errors [BDNL08]. This is another advantage of our conditional simulation procedure, since the random disturbances can have a general joint probability density. To illustrate this point the case study presented later in Section VI uses a mixture of uncorrelated Gaussian and correlated, fat tailed non-Gaussian distributions.

The u_j are modelled as step disturbances, namely

$$u(t) = u \mathbf{1}_{\{t \geq 0\}}.$$

Thus time $t = 0$ is the moment just after the random disturbance(s). The desynchronization effect of u on the frequencies at all nodes $j \in \mathcal{G}$ will be modelled from time $t = 0$ until time $t = \epsilon > 0$. This step model is valid when the disturbances represented by the u_j can be reasonably approximated as constant over the time interval $[0, \epsilon]$. (In the case study below we take $\epsilon = 0.5$ s.)

Our method in the remainder of the paper has two parts, as follows:

- 1) characterise the ‘safe region’ $K \subset \mathbb{R}^n$ of disturbances $u \in \mathbb{R}^n$ which do not give rise to frequency violations;
- 2) generate a statistically representative sample from its complement K^c .

Frequency violations will be characterised using the RoCoF, by which we mean $|\dot{\omega}|$, the magnitude of the rate of change of frequency. The analysis begins with the rate of change of the system frequency, before moving to the consideration of nodal frequencies. The latter context is particularly pertinent when nodal frequencies are considered, since generating machines are protected by RoCoF-sensitive relays which observe only the local nodal frequency.

The rest of the paper is organized as follows. In Section III we introduce step 1) in the simpler context of analysing the system frequency. In Section IV the nodal frequency dynamics are first established, and the step 1) is then carried out in this context. Step 2) is developed in Section V, and an illustrative case study for our analysis is provided in Section VI.

III. SYSTEM FREQUENCY

The *system frequency* or *center of inertia* (COI) is defined as (see, for example, [UBA14]):

$$\bar{\omega}(t) := \frac{\sum_{i \in \mathcal{G}} M_i \omega_i(t)}{\sum_{i \in \mathcal{G}} M_i}.$$

This model has been studied in [PM17] under the simplifying condition that there exist rating parameters $f_1, \dots, f_n > 0$ with $\max_i f_i = 1$ such that the inertia and damping coefficients of generator i are given respectively by

$$M_i = f_i M \quad \text{and} \quad D_i = f_i D, \quad i \in \mathcal{G}, \quad (5)$$

where M and D are those of the machine j such that $f_j = 1$. In particular, it is shown in the latter paper (cf. Eq.(18)) that the following holds under condition (5):

$$\bar{\omega}(t) = g(t) \sum_i u_i, \quad t > 0, \quad (6)$$

where $g(t) := (\sum_i D_i)^{-1} (1 - e^{-(D/M)t})$.

The swing and network dynamics in (3a)-(3b) can be enriched to incorporate the turbine control dynamics, yielding the following third-order model

$$\dot{\omega}_j = -\frac{1}{M_j} \left(D_j \omega_j - q_j - u_j + \sum_{i: (i,j) \in \mathcal{E}} f_{i,j} \right) \quad (7a)$$

$$q_j = -\frac{1}{\tau} (R_j^{-1} \omega_j + q_j) \quad (7b)$$

where q_j is the (variation of) turbine power, R_j the droop coefficient, and τ is the turbine time constant (which is uniform across different generators). If in addition to (5) we further assume that $R_j^{-1} = f_j R^{-1}$ for every $j \in \mathcal{G}$ and that the system is under-damped, i.e., $\omega_d^2 := \frac{D+R^{-1}}{M\tau} - \frac{1}{4} \left(\frac{1}{\tau} + \frac{D}{M} \right)^2 > 0$, the system frequency still obeys an equation of the same form as (6), where the function $g(t)$ is now a more involved function, namely

$$g(t) := \frac{1 - e^{-\eta t} (\cos(\omega_d t) - \frac{\gamma - \eta}{\omega_d} \sin(\omega_d t))}{(\sum_i D_i + R_i^{-1})(D + R^{-1})},$$

where $\eta := \frac{1}{2} \left(\frac{1}{\tau} + \frac{D}{M} \right)$ and $\gamma := \left(\frac{1}{\tau} - \frac{R^{-1}}{M} \right)$.

From both models, with or without turbine control, it is proved in [PM17] that the maximum RoCoF occurs at time $t \downarrow 0+$ and is equal to

$$\max_{t>0} \left| \frac{d}{dt} \bar{\omega}(t) \right| = \lim_{t \rightarrow 0+} \left| \frac{d}{dt} \bar{\omega}(t) \right| = \left| \sum_{i \in \mathcal{G}} u_i \right| |\dot{g}(0)|.$$

Hence, for both dynamics, with or without turbine control, the set of disturbances $u \in \mathbb{R}^n$ whose maximum induced RoCoF does not exceed a predetermined threshold r_{\max} is simply

$$K_{MS} = \left\{ u \in \mathbb{R}^n : \left| \sum_i u_i \right| \leq \frac{r_{\max}}{|\dot{g}(0)|} \right\}.$$

The region K_{MS} is a convex polytope in \mathbb{R}^n – that is, the intersection of a number of half-spaces, which is not necessarily bounded. (Here the subscript M refers to the Maximum RoCoF metric and S to the System frequency).

