Increasing the Resilience of Low-inertia Power Systems by Virtual Inertia and Damping

Dominic Groß, Saverio Bolognani, Bala Kameshwar Poolla, and Florian Dörfler
Automatic Control Laboratory at the Federal Institute of Technology (ETH) Zürich, Switzerland
Email: {gross,bsaverio,bpoolla,dorfler}@ethz.ch

***Abstract*—A major transition in the operation of electric power systems is the replacement of bulk generation based on conventional power plants and their synchronous generators by renewable sources connected to the grid through power electronic converters. The accompanying "loss of rotational inertia" and the fluctuations by renewable sources jeopardize the system stability, as testified by the ever-growing number of frequency incidents. As a remedy, numerous works show how virtual inertia can be emulated through various devices. In this work, we instead ask the system-level question on how to use virtual inertia to increase the resilience of low-inertia power systems. To this end, we discuss several metrics for power system robustness and show, via an insightful example, that some of these metrics are not suitable to characterize resilience. Instead, we focus on a performance metric based on system norms which accounts for network coherency as well as efficient use of control energy. We consider a linear multi-machine / multi-inverter power system model and virtual inertia devices modeled as controllable power sources with delayed feedback control actions. In practice such delays arise from signal processing, e.g., phase-locked loops, used in converter control schemes which emulate synchronous generators and their inertia. Based on this model and our performance metric we propose a computational approach to optimize the inertia allocation and ultimately increase the resilience of low-inertia power systems. A three-area power system case study is used to illustrate the results and compare different performance metrics.**

I. Introduction

The electric power system is currently undergoing a major transition towards integration of large shares of distributed generation connected via power electronic converters. Today power systems operation is based on bulk generation by synchronous machines. In contrast, future power systems will be based on renewable sources, distributed generation, and power electronics. A direct consequence of retiring synchronous generators is the loss of rotational inertia, which thus far was the main reason for the grid's stability and robustness to disturbances [1], [2]. This results in larger, and more frequent frequency deviations and jeopardizes the stability of the power grid [3], [4].

System inertia, as a global parameter, represents the capability to store and inject kinetic energy to the grid. Lower inertia is commonly associated with larger frequency fluctuations following a disturbance [5], i.e., any event that causes power imbalance in the grid: disconnection of a generator, sudden drop in power injection from a renewable (uncontrollable) source, tie line faults, grid splits, loss of loads, etc. Even as of today the deteriorating effects of low inertia levels on the system frequency and related incidents are being observed by transmission system operators worldwide [6]–[8].

To mitigate the loss of rotational inertia different technologies have been proposed to provide virtual (or synthetic) inertia. A number of control schemes have been designed in order to make power converters behave as closely as possible to synchronous machines [9], [10]. These schemes range from simple proportional-derivative droop control laws [11], to the more complex control schemes reviewed in [12] and [13] under the name of *Virtual Synchronous Generators*. All these strategies require some amount of energy storage (to provide the equivalent role as the the missing rotor kinetic energy), which could be batteries [14], super-capacitors, or flywheels.

On the other hand, specialized control schemes have been proposed for those power sources in which some kinetic energy is available, notably wind turbines [15] and diesel generators [16]. For these sources, and in particular for wind turbines interconnected to the grid via doubly fed induction generators, the power converters can be controlled in order to mimic the inertial response of a synchronous machine [17]–[20]. Interestingly, it is also possible to control these power converters so that an inertial response is induced, without relying on (possibly destabilizing) PLL measurements [21].

Given the maturity of these solutions, virtual inertia is becoming a viable technology that can be deployed throughout the grid. Given the opportunity of placing synthetic inertia in the grid, in this work we focus on the fundamental system-level question of "how to optimally use synthetic inertia". The authors in [5] observed that the detrimental effects of reduced system inertia are worsened by spatially heterogeneous inertia profiles. In other words, while the total amount of system inertia is commonly thought of as metric for robustness of the grid, the spatial distribution of inertia in the grid also plays a very crucial role.

Different authors considered this problem of optimally placing and tuning of virtual inertia controllers based on either spectral performance metrics [22], [23] or system norms [24]–[26]. In these works it is assumed that the inertia of generators can be directly modified. In contrast, we model virtual inertia devices as a feedback control that mimics the inertia and damping characteristics of synchronous generators and includes delays due to signal processing, e.g., phase locked loops, in virtual inertia devices, which are known

This research is supported by the ETH Seed Project SP-ESC 2015-07(4) of the ETH Zurich Foundation and the SNF Assistant Professor Energy Grant #160573.

to have deteriorating effects [6]. This setup allows us to consider a more realistic setting and is advantageous for optimizing system resilience via system norms. We review several performance metrics used to quantify the robustness of a power system and provide an example that highlights that three commonly used metrics (the total inertia, maximal rate of change of frequency, and smallest damping ratio) are not suitable to quantify resilience. Instead, we propose to use the $\mathcal{H}_2$ norm of the power system as measure for robustness of the power system. The $\mathcal{H}_2$ norm of a system is an integral-quadratic performance criterion that quantifies the system state excursions following a shock. The problem of optimizing system resilience via virtual inertia devices is recast as a structured $\mathcal{H}_2$ norm optimization. Whereas [23]–[25] rely on constraints which implicitly limit the peak power injection and output energy of the control devices we propose to use energy metrics which can be directly incorporated into the $\mathcal{H}_2$ norm optimization. Finally, we analyze the performance of the resulting allocations of virtual inertia for a three-area test system and several different performance metrics.

The paper is organized as follows: In Section II, we introduce the power system model. We review performance metrics used to quantify frequency stability in Section III. The main result on optimizing power system resilience via $\mathcal{H}_2$ norm optimization is presented in Section IV. The results are illustrated using a three-area test system in Section V, and the paper closes with some conclusions in Section VI.

