2018 IEEE Conference on Decision and Control (CDC)
Miami Beach, FL, USA, Dec. 17-19, 2018

Frequency violations from random disturbances: an MCMC approach

John Moriarty[1], Jure Vogrinc[1] and Alessandro Zocca[2]

***Abstract*— The frequency stability of power systems is increasingly challenged by various types of disturbance. In particular, the increasing penetration of renewable energy sources is increasing the variability of power generation while reducing system inertia against disturbances. In this paper we explore how this could give rise to rate of change of frequency (RoCoF) violations.**

Correlated and non-Gaussian power disturbances, such as may arise from renewable generation, have been shown to be significant in power system security analysis. We therefore introduce *ghost sampling* which, given any unconditional distribution of disturbances, efficiently produces samples conditional on a violation occurring. Our goal is to address 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?"

I. INTRODUCTION

Frequency stability is a prime concern of transmission system operators, since instability may cause machine desynchronization leading to large power outages [1]. Decreased levels of system inertia [2], together with the variability of renewable generation, can lead to large swings in the power system frequency [3]. While promising mitigations exist including participation from loads [4], [5], distributed energy resources [6] and virtual inertia [7], 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 [8].

In this paper, following a disturbance to the nodal power injections we model the rate of change of frequency (RoCoF) at each node over the primary control timescale. Using a novel *ghost sampling* method, we then approximate the statistical distribution of disturbances conditional on a RoCoF violation occurring in at least one node of the network. The sampler is based on the Metropolis-Hastings Markov Chain Monte Carlo (MCMC) method and the disturbances can have any unconditional distribution, including correlated and non-Gaussian cases. Our statistical modelling thus goes beyond that in current probabilistic power system reliability analyses (see, for example, [14]). This paper is also complementary to recent studies estimating the unconditional probability of rare events in power systems [11], [12], [13].

The ghost sampling algorithm is of independent interest for the simulation of rare events, as it is designed to overcome issues of slow mixing suffered by the standard Metropolis-Hastings MCMC algorithm in this context (see, for example, Section 1.11.2 of [15]).

[1]John 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

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

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) $\mathcal{E} \subseteq \{(i,j) : i < j\} \subset \mathcal{G} \times \mathcal{G}$. It is assumed that G is a reduced network in which each bus houses a generation unit, since passive loads can be eliminated via Kron reduction [16], [17]. Writing ω_j for the frequency at node $j \in \mathcal{G}$, the time evolution of nodal frequencies is modelled via linearized dynamics as

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

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 [18] 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 [19], [20] are made. In view of these assumptions, the so-called DC power flow approximation holds, namely

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

where $f_{i,j}$ describes the power flow on the line between nodes i and j, if present; $B_e = B_{i,j} \geq 0$ is the (effective) susceptance between nodes i and j, and θ_j denotes the phase angle at node $j \in \mathcal{G}$. 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, and $C \in \{-1, 0, 1\}^{n \times m}$ is the incidence matrix of G:

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

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

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

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

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

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

One advantage of our approach is that the random disturbances u_j are not required to be statistically independent. This is because 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 [21]. Further, the errors in renewable power forecasts have been shown to have significantly non-Gaussian distributions, such as heavy tails in the case of wind [22]. To illustrate this point the case study presented later in Section V 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 rest of the paper is organized as follows. In Section III we specify the nodal frequency dynamics and carry out step 1). Step 2) is developed in Section IV, and an illustrative case study is provided in Section V.

III. RoCoF violations for nodal frequencies

A common object of study is the *system frequency* or *center of inertia* (COI) (see, for example, [3]):

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

While this is of inherent interest, it does not fully capture RoCoF violations. 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 V. It is clear from the figure that a given RoCoF threshold can be simultaneously respected by the system frequency and violated by one or more nodal frequencies. From the practical perspective, too, it is nodal frequencies which trigger the operation of RoCoF-sensitive generator protection