The metric of average absolute RoCoF may alternatively be considered. Over the time interval of length ϵ following the disturbance u this is given by $\Omega(\epsilon)$, where


$$\Omega(\epsilon) := \frac{1}{\epsilon} \int_0^\epsilon |\dot{\bar{\omega}}(t)| dt = \frac{1}{\epsilon} \left| \sum_i u_i \right| \int_0^\epsilon |\dot{g}(t)| dt. \quad (8)$$

The set of disturbances whose average induced RoCoF over $t \in [0, \epsilon]$ does not exceed the threshold r_{\max} is thus


$$K_{AS} = \left\{ u \in \mathbb{R}^n : \left| \sum_i u_i \right| \leq \frac{\epsilon r_{\max}}{\int_0^\epsilon |\dot{g}(t)| dt} \right\}.$$

Since the turning points of g can be calculated analytically for both models, the evaluation of the integral is straightforward. The set K_{AS} of disturbances is also a convex polytope.

In summary, the ‘safe region’ of disturbances $u \in \mathbb{R}^n$ inducing a rate of change in the system frequency less than a threshold r is a convex polytope, for either the maximum or average absolute RoCoF metric, and with or without turbine control. In the next section we are able to establish a similar result for the set of nodal frequencies under swing dynamics.



(a) Time evolution over the first 2s of the nodal frequency deviations $\omega_1(t)$, $\omega_2(t)$, and $\omega_3(t)$ and the system frequency $\bar{\omega}(t)$ (Hz) in the case study of Section VI after a random disturbance.



(b) Corresponding evolution over the first 2s of frequency speed deviations $\dot{\omega}_1(t)$, $\dot{\omega}_2(t)$, $\dot{\omega}_3(t)$, and $\dot{\bar{\omega}}(t)$ (Hz/s) for the same random disturbance as in Fig. 1(a). The dashed horizontal lines represent the RoCoF threshold $r_{\max} = 1$ Hz/s.

Fig. 1: Post-disturbance traces of some nodal frequencies for the case study about the IEEE 39 New England interconnection system presented of Section VI.

IV. RoCoF VIOLATIONS FOR NODAL FREQUENCIES

While the system frequency is of inherent interest, it cannot capture de-synchronization within the system. To illustrate this point, Fig. 1 plots three nodal frequency traces together with the system frequency following a random disturbance. These traces are simulated from the system described in the case study of Section VI. It is clear from the figure that, under either the maximum or average absolute RoCoF metric, it is possible for a given threshold to be simultaneously respected by the system frequency and violated by one or more nodal frequencies. Further, as mentioned above, from the practical perspective it is nodal frequencies which trigger the operation of RoCoF-sensitive generator protection relays which can, in turn, lead to significantly more serious frequency violations. The main aim of this paper is therefore sampling the typical causes

of nodal frequency violations, and we focus on this from now on. We remark that the generator inertia and damping coefficients M_i and D_i respectively may be arbitrary, so that our setting is more general than that presented in Section III (cf. condition (5)).

Let $M \in \mathbb{R}^{n \times n}$ and $D \in \mathbb{R}^{n \times n}$ be the diagonal matrices containing the generator inertias and damping factors respectively, and $u \in \mathbb{R}^n$ the random vector of disturbances. Together with the notation from the previous section, the swing equations (3a) and (3b) read

$$\begin{aligned} \begin{bmatrix} \dot{\omega} \\ \dot{f} \end{bmatrix} &= \begin{bmatrix} -M^{-1}D & -M^{-1}C \\ BC^T & \mathbb{O} \end{bmatrix} \cdot \begin{bmatrix} \omega \\ f \end{bmatrix} + \begin{bmatrix} M^{-1} \\ 0 \end{bmatrix} u, \\ \begin{bmatrix} \omega(0) \\ f(0) \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \end{aligned}$$

By differentiation we obtain

$$\begin{aligned} \ddot{\omega} &= -M^{-1}D\dot{\omega} - M^{-1}CBC^T \omega \\ &= -M^{-1}D\dot{\omega} - M^{-1}L\omega, \\ \dot{\omega}(0) &= M^{-1}u, \end{aligned}$$

where $L := CBC^T$ is the weighted Laplacian matrix of the graph G . Ignoring the line flows, we thus obtain a homogeneous dynamical system of the form $\dot{x} = Ax$, with

$$\begin{aligned} x &= \begin{bmatrix} \dot{\omega} \\ \omega \end{bmatrix}, \quad A = \begin{bmatrix} -M^{-1}D & -M^{-1}L \\ I & \mathbb{O} \end{bmatrix}, \\ x(0) &= \begin{bmatrix} M^{-1}u \\ 0 \end{bmatrix}, \end{aligned}$$

whose solution is

$$x(t) = \exp(tA)x(0). \quad (9)$$

Henceforth the maximum RoCoF metric will be used to characterise frequency violations: the average metric may be applied in a similar manner but we reserve this for future work. The set of disturbances u which do not induce RoCoF violations will again be referred to as the ‘safe region’.