II. Model

A. Interconnected synchronous machines

A commonly used model to assess dynamic phenomena in power systems is the swing equation. It models each generator i with two dynamic states, angle θ_i and frequency ω_i. The dynamics of a generator are assumed to be dominated by the rotational inertia, and all voltages are assumed constant; see [27], [28] for a detailed derivation. We further linearize around a steady state and assume that the mechanical input to the generator stays constant over the time scale of interest.

The differential equation describing the dynamics is then

$$m_i\dot{\omega}_i = -d_i\omega_i + p_{\mathrm{mech},i} + p_{\mathrm{el},i}\ , \tag{1}$$

where $m_i = J_i\omega_0$ is the inertia and d_i the damping coefficient of generator i. The term $p_{\mathrm{mech},i}$ represents changes in the mechanical torque on the machine, while $p_{\mathrm{el},i}$ represents changes in the electrical torque, including line flows to neighboring buses, bus power injections, and local disturbances. In the remainder we shall use the vector ω_G to denote the frequencies of the synchronous generators. We also use a small inertia time constant and damping to model each load, thereby adding dynamic states for the load buses as well. Specifically, a small inertia constant m_i at the load buses may be used to model fast initial transients in the angle and frequency of the load buses after a disturbance. Using $m_i = 0$ for the load buses results in the well-known frequency dependent load model with damping [29]. Moreover, letting both $m_i = 0$ and $d_i = 0$ for the load buses results in an implicit formulation of the frequency divider [30]. We remark that the model in [29] can be interpreted as the result of applying the singular perturbation approximation $m_i \to 0$ to each load bus [31]. This suggests that the differences between the two models become negligible for a small enough m_i at the load buses. It should also be noted that, in some cases, using the load inertia $m_i = 0$ will result in a more involved formulation of the $\mathcal{H}_2$ norm optimization problem presented in Section IV. In the remainder, we shall consider the more general case of a small load bus inertia constants $m_i > 0$. We shall revisit the case of load bus inertia constants $m_i = 0$ in Section V-E and use a numerical example to show that the difference between the two models is negligible.

The load buses and generator buses are connected via power lines, described by the graph Laplacian L (the bus susceptance matrix of the grid). Under small-signal DC power flow assumptions, the electric torque term p_{el} can then be linearized as

$$p_{\mathrm{el}} = -L\theta + p\,,$$

where the i-th element p_i of p represents the change in electric power injection at bus i.

The system dynamics can be written compactly as

$$\begin{bmatrix}\dot{\theta}\\ \dot{\omega}\end{bmatrix} = \underbrace{\begin{bmatrix}0 & I\\ -M^{-1}L & -M^{-1}D\end{bmatrix}}_{A_0}\underbrace{\begin{bmatrix}\theta\\ \omega\end{bmatrix}}_{x_0} + \underbrace{\begin{bmatrix}0\\ M^{-1}\end{bmatrix}}_{B_0}(p_{\mathrm{mech}} + p)\,, \tag{2}$$

where the diagonal matrices M and D contain the inertia and damping coefficients m_i and d_i, respectively. We denote the state of system (2), which contains all voltage angles and bus frequencies, by $x_0 = (\theta,\ \omega)$ and define the system output $\omega = Cx_0$.

B. Governors and primary droop control

The swing equation model (2) is suitable for small-signal disturbance analysis. As we investigate frequency stability, we need to extend the model of each synchronous machine with a governor model and speed droop control which is the main frequency control mechanism in the time scales of interest. Conventionally, the control action through the turbine and governor are modeled by a first-order low-pass filter [28] of the form

$$p_{\mathrm{gov},i} = -\frac{K_{\mathrm{g},i}}{T_{\mathrm{g},i}\,s+1}\omega_i\,, \tag{3}$$

with time constant $T_{\mathrm{g},i} > 0$ and control gain $K_{\mathrm{g},i} > 0$. As a state-space representation for each governor, we adopt the form

$$A_{\mathrm{g},i} = -\frac{1}{T_{\mathrm{g},i}},\ B_{\mathrm{g},i} = \frac{1}{T_{\mathrm{g},i}},\ C_{\mathrm{g},i} = -K_{\mathrm{g},i}\,, \tag{4}$$

which constitute the diagonal elements of the aggregate state space representation matrices A_{g}, B_{g}, and C_{g}. The state x_{g} of the aggregated system is a low-pass filtered version of the generator frequencies (the input of the system), while the output p_{g} is fed into the mechanical power control term p_{mech} in (2).

C. Virtual inertia as a feedback control loop

Virtual inertia devices are modeled as local feedback control loops that connect the frequency and power injection at the terminals of a converter. While there are different implementations, we adopt here the following often-termed *grid-following* implementation: Each virtual inertia device receives bus frequency measurement ω_i as input, and feeds power $p_{\mathrm{v},i} = p_{\mathrm{m},i} + p_{\mathrm{d},i}$ into the system according to the transfer functions

$$\begin{aligned} p_{\mathrm{v},i} &= \frac{\tilde{m}_i s + \tilde{d}_i}{(T_{1i}s+1)(T_{2i}s+1)}\omega_i \,, \\ &= \underbrace{\frac{\tilde{m}_i s}{(T_{1i}s+1)(T_{2i}s+1)}\omega_i}_{=p_{\mathrm{m},i}} + \underbrace{\frac{\tilde{d}_i}{(T_{1i}s+1)(T_{2i}s+1)}\omega_i}_{=p_{\mathrm{d},i}} . \end{aligned} \tag{5}$$

We call $\tilde{m}_i \geq 0$ the virtual inertia, as it reacts proportional to the derivative of the frequency, and $\tilde{d}_i \geq 0$ the virtual damping, as it reacts proportional to the frequency itself.

The transfer functions have two time constants $T_{1i} > 0$ and $T_{2i} > 0$ – one is needed for causality of this PD-control, and the other one models the time constant of the PLL. In fact, ω_i is the frequency of the voltage at bus i, which cannot be measured without a time delay.