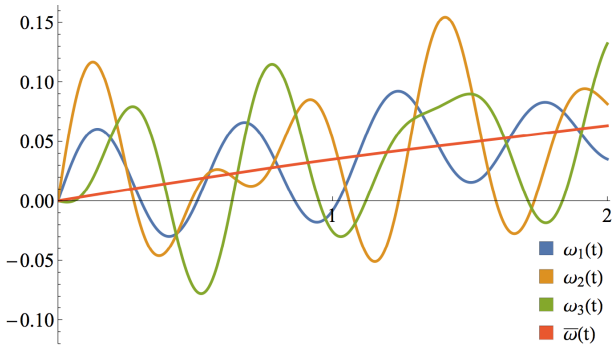

(a) Time evolution over the first 2s of the nodal frequency deviations $\omega_1(t)$, $\omega_2(t)$, and $\omega_3(t)$ and the system frequency $\bar{\omega}(t)$ (Hz) in the case study of Section V after a random disturbance.

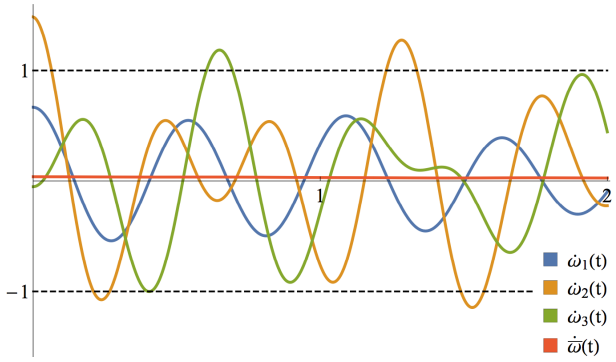

(b) Corresponding evolution over the first 2s of frequency speed deviations $\dot{\omega}_1(t)$, $\dot{\omega}_2(t)$, $\dot{\omega}_3(t)$, and $\dot{\bar{\omega}}(t)$ (Hz/s) for the same random disturbance as in Fig. 1(a). The dashed horizontal lines represent the RoCoF threshold $r_{\max} = 1$Hz/s.

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

relays which can, in turn, lead to significantly more serious frequency violations. We will therefore focus on nodal frequency violations in this paper.

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. This may be written as the homogeneous dynamical

system $\dot{x} = Ax$, with

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

whose solution is

$$x(t) = \exp(tA)\, x(0). \tag{5}$$

For a fixed node j the maximum RoCoF $|\dot{\omega}_j|$ does not in general occur at time 0, in contrast to the system frequency (see for example [14]). This can be confirmed by examining $\dot{\omega}_3$ in Fig. 1(b). 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 (comments on the choice of N are provided in the case study of Section V). Define the 'node-j safe region' $K^{(j,N)}$ by

$$K^{(j,N)} = \bigcap_{n=0}^{N} K_n^{(j,N)}, \text{ where} \tag{6}$$
$$K_n^{(j,N)} = \left\{ u \in \mathbb{R}^n : \left| \dot{\omega}_j \left(\frac{n}{N}\epsilon \right) \right| \leq r_{\max} \right\}. \tag{7}$$

It follows from (5) that $K_n^{(j,N)}$ is given by the set

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

where $\exp(tA)_j$ denotes the j-th row of the matrix $\exp(tA)$. As the intersection of two half-spaces, this set is a convex polytope. Hence $K^{(j,N)}$ and the 'all-nodes safe region' $K^{(N)}$ are also convex polytopes, where

$$K^{(N)} = \bigcap_{j \in \mathcal{G}} K^{(j,N)}. \tag{8}$$

(Note that, clearly, different thresholds $r_{\max}^{(j)}$ 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.

IV. Ghost sampling algorithm

The goal of this section is to introduce a method 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 [15], [23]. However the event K^c is in principle rare, which may cause problems of computational inefficiency. Below we describe the *ghost sampling (GS) algorithm*, which is a particular instance of the Metropolis-Hastings (MH) algorithm, more precisely of a Symmetric Random Walk Metropolis (SRWM) algorithm [23], [24]. We will show in the case study of Section V that it enables an efficient approach to statistical questions 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?".

Given a symmetric (meaning $q(x) = q(-x)$) density function $q\colon \mathbb{R}^d \to \mathbb{R}$ (termed *the proposal density*) and initiating at $X_0 \in K^c \subset \mathbb{R}^d$, the SRWM successively generates samples $X_1, X_2, \dots$ according to the procedure in Algorithm 1.