As confirmed by Fig. 1(b), and in contrast to the system frequency models of Section III, for a fixed node j the maximum RoCoF $\dot{\omega}_j$ does not in general occur at time 0. Let us therefore consider sampling $\dot{\omega}_j$ at times $\frac{n}{N}\epsilon$, $n = 0, \dots, N$. Although in principle this involves no loss of generality since digital RoCoF measurements have a discrete sampling rate, we note that any lower sampling rate N/ϵ should be chosen carefully to avoid an excessive loss of sensitivity (a sensitivity analysis for N is provided in the case study of Section VI). Define the ‘node- j safe region’ $K^{(j,N)}$ by

$$K^{(j,N)} = \bigcap_{n=0}^N K_n^{(j,N)}, \quad \text{where} \quad (10)$$

$$K_n^{(j,N)} = \left\{ u \in \mathbb{R}^n : \left| \dot{\omega}_j \left(\frac{n}{N}\epsilon \right) \right| \leq r_{\max} \right\}. \quad (11)$$

It follows from (9) that $K_n^{(j,N)}$ is given by the convex polytope

$$K_n^{(j,N)} = \left\{ u \in \mathbb{R}^n : \left| \exp \left(\frac{n}{N}A \right)_j \begin{bmatrix} M^{-1}u \\ 0 \end{bmatrix} \right| \leq r_{\max} \right\},$$

where $\exp(tA)_j$ denotes the j -th row of the matrix $\exp(tA)$. Hence $K^{(j,N)}$ and the ‘all-nodes safe region’ $K^{(N)}$ are also convex polytopes, where

$$K^{(N)} = \bigcap_{j \in \mathcal{G}} K^{(j,N)}. \quad (12)$$

(Note that, clearly, different thresholds r_{\max} could be chosen per node to enable modelling of differing protection relay settings for differing types of generating machine, or to enable to modelling of DC links, and the safe region would again be a convex polytope).

Having characterised the safe region, we now turn to the problem of generating a representative sample from its complement. In the next section we describe how the Metropolis-Hastings MCMC algorithm, a commonly used technique for generating random samples, may be efficiently adapted for this purpose.

V. GHOST SAMPLING

Recalling from (4) that the entries u_j of the random disturbance u are modelled as continuous random variables with a joint probability density function π , the goal in this section is to sample efficiently from the conditional joint density, or *target*,

$$\frac{\pi(u)\mathbf{1}_{K^c}(u)}{\pi(K^c)}, \quad \text{where} \quad \pi(K^c) = \int_{K^c} \pi(v)dv.$$

Since MCMC sampling methods do not require strong assumptions on the target density they are ideally suited to such problems [BGJM11], [Tie94]. However the event K^c is in principle rare, which may cause problems of computational inefficiency. Below we describe the *ghost sampler*, a particular Metropolis-Hastings (MH) algorithm designed to be efficient in this context.

The ghost sampler is defined in Section V-A, and in Section V-B it is shown that the generated samples may be used to approximate statistics of the corresponding target distribution. We will show in the case study of Section VI that this enables important statistical questions to be addressed such as ‘‘which generator is most likely to be disconnected due to a RoCoF violation?’’ or ‘‘what is the probability of two simultaneous RoCoF violations being caused, given that a violation occurs?’’.

A. Ghost sampling algorithm

Ghost sampling lies in the class of Metropolis-Hastings algorithms [RR04], [Tie94]. That is, beginning at $X_0 \in \mathbb{R}^n$, for every $i = 1, 2, \dots$ we generate a proposal Y_{i+1} distributed according to a density $q(X_i, y)dy$, and evaluate the acceptance probability

$$\alpha(X_i, Y_{i+1}) = \min \left(1, \frac{\pi(Y_{i+1})\mathbf{1}_{K^c}(Y_{i+1})q(Y_{i+1}, X_i)}{\pi(X_i)\mathbf{1}_{K^c}(X_i)q(X_i, Y_{i+1})} \right), \quad (13)$$

which is interpreted as one if $\pi(X_i)\mathbf{1}_{K^c}(X_i)q(X_i, Y_{i+1}) = 0$. With probability $\alpha(X_i, Y_{i+1})$ the proposal is accepted and we set $X_{i+1} = Y_{i+1}$, otherwise it is rejected and $X_{i+1} = X_i$. The aim is to generate a Markov chain X_1, X_2, \dots with stationary distribution equal to $\frac{\pi\mathbf{1}_{K^c}}{\pi(K^c)}$ which satisfies the law of large numbers (LLN), meaning that sample averages