One possible state-space realization of the controller (5) is

$$\begin{aligned} \tilde{A}_i &= \begin{bmatrix} -\frac{T_{1i}+T_{2i}}{T_{1i}T_{2i}} & -\frac{1}{T_{1i}T_{2i}} \\ 1 & 0 \end{bmatrix}, \\ \tilde{B}_i &= \begin{bmatrix} \frac{1}{T_{1i}T_{2i}} \\ 0 \end{bmatrix}, \quad \tilde{C}_i = \begin{bmatrix} \tilde{m}_i & 0 \\ 0 & \tilde{d}_i \end{bmatrix} . \end{aligned} \tag{6}$$

The controller (6) has two states: $\tilde{\omega}_i$ and $\dot{\tilde{\omega}}_i$. The state $\tilde{\omega}_i$ can be interpreted as low-pass-filtered measurement of the frequency ω_i, while the state $\dot{\tilde{\omega}}_i$ can be seen as low pass filtered derivative of ω_i. The output of the controller (6) are the electric powers $p_{\mathrm{m},i}$ and $p_{\mathrm{d},i}$ which are injected by the virtual inertia device at bus i and therefore acts via the term $p_{\mathrm{el},i}$ in (2). We finally define the aggregate representation matrices $\tilde{A}$, $\tilde{B}$, $\tilde{C}$, which have the blocks $\tilde{A}_i$, $\tilde{B}_i$, $\tilde{C}_i$, respectively, on their diagonal. We denote the state of the aggregated virtual inertia devices by $\tilde{x}$.

D. Interconnected closed-loop power system

The interconnection of the dynamical models of synchronous machines, governors, and virtual inertia devices, is schematically represented in Figure 1. In the block-diagram we have also considered a disturbance input η in the electric power injection accounting for, e.g., loss of load or generation.

III. Performance metrics and design constraints

A. Quantification of frequency stability

The amount of inertia available in the grid is commonly thought of as a direct indicator for the rate of change of frequency (RoCoF) in the instants that immediately follow a large disturbance, as depicted in upper panel of Figure 2. Moreover, as the governor response of the generators (primary frequency regulation) does not act until seconds after an incident, the RoCoF directly affects the lowest frequency reached by the grid (the frequency *nadir*, in Figure 2). A low frequency nadir can lead to the disconnection of generators, load shedding intervention, and is therefore dangerous for system stability. Moreover, a steep RoCoF itself triggers circuit breakers and safety equipment. Based on this understanding, the inertia and RoCoF are typically adopted as the main metrics to evaluate the robustness of the system in terms of frequency stability.

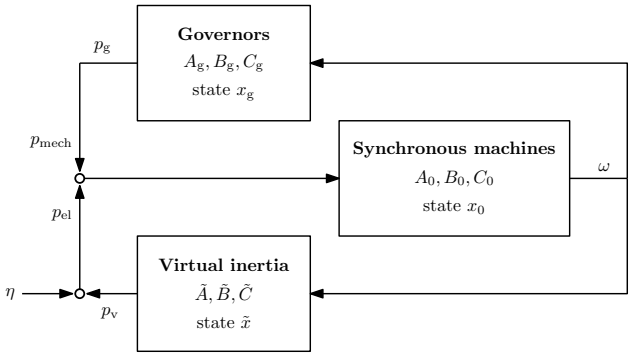

Fig. 1. Schematic representation of the interconnection of the dynamical models for synchronous machines, governors, and virtual inertia.

Another approach for the assessment of frequency stability consists of evaluating a *signal norm* for the post-fault frequency response. As depicted in Figure 2, the total area between the frequency evolution and the steady-state post-disturbance frequency can be interpreted as the *energy imbalance* caused by the disturbance, i.e., it describes how promptly and efficiently the system is restored to its nominal operating conditions after a large-scale disturbance.

It should be noted that upper panel in Figure 2 shows the post-fault behavior for a power system based on synchronous machines, today's control strategies, and faults expected in today's grid operation. However, in low-inertia power systems the inherently different properties of renewable generation may lead to different scenarios. Interestingly, signal norms are also informative in such scenarios. When a large number of renewable sources are connected to the grid, power imbalances will not only be caused by rare large events, such as the loss

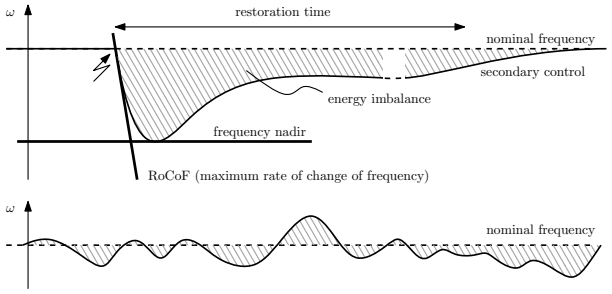

Fig. 2. Schematic representation of a power system frequency response to large disturbances, such as generator faults and network splits (upper panel) and to smaller persistent disturbances, such as fluctuating power demand or power generation from renewables (lower panel).

of a generator, but also by the temporal fluctuations of loads and uncontrollable sources such as wind and solar. The size of this latter class of events is expected to be smaller, but their occurrence is quite more frequent. For analysis purposes, one should consider a persistent disturbance on the power infeed of the buses where renewable sources are connected. In this case, the aforementioned signal norm would describe the amplification gain between these disturbances and the resulting fluctuations in grid frequency, as depicted in the bottom panel of Figure 2.

B. Performance metrics

In this section, we consider all these metrics, in order to build a deeper understanding of the effects of lower grid inertia and of the related frequency stability phenomena. Based on the model presented in Section II, we formally define a set of performance metrics that we shall use to assess the frequency stability of the grid, when subjected to disturbances.

As schematically represented in Figure 2, different metrics can be defined on the time-domain response of the system following a specified disturbance input η_0. In particular, given a negative step disturbance η_0, e.g., a sudden load increase or generation drop, at time $t = 0$, we define the following indices on the time domain evolution of bus frequencies.