Algorithm 1: SRWM algorithm (i-th step)

Input : The i-th sample $X_i \in K^c \subset \mathbb{R}^d$

Generate a SRWM proposal Y_{i+1} distributed according to the density $q(y - X_i)dy$;
Evaluate the acceptance probability:

$$\alpha(X_i, Y_{i+1}) = \min\left(1, \frac{\pi(Y_{i+1})\mathbf{1}_{K^c}(Y_{i+1})}{\pi(X_i)\mathbf{1}_{K^c}(X_i)}\right), \tag{9}$$

interpreted as one if $\pi(X_i)\mathbf{1}_{K^c}(X_i) = 0$:
Generate a uniform random variable U on $[0, 1]$;
if $U \leq \alpha(X_i, Y_{i+1})$ **then**
  $X_{i+1} = Y_{i+1}$;
else
  $X_{i+1} = X_i$;
end
return X_{i+1}

The aim is to generate a Markov chain $X_1, X_2 \dots$ with stationary distribution equal to $\frac{\pi\mathbf{1}_{K^c}}{\pi(K^c)}$. Typical examples of symmetric proposal densities are $Y_{i+1} \sim N(X_i, \sigma^2 I_d)$ 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 (9) 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 SRWM algorithm arises when the target density is multi-modal (see Section 1.11.2 in [15]). In the present application to rare event sampling, where the set K of "common" outcomes is 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. Starting in the vicinity of one mode of the target density, "jumping across the set K" could be a rare event with respect to the proposal density q, resulting in slow mixing of the SRWM sampler.

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. (A more general case, together with proofs, can be found in our technical report [25].) 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 an SRWM proposal Y_{i+1} and denote $\varphi_i := Y_{i+1} - X_i$. Then with probability 1, $Y_{i+1} \neq X_i$ 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)). In the former case, when there

are two numbers $t_2 > t_1 > 0$ such that $X_i + t\varphi_i \in \delta K$ and $t_1 < 1$, then Y_{i+1} is either inside K or on the opposite side of K with respect to X_i and we 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 $\alpha(X_i, Z_{i+1})$ (cf. (10)) and setting $X_{i+1} = Z_{i+1}$, otherwise rejecting the proposal and setting $X_{i+1} = X_i$. Pseudocode is provided in Algorithm 2.

Algorithm 2: Ghost Sampling algorithm (i-th step)

Input : The i-th sample $X_i \in K^c \subset \mathbb{R}^d$

Generate a SRWM proposal Y_{i+1} distributed according to the density $q(y - X_i)dy$;
Calculate $\varphi_i = Y_{i+1} - X_i$;
Calculate all the intersection points (which are at most two) $T := \{t > 0 : X_i + t\varphi \in \delta K\}$;
if $T = \{t_1, t_2\}$ **and** $\min\{t_1, t_2\} < 1$ **then**
  $Z_{i+1} = Y_{i+1} + (t_2 - t_1)\varphi$;
else
  $Z_{i+1} = Y_{i+1}$;
end
Evaluate the acceptance probability:

$$\alpha(X_i, Z_{i+1}) = \min\left(1, \frac{\pi(Z_{i+1})\mathbf{1}_{K^c}(Z_{i+1})}{\pi(X_i)\mathbf{1}_{K^c}(X_i)}\right), \quad (10)$$

interpreted as one if $\pi(X_i)\mathbf{1}_{K^c}(X_i) = 0$:
Generate a uniform random variable U on $[0, 1]$;
if $U \le \alpha(X_i, Z_{i+1})$ **then**
  $X_{i+1} = Z_{i+1}$;
else
  $X_{i+1} = X_i$;
end
return X_{i+1}