$\frac{1}{n} \sum_{i=1}^n f(X_i)$ for large n provide good estimates for the actual conditional expectations

$$\frac{\int_{K^c} f(v) \pi(v) dv}{\pi(K^c)} = \mathbb{E}_\pi[f(X) | X \notin K].$$

Commonly there is an underlying symmetric density $q: \mathbb{R}^n \rightarrow \mathbb{R}$ (that is, with $q(x) = q(-x)$) and the proposal density used in MH algorithm is $q(x, y) = q(|y - x|)$ where we abuse notation slightly by denoting both with q . In this case the algorithm is called *Symmetric Random Walk Metropolis* (SRWM), $q(x, y) = q(y, x)$ holds and the q terms in (13) cancel out. Typical examples are $Y_{i+1} \sim N(x, \sigma^2 I_n)$ or $Y_{i+1} \sim X_I + U([- \delta, \delta]^d)$, that is, the proposal is drawn from a normal (resp. uniform) distribution centred at X_i . Note from (13) that knowledge of $\pi \mathbf{1}_{K^c}$ suffices and the normalising constant $\pi(K^c)$ need not be known.

A well-known difficulty with the MH algorithm arises when the target density is multi-modal (see Section 1.11.2 in [BGJM11]). In the present application to rare event sampling, where the ‘‘common’’ events are removed from π , we may be left with a target density $\frac{\pi \mathbf{1}_{K^c}}{\pi(K^c)}$ with multiple, well-separated local modes. The difficulty arises since a large proportion of the proposals X_i will lie in K and thus be rejected (since then $\alpha(X_i, Y_{i+1}) = 0$), rendering the method inefficient.

The *ghost proposal* is designed to circumvent these issues by moving ‘through’ K , and is now described in the case when K is a bounded convex polytope (clearly K should also have nonzero volume).

Fix an SRWM algorithm with proposal density q and target $\frac{\pi \mathbf{1}_{K^c}}{\pi(K^c)}$. Denote the boundary of K by δK , and let the current state of the chain be $X_i \notin K$. First, generate a SRWM proposal Y_{i+1} and denote $\varphi_i := \frac{Y_{i+1} - X_i}{|Y_{i+1} - X_i|}$. Then with probability 1 we have $Y_{i+1} \neq X_i$ and the ray from X_i passing through Y_{i+1} intersects δK either twice (cf. Fig. 2(b)-(c)) or not at all (cf. Fig. 2(a)). If there are two numbers $t_2 > t_1 > 0$ such that $X_i + t\varphi_i \in \delta K$, modify the proposal to $Z_{i+1} = Y_{i+1} + (t_2 - t_1)\varphi$ (cf. Fig. 2(d)). If there are no such points, set $Z_{i+1} = Y_{i+1}$ (cf. Fig. 2(a)). Finally perform a MH step, accepting the proposal Z_{i+1} with probability $\alpha(X_i, Z_{i+1})$ given by (13) and setting $X_{i+1} = Z_{i+1}$, otherwise rejecting the proposal and setting $X_{i+1} = X_i$. This procedure is depicted in Fig. 2 and pseudocode is provided in Algorithm 1.

The following example illustrates how ghost sampling improves upon the standard MH algorithm in the present context. Set $K = \{(x, y) \in \mathbb{R}^2 : |x| + |y| < 7\}$, let π be the two-dimensional Gaussian density with zero mean and covariance matrix $\text{diag}(4, 1)$ and let q be the density of a standard two-dimensional Gaussian random variable.

A MH chain starting to the left of the diamond-shaped set K (cf. Fig. 3) will have difficulties crossing to the right side of the diamond, since a direct move to the other side is unlikely and any sequence of steps towards the right side is likely to suffer rejections because the values of π are much smaller around the top and bottom vertices of K than around the left and right vertices. The ghost sampler, however, is likely to make a direct move between the left and right sides. This can be seen in Fig. 3, where scatter plots are given of values taken by the standard MH chain (in red) and the ghost sampler (blue). In this way the ghost sampler is designed to more efficiently explore the rare event K^c .




Fig. 2: Illustration of the key ideas underlying the ghost sampling method.

Algorithm 1: Ghost sampler (i -th step)

input : $X_i \in K^c$

- 1 Generate SRWM proposal Y_{i+1} ;
- 2 Calculate direction $\varphi_i = \frac{Y_{i+1} - X_i}{|Y_{i+1} - X_i|}$;
- 3 Calculate all points $T := \{t > 0 : X_i + t\varphi \in \delta K\}$;
- 4 **if** $T = \{t_1, t_2\}$ **then**
- 5 | $Z_{i+1} = Y_{i+1} + (t_2 - t_1)\varphi$;
- 6 **else**
- 7 | $Z_{i+1} = Y_{i+1}$;
- 8 **end**
- 9 Generate a uniform random variable U on $[0, 1]$;
- 10 **if** $U \leq \alpha(X_i, Z_{i+1})$ **then**
- 11 | $X_{i+1} = Z_{i+1}$;
- 12 **else**
- 13 | $X_{i+1} = X_i$;
- 14 **end**
- 15 **return** X_{i+1}




Fig. 3: Example values taken by the ghost sampler (blue) and standard MH chain (red).