- Rate of change of frequency (RoCoF):

$$|\dot{\omega}|_{\max} := \max_i \left(\max_{t \geq 0} |\dot{\omega}_i(t)| \right) . \tag{7}$$

- Frequency nadir:

$$|\underline{\omega}| := \max_i \left| \min_{t \geq 0} \omega_i(t) \right| . \tag{8}$$

Independently of the particular disturbance, the damping ratio describes how fast oscillations in the power system are vanishing. The damping ratio of a power system is defined as the smallest damping ratio of its eigenvalues λ_k. A higher numeric value hence corresponds to better performance.

- Damping ratio of a power system:

$$\zeta_{\min} := \min_k \frac{-\sigma_k}{\sqrt{(\sigma_k)^2 + (\omega_k)^2}} , \tag{9}$$

  where $\lambda_k = \sigma_k + \mathrm{i}\omega_k$ is the k-th eigenvalue of the closed-loop power system model.

The total inertia of a power system is commonly used as measure for its resilience and given by the sum of the inertia coefficients of the generators and the virtual inertia devices.

- Total inertia:

$$H_{\text{total}} := \sum_i H_i + \sum_i \tilde{H}_i, \tag{10}$$

  where $H_i = \frac{m_i \omega_0}{2\, S_{\text{rated},i}}$ and $\tilde{H}_i = \frac{\tilde{m}_i \omega_0}{2\, S_{\text{rated},i}}$ are the rated inertia constants in seconds.

For the same step disturbance as considered above, we also define the following indices to quantify the control effort that the virtual inertia devices need to exert.

- Peak virtual inertia power injection:

$$\bar{p}_{\text{v}} := \max_i \left(\max_{t \geq 0} |p_{\text{v},i}(t)| \right) . \tag{11}$$

Moreover, we define the following energy metrics, which we shall use to quantify the energy imbalance and control effort for a time τ after a fault.

- Total energy imbalance:

$$E_{\tau,\omega} := \int_0^\tau \sum_i q_i\, \omega_i^2 \, dt = \int_0^\tau \omega^\top Q\, \omega \, dt , \tag{12}$$

  where q_i are positive, possibly bus-dependent, weights on the different bus frequency deviations, and are collected in the diagonal matrix Q.
- Total virtual inertia effort:

$$E_{\tau,\text{m}} := \int_0^\tau \sum_i r_{\text{m},i}\, p_{\text{m},i}^2 \, dt = \int_0^\tau p_{\text{m}}^\top R_{\text{m}}\, p_{\text{m}} \, dt , \tag{13}$$

  where $r_{\text{m},i}$ are positive, possibly bus-dependent, weights on the control effort due to virtual inertia, and are collected in the diagonal matrix R_{m}.

- Total virtual damping effort:

$$E_{\tau,\text{d}} := \int_0^\tau \sum_i r_{\text{d},i}\, p_{\text{d},i}^2 \, dt = \int_0^\tau p_{\text{d}}^\top R_{\text{d}}\, p_{\text{d}} \, dt , \tag{14}$$

  where $r_{\text{d},i}$ are positive, possibly bus-dependent, weights on the control effort due to virtual damping, and are collected in the diagonal matrix R_{d}.

The following two system norms provide a measure of the magnitude of the system output in response to a disturbance input η (see [32]).

- $\mathcal{H}_2$ norm: The $\mathcal{H}_2$ norm can be interpreted as the energy of the response to an impulsive fault or the expected energy of the response to white noise. By defining a suitable performance output y_p (see e.g., (19)), the energy metrics (12)-(14) with $\tau = \infty$ can be directly considered in this framework as the overall output energy $\int_0^\infty \|y_p\|^2 dt$.
- $\mathcal{H}_\infty$ norm: The $\mathcal{H}_\infty$ norm corresponds to the root mean square (RMS) gain from the disturbance to the performance output y_p, which may include the frequencies ω as well as the control inputs p_{g} and p_{v}.

C. Design constraints

Virtual inertia devices have constraints on the maximum instantaneous power $\bar{p}_{\text{v}}$ that their power converters can inject in the grid. Moreover, when considering the problem of allocating virtual inertia in a grid from an economic perspective, we expect to have a budget on the maximum total size of the energy storage capacity in these devices, i.e. the control effort $E_{\tau,\text{v}} = E_{\tau,\text{m}} + E_{\tau,\text{d}}$ of the devices is limited. All these constraints need to be considered when tuning the virtual inertia gains $\tilde{m}_i$ and $\tilde{d}_i$.

In [23]–[25] it is assumed that the parameters M and D in (2) can be optimized directly and the power injections p_{v} by

virtual inertia devices are not explicitly considered. Because of this the peak power injection and output energy can only be limited implicitly by relying on a worst-case analysis of the open-loop response. According to that approach, bounds on the peak power injection can be converted into bounds on the values of $\tilde{m}_i$ and $\tilde{d}_i$ by looking at typical frequency measurements collected *before* adding the virtual inertia devices. An example from two different grids is reported in Figure 3. In both cases, it is evident that maximum frequency deviations and maximum RoCoF do not happen simultaneously. Therefore, the derivative control term $\tilde{m}_i\dot{\omega}_i$ and the proportional term $\tilde{d}_i\omega_i$ in (5) will not reach their peak value at the same time. Based on this observation, the *maximum instantaneous power* (11) can be approximately limited by magnitude constraints

$$\tilde{m}_i \leq \bar{m}_i, \qquad \tilde{d}_i \leq \bar{d}_i\,, \tag{15}$$

derived from worst-case estimates of the RoCoF $|\dot{\omega}|_{\max}$ and frequency nadir $|\underline{\omega}|$. Similarly, the system-wide *energy capacity* limits can be formulated as a budget constraint of the form

$$\sum\nolimits_i \tilde{m}_i \leq m_{\text{budget}}, \qquad \sum\nolimits_i \tilde{d}_i \leq d_{\text{budget}}. \tag{16}$$

The worst-case nature of both the *peak power* and *energy capacity* constraints can result in suboptimal solutions to the optimal inertia allocation problem. Because of the intricate dynamical coupling of the virtual inertia devices and generators through the grid, the intuition that larger control gains result in a larger peak power injection or increased use of control energy can be misleading. Specifically, by increasing some control gains the RoCoF and frequency nadir in closed-loop can be reduced which in turn can result in more efficient use of control energy.