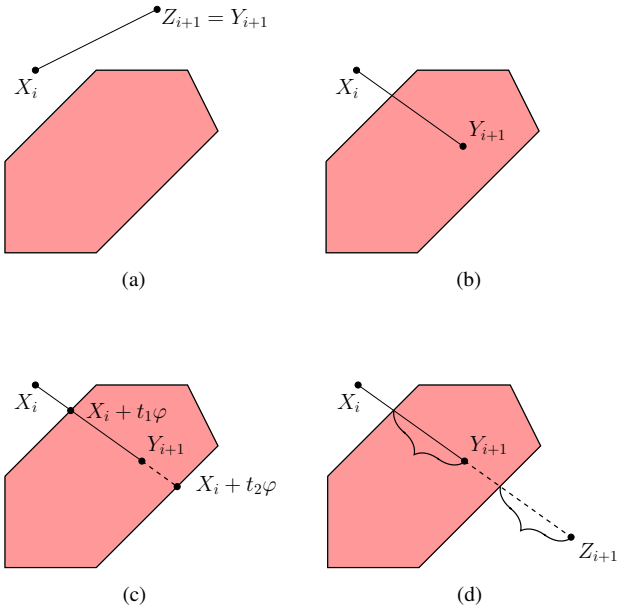

Fig. 2: Illustration of the key ideas underlying the ghost sampling algorithm.

The following example (see Figures 3 and 4) illustrates how ghost sampling improves upon the standard MH algorithm in the present context. Consider a diamond shaped region $K = \{(x, y) \in \mathbb{R}^2 : |x| + |y| < 7\}$, let π be the two-dimensional Gaussian density with zero mean and covariance matrix $\mathrm{diag}(4, 1)$ and let q be the density of a standard two-dimensional Gaussian random variable.

Starting to the left of the diamond-shaped set K, Algorithm 1 will have difficulties crossing to the right side of the diamond since, firstly, a direct move to the other side is unlikely with respect to q. Secondly, any sequence of steps through K^c towards the right side is likely to suffer rejections at the top and bottom vertices of K, where the values of π are much smaller than around the left and right vertices. The ghost sampler, however, is likely to make a direct move between the left and right sides, making it more efficient at exploring the rare event K^c.

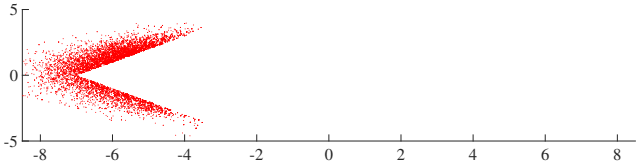

Fig. 3: Example of samples by a standard MH chain.

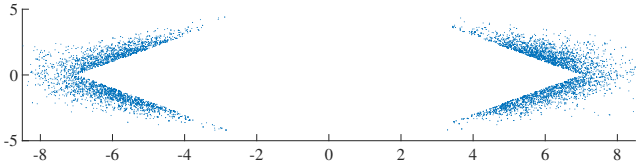

Fig. 4: Example of samples by the ghost sampling algorithm.

From the mathematical perspective it is important that the ghost sampling algorithm can be seen as a special instance of the SRWM algorithm (with some proposal density q_K that is derived from q). As a consequence, the law of large numbers (LLN) is valid for the ghost sampling algorithm, meaning that sample averages $\frac{1}{n}\sum_{i=1}^{n} f(X_i)$ for large n can provide good estimates for the actual conditional expectations

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

In the context of frequency stability in power grids, the GS algorithm can be used for networks with hundreds or possibly even thousands of nodes. If the safe region K is a known convex polytope, like the one in (8), the ghost sampler is computationally not significantly more demanding than the standard SRWM algorithm. The number of samples required to efficiently explore the state-space and to capture all the modes of rare events is of the same order as the number of generators (see Section 6 in [24]). Moreover, it has recently been shown [26] that a low-dimensional representation can often be obtained for the weather-related randomness in renewable generation using PCA techniques. Although beyond the scope of this paper, it should be possible to further reduce computational complexity by exploiting this approach.

V. Case study: IEEE 39 New England network

In this section we illustrate how inferences can be made about RoCoF violations using the ghost sampling algorithm. As a case study we consider the IEEE 39-bus New England interconnection system, which has 10 generators and 29 load nodes, see Fig. 5(a). The system parameters for our experiments are taken from the Matpower Simulation Package [27].

We consider the Kron-reduced version of the aforementioned system, which is illustrated in Fig. 5(b). The thickness of the edges in Fig. 5(b) is proportional to the equivalent susceptance between the two corresponding generator nodes.