B. Ghost sampler properties

We now state necessary mathematical properties of the ghost proposal in a more general setting, which are proved in the Appendix. More precisely, we assume only that K is closed and *ray-bounded* (see Definition 2).

For a starting point $x \in K^c$, a direction φ lying in the unit sphere \mathbb{S}^{n-1} and a distance $r > 0$, denote the total length between x and $x + r\varphi$ that lies outside K by

$$l_x^\varphi(r) := \int_0^r \mathbf{1}_{K^c}(x + t\varphi) dt. \quad (14)$$

Also denote the mapping $T_x^K: \mathbb{R}^n \rightarrow \mathbb{R}^n$ by

$$T_x^K(x + r\varphi) := x + l_x^\varphi(r)\varphi \quad (15)$$

and a modified proposal density termed the *ghost density* as

$$q_K(x, x + r\varphi) := q(l_x^\varphi(r)\varphi) \left(\frac{l_x^\varphi(r)}{r} \right)^{n-1} \mathbf{1}_{K^c}(x + r\varphi), \quad (16)$$

where $l_x^\varphi(r)/r$ is interpreted as 1 if $r = 0$.

Note that for $x, y \in K^c$ we have $q_K(x, y) = q_K(y, x)$, since

$$l_x^{\frac{y-x}{|y-x|}}(|y-x|) = l_y^{\frac{x-y}{|x-y|}}(|x-y|).$$

Intuitively the map T_x^K contracts each ray emanating from x by removing its intersection with K . It is the reverse of the ghost sampling modification, in the sense that $T_{X_i}^K(Z_{i+1}) = Y_{i+1}$.

Lemma 1. *If K is a closed set and $x \in K^c$, then the map $T_x^K: K^c \rightarrow \mathbb{R}^n$ is injective.*

Definition 2. *A closed set K is said to be ray-bounded if the map $T_x^K: K^c \rightarrow \mathbb{R}^n$ is surjective for all $x \in K^c$.*

In the context of convex polytopes ray boundedness simply means that any ray starting outside K and intersecting it will also exit it. For instance, the set $\{(x, y) \in \mathbb{R}^2: |x| < 1\} \subset \mathbb{R}^2$ is ray-bounded but not bounded.

For each $x \in K^c$ and measurable set $A \subset \mathbb{R}^n$ we define

$$Q(x, A) := \int_{\mathbb{R}^n} q(x, y) \mathbf{1}_A(y) dy,$$

$$Q_K(x, A) := \int_{\mathbb{R}^n} q_K(x, y) \mathbf{1}_A(y) dy.$$

The following lemma shows that in the case of ray-bounded polytope procedure described in Section V-A indeed results in the ghost density defined in (16).

Lemma 3. *For every $x \in K^c$ and every measurable set $A \subset \mathbb{R}^n$ we have*

$$Q_K(x, A) = Q(x, T_x^K(A \cap K^c)).$$

Whenever $x \in K^c$ and K is ray-bounded, Lemma 3 gives that $Q_K(x, K^c) = Q(x, T_x^K(K^c)) = Q(x, \mathbb{R}^n) = 1$, which means that Q_K is a *Markov kernel* on K^c . It then follows from [Tie94, Section 2.3.1] that the measure with density $\frac{\pi \mathbf{1}_{K^c}}{\pi(K^c)}$ is the unique stationary probability measure of the MH algorithm with the ghost proposal density q_K .

To complete the justification of ghost sampling, we are also able to show the LLN:

Theorem 4. *Let K be closed and ray-bounded. Let X_1, X_2, \dots be a Markov chain generated by the ghost proposal q_K derived from a SRWM proposal with density q which is strictly positive on \mathbb{R}^n . Then the strong law of large numbers holds, that is, for every π -integrable function f we have:*

$$\frac{1}{n} \sum_{i=1}^n f(X_i) \xrightarrow[n \rightarrow \infty]{a.s.} \mathbb{E}_\pi[f(X) | X \notin K].$$

VI. CASE STUDY: IEEE 39 NEW ENGLAND NETWORK

In this section we illustrate how the ghost sampler enables inference about RoCoF violations. As a case study we consider the IEEE 39-bus New England interconnection system, which has 10 generators and 29 load nodes, see Fig. 4(a). The system parameters for our experiments are taken from the Matpower Simulation Package [ZMST11].

We consider the Kron-reduced version of the aforementioned system, which is illustrated in Fig. 4(b). The thickness of the edges in Fig. 4(b) is proportional to the equivalent susceptance between the two corresponding generator nodes.




Fig. 4: (a) Line diagram of the IEEE 39-bus system and (b) the Kron-reduced version of the IEEE 39-bus system with only the 10 generator nodes.

The ghost sampler is capable of sampling from any continuous distribution one may want to consider for the power disturbances, in particular those featuring heavy-tailed

or correlated components. Aiming to illustrate its potential, we thus consider a mixed distribution that prescribes the disturbances u_1 and u_2 in generators 1 and 2 to be correlated and heavy-tailed, while the remaining generators are assumed to have i.i.d. Gaussian disturbances.