In contrast, an important feature of modeling the virtual inertia devices as a feedback control (as presented in Section II) is that the the power injections p_{v} by virtual inertia devices are explicitly included in our model. This is a significant advantage over the methods proposed in [23]–[25] because it allows to directly consider the *closed-loop* gain from the disturbance input to the control input instead of relying on *open-loop* bounds. In particular, the control energy used to return the system to its nominal operating point can be explicitly included in the performance metrics via (13) and (14). This approach allows us to consider both resilience as well as power and energy limits in a single consistent problem formulation.

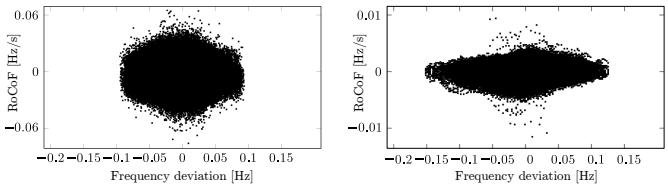

Fig. 3. Scatter plot of frequency data measurements from Ireland (left panel, courtesy of F. Milano, University College Dublin) and from continental Europe (right panel, courtesy of RTE France).

The two approaches are compared and constrasted in the simulations presented in Section V. A discussion on how to use the energy metrics (i.e. (13) and (14)) to limit the output energy and maximum instantaneous output power of the virtual inertia devices can be found in Section V-C.

D. Discussion of performance metrics

To illustrate how well different performance metrics capture the resilience of the system, we consider two allocations of virtual inertia for a three-area test system proposed in [23]. The test system consists of 6 generator buses and 5 load buses and is shown in Figure 7. To obtain a system that more closely resembles transmission grids with low-inertia (see e.g. [7]) we reduced the inertia of the generators to one third of the original value. Moreover, in comparison to the three-area test system power grids in which low levels of inertia are problematic today, e.g. the Nordic Grid [7], typically have longer transmission lines and weaker coupling. To account for this, we weakened the inductive coupling through the grid. Because bus 11 is not connected to a load or a generator we can reduce the system, without any loss of generality, to an eleven-bus (i.e. buses labeled 1-10, and 12) system via Kron reduction [33].

We consider two allocations that use the same amount of virtual inertia, i.e., $\sum_i \tilde{H}_i = 27.23\,\text{s}$, but in different locations. No additional damping is used in these examples, i.e., $\tilde{d}_i = 0$. The inertia at each node, i.e., $H_i + \tilde{H}_i$, is depicted in Figure 4 for each of the two allocations. Figure 5 shows the response of the synchronous generator frequencies ω_G to a 100MW load step at bus 7 (the topology of the system is depicted in Figure 7). It can be seen that the first allocation results in substantially larger frequency deviations compared to the second allocation. The corresponding performance metrics can be found in Table I.

While the frequency nadir $|\underline{\omega}|$, and system norms support the observation that the second allocation is more resilient, the total inertia H_{total} is identical for both allocations. Similarly, the RoCoF $|\dot{\omega}|_{\max}$ and damping ratio $\zeta_{\min}$ are almost identical for both allocations. Moreover, it can be seen that the second allocation results in a significantly smaller peak power injection $\bar{p}_{\text{v}}$. The plot of the eigenvalues of the closed-loop power system in Figure 6 does not shed much light on the

TABLE I
COMPARISON OF PERFORMANCE METRICS FOR A LOAD STEP AT NODE 7

	Allocation 1	Allocation 2
H_{total}	40.85 s	40.85 s
$\zeta_{\min}$	0.1190	0.1206
$\lvert\dot{\omega}\rvert_{\max}$ [α]	0.8149 Hz/s	0.8135 Hz/s
$\lvert\underline{\omega}\rvert$ [α]	84.8 mHz	65.1 mHz
$\mathcal{H}_2$ gain[β]	1.5337	0.6522
$\mathcal{H}_\infty$ gain[β]	0.7454	0.2782
$\bar{p}_{\text{v}}$	118.38 MW	7.0446 MW

[α] generator frequencies.
[β] generator frequencies ω_G and virtual inertia control p_{v}.

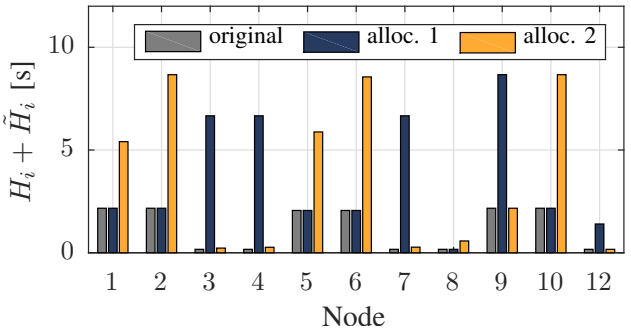

Fig. 4. Inertia of the original system and two allocations with the same amount of virtual inertia.

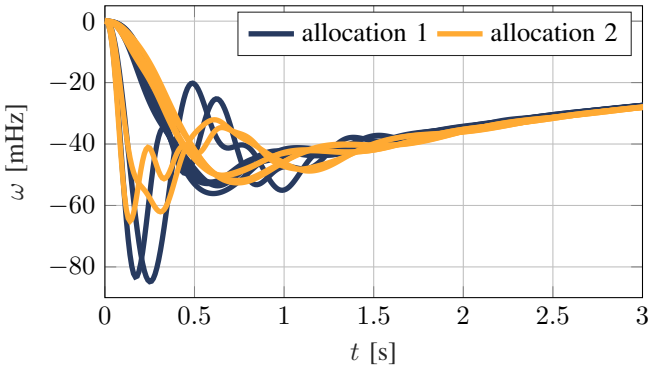

Fig. 5. Time domain plots for frequency variation at the generator nodes after 100MW load increase at bus 7.