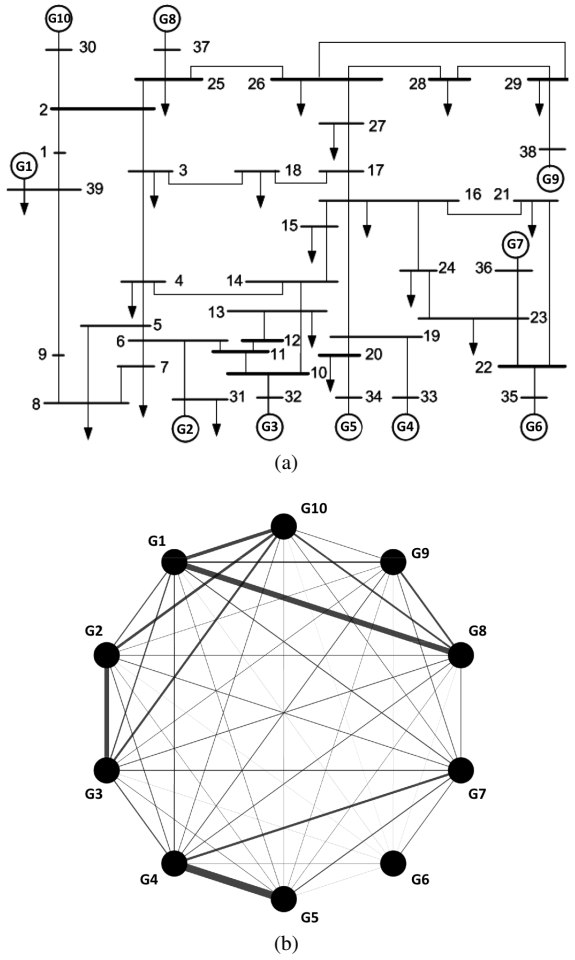

Fig. 5: (a) Line diagram of the IEEE 39-bus system and (b) the Kron-reduced version of the IEEE 39-bus system with only the 10 generator nodes.

The GS algorithm is appropriate for 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, each with standard deviation 65 times smaller than the associated generator's nominal power injection. The disturbances u_1 and u_2 are modelled by the joint density

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

We declare a violation if the RoCoF at any node exceeds the threshold $r_{\max} = 1\text{Hz/s}$, which corresponds to the safe region $K^{(N)}$ introduced in Section III, see (8). 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 GS algorithm 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 [28].

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 provides information on 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: Empirical 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 comparable for $N > 5$. Recalling (5), 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.

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

TABLE II: Empirical statistics for the IEEE 39-bus system, given that a RoCoF violation occurs: the probability p_d of multuple RoCoF violations, the average number $\overline{d}$ of violations and the corresponding average level $\overline{L}$ of disconnected generation (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. 5, 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 subsequent RoCoF violations, at around $t = 0.4s$ and

$t = 0.6s$, 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.

VI. CONCLUSIONS

This work aims to provide a 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 nadir and average RoCoF. It would be natural to look also at line overloads caused by power fluctuations and complement in this way the work done in [11]. 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, e.g., convexity.

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 L. Guo, J. Bialek, and S.H. Low for helpful discussions on the model.

REFERENCES

[1] G. Andersson *et al.*, "Causes of the 2003 major grid blackouts in north america and europe, and recommended means to improve system dynamic performance," *IEEE IEEE Trans. Power Syst.*, vol. 20, no. 4, pp. 1922–1928, 2005.

[2] 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 Energy Mag.*, vol. 13, no. 1, pp. 60–74, 2015.

[3] A. Ulbig, T. Borsche, and G. Andersson, "Impact of low rotational inertia on power system stability and operation," *IFAC Proceedings*, vol. 19, no. 3, pp. 7290–7297, 2014.

[4] C. Zhao, U. Topcu, N. Li, and S. Low, "Design and Stability of Load-Side Primary Frequency Control in Power Systems," *IEEE Trans. Autom. Control*, vol. 59, no. 5, pp. 1177–1189, 2014.