More specifically, we model the disturbances u_3, \dots, u_{10} as independent Gaussian random variables with zero mean and standard deviations of the size $1/65$ times the nominal power injections of the associated generators, while the disturbances u_1 and u_2 are modelled by a correlated heavy-tailed density

$$(u_1, u_2) \sim \frac{1}{1 + (30(u_1 - u_2/2))^4} \cdot \frac{1}{1 + (30(u_2 - u_1/2))^4}.$$

RoCoF violations are characterised using the maximum RoCoF metric with threshold $r_{\max} = 1\text{Hz/s}$, which corresponds to the safe region $K^{(N)}$ introduced in the previous section, see (12). The duration considered is $\epsilon = 0.5\text{s}$ and the sensitivity of the results is examined with respect to N , taking $N = 1, 5, 20, 50, 100$. The ghost sampler uses a Gaussian proposal $N(0, \sigma^2 I)$, whose standard deviation $\sigma^2 = 10^{-3}$ has been tuned so that approximately 15% of the proposed moves are accepted, as suggested in [NRY12].

For each value of N , 10^6 disturbances u from the complement of $K^{(N)}$ were sampled after discarding an initial burn-in period. For each generator, Table I reports its probability of disconnection due to a nodal RoCoF violation. Note that generator 10 was never disconnected in our experiments and so is not shown.

N	G1	G2	G3	G4	G5	G6	G7	G8	G9
1	28.9	80.3	0.5	0.6	0.9	0.4	1.6	6.5	1.6
5	27.6	81.5	12.4	1.1	2.1	0	1.2	9.8	1.7
20	27.5	79.5	11.5	1.9	3.0	0.1	2.4	15.5	2.0
50	28.5	78.8	12.2	1.1	2.7	0.1	2.4	17.1	2.6
100	28.6	79.8	12.2	1.7	2.4	0.1	1.9	15.6	2.0

TABLE I: Conditional probabilities (in %) of nodal RoCoF violations at each generator, given that a RoCoF violation occurs. Results are shown for different time discretizations N of the interval $[0, 0.5\text{s}]$.

Despite noise due to random sampling, the estimates in Table I are consistent for $N > 5$. Recalling (9), the appropriate choice of N is also informed by the spectral properties of the matrix A . In particular, the highest frequency component of the fluctuations is the eigenfunction corresponding to the eigenvalue with largest imaginary part. Table II reports some other relevant statistics for the IEEE 39-bus system under the considered disturbance model, namely the probability p_d of multiple RoCoF violations, the average number \bar{d} of violations, and the corresponding average level \bar{L} of lost generation.

	$N = 1$	$N = 5$	$N = 20$	$N = 50$	$N = 100$
p_d	15.2%	22.4%	24.0%	25.0%	25.0%
\bar{d}	1.21	1.37	1.44	1.46	1.44
\bar{L}	596	701	735	744	736

TABLE II: Some statistics for the IEEE 39-bus system: the probability p_d of disconnecting more than one generator, the average number \bar{d} of disconnected generators and the average amount \bar{L} of lost load (in MW).

Our case study results highlight the importance of modelling the desynchronization in nodal frequency. It is clear from Table I that the majority of RoCoF violations occur at generator 2, which has a heavy-tailed disturbance model. From Fig. 4, generator 2 is connected via a relatively high susceptance line to generator 3, which has a Gaussian disturbance model. Thus RoCoF violations due to a large disturbance at the former generator are capable of inducing subsequent violations at the latter within our considered timescale. This network effect is clearly visible in Fig. 1, where a large initial disturbance at generator 2 is followed by a subsequent RoCoF violation, at around $t = 0.4\text{s}$, at generator 3. The same relationship can be seen between generators 1 and 8. These observations highlight the importance of the (reduced) network structure and line susceptances in the modelling of frequency violation patterns and system vulnerabilities.

VII. CONCLUSIONS

This work aims to provide the mathematical framework to understand how unusually large power disturbances cause frequency violations, in particular in terms of RoCoF. We describe the time evolution of the nodal frequencies as a system of coupled swing equations with a random step disturbance at time 0. A novel MCMC method is introduced, called the ghost sampler, to sample disturbances conditionally on a RoCoF violation occurring, i.e., outside the so-called “safe region”. An illustrative case study is presented, and it would be of interest to develop this further, for example using empirical probability distributions for heavy-tailed and correlated renewable forecast errors.

Future work will explore further metrics capturing frequency violations, such as the nadir. It would be natural to look also at line overloads caused by power fluctuations and complement in this way the work done in [OM17]. Lastly, we believe that the MCMC ghost sampler has potentially wide applicability beyond power systems reliability in settings where one has to sample rare events. This is particularly so in view of the fact that many of the conditions for the region K can be relaxed.