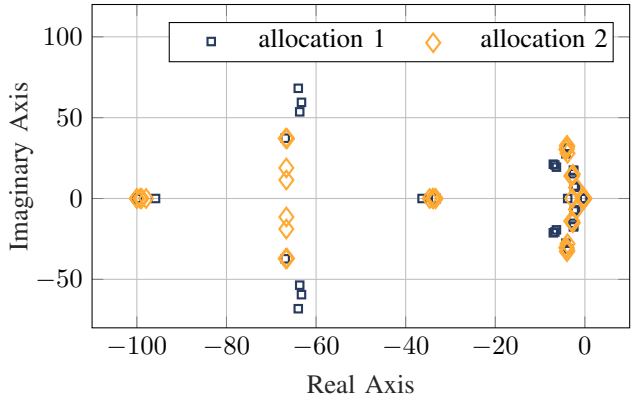

Fig. 6. Spectrum of the two allocations with the same amount of virtual inertia, same damping ratio, and identical RoCoF.

performance of the two allocations, as both achieve similar worst-case damping ratios and damping asymptotes.

Overall, we conclude that the total inertia, the RoCoF, and the damping ratio are not suitable to quantify robustness, which points towards the need to investigate other metrics. The frequency nadir characterizes frequency stability, but is only well defined for steps in the disturbance input and does not account for the cost of control energy. Moreover, computing the frequency nadir and RoCoF is computationally expensive because it requires to compute the step response from every disturbance input to every bus frequency.

In stark contrast, the $\mathcal{H}_2$ norm and $\mathcal{H}_\infty$ norm are different. These norms directly compute the gain from the disturbance inputs to a performance output, which may consider both frequency stability and efficient use of control energy, for a whole set of disturbance signals. This makes system norms a very attractive tool for analysis and optimization of power system resilience. Both norms are computationally tractable with the $\mathcal{H}_2$ norm additionally being a smooth function of the control gains and thus attractive for optimization algorithms. In the next section we show how to optimize the allocation of virtual inertia devices to minimize the $\mathcal{H}_2$ norm and improve system resilience.

IV. Control Design via $\mathcal{H}_2$ Optimization

In this section we present a computational approach to answer the question of "how and where to optimally use virtual inertia". It consists of recasting the problem as that of minimizing a system norm which takes the form of an input-output gain. Here, the disturbances η acting on the system form the set of inputs, and the frequency deviations of the synchronous machines and control effort used by the virtual inertia device, governors are the performance outputs y_p. Such an input-output gain is also referred to as the $\mathcal{H}_2$ system gain as discussed in Section III.

A. Closed-loop system and performance outputs

To this end, we consider the interconnected grid model presented in Figure 1. More precisely, we combine the state-space representations of the synchronous machines and load buses (2), the governors (4), and the virtual inertia filters (6), but keep the virtual inertia gains as explicit feedback inputs. While the frequencies ω in (2) are asymptotically stable, the angles θ are not. However, the input-output behavior of (2) with output ω can be equivalently expressed in terms of the state vector $x_\delta = (\delta, \omega)$ where $\delta_i = \theta_1 - \theta_i$ corresponds to the angle relative to the angle of bus 1. After applying a corresponding similarity transformation and removing the remaining unstable mode with zero eigenvalue corresponding to the absolute angle θ_1, we obtain a stable system $(A_\delta, B_\delta, C_\delta)$ with output $\omega = C_\delta x_\delta$. The overall system is then given by

$$\begin{bmatrix} \dot{x}_\delta \\ \dot{x}_\mathrm{g} \\ \dot{\tilde{x}} \end{bmatrix} = \underbrace{\begin{bmatrix} A_\delta & B_{\delta,\mathrm{g}} & 0 \\ B_{\mathrm{g},\delta} & A_\mathrm{g} & 0 \\ \tilde{B}C_\delta & 0 & \tilde{A} \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} x_\delta \\ x_\mathrm{g} \\ \tilde{x} \end{bmatrix}}_{x} + \underbrace{\begin{bmatrix} B_\mathrm{mv} \\ 0 \\ 0 \end{bmatrix}}_{B} u + \underbrace{\begin{bmatrix} B_\delta \Pi \\ 0 \\ 0 \end{bmatrix}}_{G} \eta, \tag{17}$$

where η is the vector of disturbances, G is the disturbance gain matrix which encodes (via Π) the location of these disturbances, and the matrix $B_\mathrm{mv} = B_\delta \otimes \begin{bmatrix} 1 & 1 \end{bmatrix}$, where $\otimes$ denotes the Kronecker product, is used to construct the control input matrix B. Moreover, the matrix Π_G is used to obtain the generator frequencies ω_G from the system state, i.e. $\omega_G = \Pi_G C_\delta x_\delta$ and the matrices $B_{\delta,\mathrm{g}} = B_\delta \Pi_G^\top C_\mathrm{g}$ and $B_{\mathrm{g},\delta} = B_\mathrm{g} \Pi_G C_\delta$ are used to model the interconnection between the synchronous machines and the governors.

The output of the virtual inertia devices fed into the interconnected system in Figure 1 is given by

$$u = \underbrace{\begin{bmatrix} 0 & 0 & \tilde{m}_1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \tilde{d}_1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \tilde{m}_n & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & \tilde{d}_n \end{bmatrix}}_{\tilde{K}} \underbrace{\begin{bmatrix} x_\delta \\ x_\mathrm{g} \\ \tilde{x}_1 \\ \vdots \\ \tilde{x}_n \end{bmatrix}}_{x}, \tag{18}$$

where $\tilde{K}$ is the matrix of virtual inertia parameters (proportional and derivative gains).