[5] T. Vincent, K. Poolla, S. Mohagheghi, and E. Bitar, "Stability guarantees for primary frequency control with randomized flexible loads," in *2016 American Control Conference (ACC)*, Boston, MA, 2016, pp. 2328–2333.

[6] S. Guggilam, C. Zhao, E. Dall'Anese, Y. Chen, and S. Dhople, "Engineering Inertial and Primary-frequency Response for Distributed Energy Resources," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Melbourne, VIC, 2017, pp. 5112-5118.

[7] B. Poolla, S. Bolognani, and F. Dörfler, "Optimal Placement of Virtual Inertia in Power Grids," *IEEE Trans. Autom. Control*, vol. 62, no. 12, pp. 6209–6220, 2017.

[8] 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.

[9] 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*, vol. 3, no. 2, pp. 119–126, 2018.

[10] E. Tegling, B. Bamieh, and D. Gayme, "The Price of Synchrony: Evaluating the Resistive Losses in Synchronizing Power Networks," *IEEE Trans. Control Netw. Syst.*, vol. 2, no. 3, pp. 254–266, 2015.

[11] A. Owen and Y. Maximov, "Importance sampling the union of rare events with an application to power systems analysis," *arXiv:1710.06965*, 2017.

[12] T. Nesti, A. Zocca, and B. Zwart, "Emergent failures and cascades in power grids: a statistical physics perspective," *Phys. Rev. Lett. 120*, 258301, 2018.

[13] ——, "Line failure probability bounds for power grids," in *IEEE Power & Energy Society General Meeting 2017*, IEEE 2017, pp. 1–5.

[14] 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*, IEEE 2017, pp. 324–331.

[15] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, Eds., *Handbook of Markov chain Monte Carlo*, ser. Handbooks of Modern Stat. Methods. Chapman & Hall, Boca Raton, 2011.

[16] F. Dörfler and F. Bullo, "Kron reduction of graphs with applications to electrical networks," *IEEE Trans. Circuits Syst. I*, vol. 60, no. 1, pp. 150–163, 2013.

[17] ——, "Synchronization of Power Networks: Network Reduction and Effective Resistance," *IFAC Proceedings*, vol. 43, no. 19, pp. 197–202, 2010.

[18] P. Kundur, *Power System Stability And Control*, ser. EPRI Power System Engineering Series. McGraw-Hill, 1994.

[19] 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*. IEEE, 2005, pp. 2457–2462.

[20] A. Wood, B. Wollenberg, and G. Sheble, *Power generation, operation, and control*, 3rd ed. John Wiley & Sons, 2014.

[21] 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 Trans. Smart Grid*, vol. 6, no. 3, pp. 1233–1242, 2015.

[22] H. Bludszuweit, J. Domínguez-Navarro, and A. Llombart, "Statistical analysis of wind power forecast error," *IEEE Trans. Power Syst.*, vol. 23, no. 3, pp. 983–991, 2008.

[23] L. Tierney, "Markov chains for exploring posterior distributions," *Ann. Statist.*, vol. 22, no. 4, pp. 1701–1762, 1994, with discussion and a rejoinder by the author.

[24] G. Roberts and J. Rosenthal, "General state space Markov chains and MCMC algorithms," *Probab. Surv.*, vol. 1, pp. 20–71, 2004.

[25] J. Moriarty, J. Vogrinc and A. Zocca, "Frequency violations from random disturbances: an MCMC approach," Technical report *arXiv:1803.08522*, 2018.

[26] J. Warrington, D. Drew, J. Lygeros, "Low-dimensional space- and time-coupled power system control policies driven by high-dimensional ensemble weather forecasts", *IEEE Control Syst. Lett.*, vol. 2, no. 1, pp. 1-6, 2018.

[27] R. Zimmerman, C. Murillo-Sánchez, and R. Thomas, "Matpower: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, 2011.

[28] P. Neal, G. Roberts, and W. Yuen, "Optimal scaling of random walk metropolis algorithms with discontinuous target densities," *The Annals of Applied Probability*, vol. 22, no. 5, pp. 1880–1927, 2012.

[29] S. Meyn and R. Tweedie, *Markov chains and stochastic stability*, 2nd ed. CUP, Cambridge, 2009.