ACKNOWLEDGEMENTS

JM and JV were supported by EPSRC grant EP/P002625/1. AZ is supported by NWO Rubicon grant 680.50.1529. The authors thank Linqi Guo, Janusz Bialek, and Steven H. Low for helpful discussions on the model.

REFERENCES

- [ADF⁺05] G. Andersson, P. Donalek, R. Farmer, N. Hatziairgyriou, I. Kamwa, P. Kundur, N. Martins, J. Paserba, P. Pourbeik, J. Sanchez-Gasca, et al. Causes of the 2003 major grid black-outs in north america and europe, and recommended means to improve system dynamic performance. *IEEE transactions on Power Systems*, 20(4):1922–1928, 2005.
- [BDNL08] H. Bludszweit, J.A. Domínguez-Navarro, and A. Llombart. Statistical analysis of wind power forecast error. *IEEE Transactions on Power Systems*, 23(3):983–991, 2008.
- [BGJM11] S. Brooks, A. Gelman, G.L. Jones, and X.-L. Meng, editors. *Handbook of Markov chain Monte Carlo*. Handbooks of Modern Stat. Methods. Chapman & Hall, Boca Raton, 2011.
- [CMSR13] K. Creighton, M. McClure, R. Skillen, and A. Rogers. Increased wind generation in ireland and northern ireland and the impact on rate of change of frequency. In *Proc. of the 12th Wind Integration Workshop*, 2013.

- [DB10] F. Dörfler and F. Bullo. Synchronization of Power Networks: Network Reduction and Effective Resistance. *IFAC Proceedings*, 43(19):197–202, 2010.
- [DB13] F. Dörfler and F. Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(1):150–163, 2013.
- [GZD⁺17] S.S. Guggilam, C. Zhao, E. Dall’Anese, Y.C. Chen, and S.V. Dhople. Engineering Inertial and Primary-frequency Response for Distributed Energy Resources. *Preprint at arXiv:1706.03612*, 2017.
- [Kun94] P. Kundur. *Power System Stability And Control*. EPRI Power System Engineering Series. McGraw-Hill, 1994.
- [LZW⁺15] X. Li, X. Zhang, L. Wu, P. Lu, and S. Zhang. Transmission line overload risk assessment for power systems with wind and load-power generation correlation. *IEEE Transactions on Smart Grid*, 6(3):1233–1242, 2015.
- [MT09] S. Meyn and R.L. Tweedie. *Markov chains and stochastic stability*. CUP, Cambridge, 2nd edition, 2009.
- [NRY12] P. Neal, G. Roberts, and W.K. Yuen. Optimal scaling of random walk metropolis algorithms with discontinuous target densities. *The Annals of Applied Probability*, 22(5):1880–1927, 2012.
- [NZZ17a] T. Nesti, A. Zocca, and B. Zwart. Emergent failures and cascades in power grids: a statistical physics perspective. *Preprint at arXiv:1709.10166*, 2017.
- [NZZ17b] T. Nesti, A. Zocca, and B. Zwart. Line failure probability bounds for power grids. In *Power & Energy Society General Meeting, 2017 IEEE*, pages 1–5. IEEE, 2017.
- [OM17] A.B. Owen and Y. Maximov. Importance sampling the union of rare events with an application to power systems analysis. *Preprint at arXiv:1710.06965*, 2017.
- [PBD17] B.K. Poolla, S. Bolognani, and F. Dörfler. Optimal Placement of Virtual Inertia in Power Grids. *IEEE Transactions on Automatic Control*, 62(12):6209–6220, 2017.
- [PM17] F. Paganini and E. Mallada. Global performance metrics for synchronization of heterogeneously rated power systems: The role of machine models and inertia. In *55th Annual Allerton Conference on Communication, Control, and Computing*, pages 324–331. IEEE, 2017.
- [PMVB05] K. Purchala, L. Meeus, D. Van Dommelen, and R. Belmans. Usefulness of DC power flow for active power flow analysis. In *IEEE Power Engineering Society General Meeting*, pages 2457–2462. IEEE, 2005.
- [RR04] G.O. Roberts and J.S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probab. Surv.*, 1:20–71, 2004.
- [SBA⁺18] B. Schäfer, C. Beck, K. Aihara, D. Witthaut, and M. Timme. Non-Gaussian power grid frequency fluctuations characterized by Lévy-stable laws and superstatistics. *Nature Energy*, 3(2):119–126, 2018.
- [TBG15] E. Tegling, B. Bamieh, and D.F. Gayme. The Price of Synchrony: Evaluating the Resistive Losses in Synchronizing Power Networks. *IEEE Transactions on Control of Network Systems*, 2(3):254–266, 2015.
- [Tie94] L. Tierney. Markov chains for exploring posterior distributions. *Ann. Statist.*, 22(4):1701–1762, 1994. With discussion and a rejoinder by the author.
- [UBA14] A. Ulbig, T.S. Borsche, and G. Andersson. Impact of low rotational inertia on power system stability and operation. *IFAC Proceedings*, 19(3):7290–7297, 2014.
- [VPMB16] T.L. Vincent, K. Poolla, S. Mohagheghi, and E. Bitar. Stability guarantees for primary frequency control with randomized flexible loads. In *2016 American Control Conference (ACC)*, pages 2328–2333. IEEE, jul 2016.
- [WEBK15] W. Winter, K. Elkington, G. Bareux, and J. Kostevc. Pushing the limits: Europe’s new grid: Innovative tools to combat transmission bottlenecks and reduced inertia. *IEEE Power and Energy Magazine*, 13(1):60–74, 2015.
- [WWS14] A.J. Wood, B.F. Wollenberg, and G.B. Sheble. *Power generation, operation, and control*. John Wiley & Sons, 3rd edition, 2014.
- [ZMST11] R.D. Zimmerman, C.E. Murillo-Sánchez, and R.J. Thomas. Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE TPS*, 26(1):12–19, 2011.
- [ZTLL14] C. Zhao, U. Topcu, N. Li, and S.H. Low. Design and Stability of Load-Side Primary Frequency Control in Power Systems. *IEEE Transactions on Automatic Control*, 59(5):1177–1189, 2014.