Next, we introduce a performance output y_p which contains the signals that we wish to include in the $\mathcal{H}_2$ gain analysis. With reference to (12), (13), (14), this output can be constructed as

$$y_p = \underbrace{\begin{bmatrix} Q^{\frac{1}{2}} C_\delta & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{C} \underbrace{\begin{bmatrix} x_\delta \\ x_\mathrm{g} \\ \tilde{x} \end{bmatrix}}_{x} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ R^{\frac{1}{2}} \end{bmatrix}}_{F} u, \tag{19}$$

where Q penalizes the frequency deviations ω and

$$R = \begin{bmatrix} r_{\mathrm{m},1} & 0 & \dots & 0 & 0 \\ 0 & r_{\mathrm{d},1} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & r_{\mathrm{m},n} & 0 \\ 0 & 0 & \dots & 0 & r_{\mathrm{d},n} \end{bmatrix} \tag{20}$$

is a penalty on the control effort of the virtual inertia, i.e., p_m, and control effort of the virtual damping, i.e., p_d, as discussed in Section III.

The performance output (19) is then used to formulate the following performance metric which combines the metrics (12), (13), and (14):

$$\int_0^\infty y_p^\top y_p \, dt = \int_0^\infty \omega^\top Q \omega + p_\mathrm{m}^\top R_\mathrm{m} p_\mathrm{m} + p_\mathrm{d}^\top R_\mathrm{d} p_\mathrm{d} \, dt \,. \tag{21}$$

By explicitly closing the loop between (17) and (18) this results in the following dynamic system $\mathcal{G}$:

$$\begin{aligned} \dot{x} &= (A + B\tilde{K})x + G\eta \,, \\ y_p &= (C + F\tilde{K})x \,. \end{aligned} \tag{22}$$

B. Optimization of the $\mathcal{H}_2$ norm

To compute the $\mathcal{H}_2$ norm between the disturbance input η and the performance output y_p of system (22), let the so-called observability Gramian $P_{\tilde{K}}$ denote the solution of the Lyapunov equation

$$P(A + B\tilde{K}) + (A + B\tilde{K})^\top P + C^\top C + \tilde{K}^\top F^\top F \tilde{K} = 0, \tag{23}$$

parameterized in $\tilde{K}$ for the given system matrices A, B, C, F. Based on the observability Gramian $P_{\tilde{K}}$, the norm $\|\mathcal{G}\|_2^2$ is given by [32]

$$\|\mathcal{G}\|_2^2 = \mathrm{trace}(G^\top P_{\tilde{K}} G). \tag{24}$$

Thus, the optimization problem to compute the optimal allocation with respect to the $\mathcal{H}_2$ norm $\|\mathcal{G}\|_2^2$ is obtained as

$$\begin{aligned} \min_{\tilde{K}} \quad & \mathrm{trace}(G^\top P_{\tilde{K}} G) \\ \text{s.t.} \quad & \tilde{K} \in \mathcal{S} \cap \mathcal{C}. \end{aligned} \tag{25}$$

The set $\mathcal{C}$ can be used to incorporate the peak power and energy capacity constraints introduced in Section III-C. Moreover, the set $\mathcal{S}$ is used to encode the structural constraint on $\tilde{K}$, i.e., the purely local feedback structure of the virtual inertia control in (5). We shall refer to (25) as *constrained $\mathcal{H}_2$ norm optimization* problem if the constraint $\mathcal{C}$ is present and as *$\mathcal{H}_2$ norm optimization* problem if the constraint $\mathcal{C}$ is removed. Note that evaluating the cost function requires solving the Lyapunov equation (23).

Remark 1: In general, the optimization problem (25) is non-convex and may be of very large-scale, but its structure can be exploited to obtain efficient solution methods. In particular applying the ideas of [34] to the problem at hand can be used to compute the gradient and Hessian of the cost function in (25) and speed up the computations of standard optimization methods.

V. Numerical case study

In this section, we investigate the virtual inertia placement problem for a case study described in Section V-A. Next, optimal virtual inertia allocation profiles are computed using the $\mathcal{H}_2$ norm optimization presented in Section IV. To compare the two approaches to limit the peak power and output energy discussed in Section III-C one allocation is optimized using the constraints (15) and (16) while a second allocation is optimized using the energy metrics (13) and (14). In Section V-B, we compare the performance of these allocations based on different system metrics and illustrate how the weights in the performance output can be used to trade-off resilience and energy efficiency. In Section V-D we illustrate the time domain evolution of some of the performance outputs. Finally, we conclude by discussing limitations of grid-following virtual inertia devices in Section V-F.

A. Description of the case study

The three-area test system is adapted from [23], and its topology is shown in Figure 7. The system is an aggregation of twelve buses, which can be classified as: buses with synchronous generation and buses without synchronous generation. The reactance of the step-up transformers is $0.15\,\mathrm{p.u.}$ and the line impedance is $(0.0001 + 0.001i)\,\mathrm{p.u./km}$. We do not initially constrain ourselves concerning possible sites for virtual inertia placement. Rather, we assume that each of the remaining eleven buses can be assigned a virtual inertia device with identical time constants $T_1 = 0.01$, $T_2 = 0.03$ resulting in a processing delay of roughly $100\mathrm{ms}$ as in [6]. As discussed in Section II-B a governor model is added to each generator and the gain and time delay are adjusted such that frequency nadir occurs after roughly 7s. For the load buses we assume a physical inertia time constant of $m_i = 10^{-3}$. This results

in a fast initial frequency transient at the load buses after a step disturbance that decays within 50ms, afterwards the frequency of the load buses is dominated by the frequency swings of the generators and the power injection of the virtual inertia devices. The damping coefficients of the load buses are 1.5p.u.. The test system, with the parameters according to [23], has both high levels of inertia and a uniform inertia distribution across the three areas.