Proof of Lemma 1: For any $r > 0$ and $\varphi \in \mathbb{S}^{n-1}$ we have $l_x^\varphi(r) > 0$ since K^c is an open set, hence x is the only point that maps to x . Then let $T_x^K(x + r_1\varphi_1) = T_x^K(x + r_2\varphi_2)$ for some $r_1, r_2 > 0$ and $\varphi_1, \varphi_2 \in \mathbb{S}^{n-1}$. This implies $x + l_x^{\varphi_1}(r_1)\varphi_1 = x + l_x^{\varphi_2}(r_2)\varphi_2$, so that $\varphi_1 = \varphi_2 = \varphi$ say, and $0 = l_x^\varphi(r_1) - l_x^\varphi(r_2) = \int_{r_2}^{r_1} \mathbf{1}_{K^c}(x + t\varphi) dt$ must hold, which is only possible if $r_1 = r_2$ (again since the set K^c is open). \square

Proof of Lemma 2: Denote with S the surface of \mathbb{S}^{n-1} . Changing (16) to polar coordinates and then using the substitution $u := l_x^\varphi(r)$ (with $du = \mathbf{1}_{K^c}(x + r\varphi) dr$), we have:

$$\begin{aligned} Q_K(x, A) &= \int_{\mathbb{R}^n} q_K(x, y) \mathbf{1}_A(y) dy \\ &= \frac{1}{S} \int_{\mathbb{S}^{n-1}} \left(\int_0^\infty q_K(x, x + r\varphi) \mathbf{1}_A(x + r\varphi) r^{n-1} dr \right) d\varphi \\ &= \frac{1}{S} \int_{\mathbb{S}^{n-1}} \int_0^\infty q(l_x^\varphi(r)\varphi) (l_x^\varphi(r))^{n-1} \mathbf{1}_{A \cap K^c}(x + r\varphi) dr d\varphi \\ &= \frac{1}{S} \int_{\mathbb{S}^{n-1}} \left(\int_0^\infty q(u\varphi) u^{n-1} \mathbf{1}_{T_x^K(A \cap K^c)}(x + u\varphi) du \right) d\varphi \\ &= \int_{\mathbb{R}^n} q(x, y) \mathbf{1}_{T_x^K(A \cap K^c)}(y) dy = Q(x, T_x^K(A \cap K^c)). \quad \square \end{aligned}$$

Definition 5. A Markov kernel P on a space S is ν -irreducible with respect to a measure ν (on S) if for every $x \in S$ and every measurable $A \subset S$ with $\nu(A) > 0$ there exists an $n \in \mathbb{N}$ such that $P^n(x, A) > 0$.

For the approximation of π integrable functions using our MH procedure, a sufficient condition to establish the LLN is the $\pi \mathbf{1}_K$ -irreducibility of Q_K on K^c (see, for example [Tie94, Corollary 2] and [MT09, Theorem 17.1.7]). Hence, to establish Theorem 4 it is enough to show the following:

Theorem 6. Let K be closed and ray-bounded. If the underlying proposal density q is strictly positive then the ghost sampling kernel Q_K is $\pi \mathbf{1}_K$ -irreducible.

Proof. Fix an arbitrary $x \in K^c$ and measurable $A \subset \mathbb{R}^n$ such that $\pi \mathbf{1}_{K^c}(A) = \pi(A \cap K^c) > 0$. Define $B = A \cap K^c \cap \{x \in \mathbb{R}^n; \pi(x) > 0\}$. Clearly $\pi(B) > 0$.

Because K^c is open we have $l_x^\varphi(r) > 0$ for all $r > 0$ and all $\varphi \in \mathbb{S}^{n-1}$. By equation (16) q_K is strictly positive on K^c . Since B must have positive Lebesgue measure we then have $Q_K(x, B) > 0$ which implies $Q_K(x, A) \geq Q_K(x, B) > 0$, so we may take $n = 1$ in Definition 5. \square