Next, to study the effects of low inertia and heterogeneous inertia profiles, we reduce the inertia of the synchronous generators from approximately $H_i = 6.5\text{s}$ to $H_i = 1\text{s}$ for bus 5, 6, and 7 and to $H_i = 4\text{s}$ for bus 1, 2, and 9. This results in a total inertia $H_{\text{total}} = 15\text{s}$ and a ratio between kinetic energy and power production which is approximately half of that found today in low-inertia conditions in the Nordic Grid [7]. Moreover, to account for the loss of primary control of rotational synchronous generation, the gain of the speed drop control at the buses 5,6, and 7 is reduced to one sixth of the original value. For the $\mathcal{H}_2$-based optimization we consider a scenario in which the disturbance input η is equally likely to act on the load buses (i.e., the buses labeled 3, 4, 7, 8, 12), i.e., the matrix Π in the system model (17) maps each entry of $\eta \in \mathbb{R}^5$ to one of the five load buses.

We assume an inertia budget $m_{\text{budget}} = 133.7856\text{MW}\,\text{s}^2$, which, if fully used, restores the ratio between kinetic energy storage and power production to typical levels of the Nordic Grid without wind production. The budget constraint on the damping coefficients is assumed to be $d_{\text{budget}} = 37.5\text{MW}\,\text{s}$, this corresponds to fully restoring the damping to its previous level. Moreover, we simulate load increases of up to 100MW at every load bus to obtain the worst-case RoCoF and frequency nadir to construct the magnitude constraints discussed in Section III-C. Assuming that the peak virtual inertia response to these faults should not exceed 100MW and the virtual damping response should not exceed 10MW, the open-loop worst-case analysis reported in Section III-C yields the bounds $\bar{m}_i = 13.605\text{MW}\,\text{s}^2$ and $\bar{d}_i = 6.5\text{MW}\,\text{s}$ for each bus. We consider a cost function that penalizes the generator frequency deviations ω_G, and the virtual inertia and virtual damping power injections p_{m} and p_{d}. In particular, for ω_G in Hz we use the weight $q_i = 1$ in (12) and vary r_{m} and r_{d} to achieve different trade-offs between system robustness and the peak power and energy storage requirements of the virtual inertia devices.

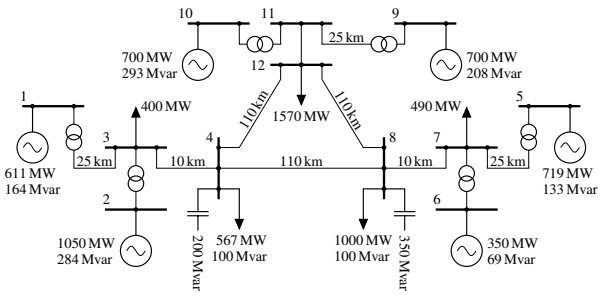

Fig. 7. Topology of the twelve-bus test system with six generators and five load buses.

B. Performance of optimal allocations

In this section we compare the performance of the original system, i.e., before reducing the inertia and primary control gain, the low-inertia system without virtual inertia or virtual damping, and the low-inertia system with two allocations obtained by using the algorithm presented in Section IV. To illustrate the effect of constraints on the parameters $\tilde{m}_i$ and $\tilde{d}_i$ we solved two variants of the optimization problem (25): the first allocation is optimized using the constraint $\mathcal{C}$ which encodes both the budget and capacity constraints presented in Section III-C, but without penalizing the control effort, i.e., $r_{\text{m}} = r_{\text{d}} = 0$. For the second allocation we removed the constraint $\mathcal{C}$ and used the penalties $r_{\text{m}} = 0.0083$ and $r_{\text{d}} = 0.08$.

The resulting inertia and damping allocations at each node are depicted in Figure 8, respectively. Note that the lower panel only contains machine damping, load damping, and virtual damping but does not include the speed droop control gains which have been lowered for the low inertia test case.

The penalties $r_{\text{m}} = 0.0083$ and $r_{\text{d}} = 0.08$ on the virtual inertia and virtual damping control effort in the $\mathcal{H}_2$ optimization were chosen to yield a control effort (specifically, a peak power injection $\bar{p}_{\text{v}}$) which is comparable to the constrained $\mathcal{H}_2$

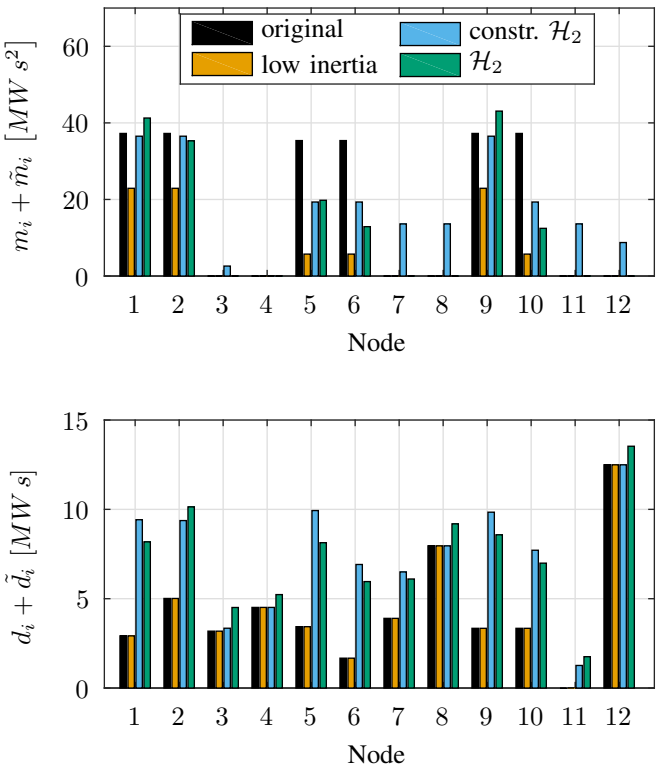

Fig. 8. Damping and inertia of the original three-area test system with high inertia, the test system with low inertia, as well as for the low inertia test system with virtual inertia allocations obtained via constrained and unconstrained $\mathcal{H}_2$ optimization.

optimization. The tuning procedure for these two parameters is discussed in detail in Section V-C.

Observe that the results of the $\mathcal{H}_2$ optimization predominantly allocates virtual inertia at the peripheral generation buses of the grid (see Figure 7 and Figure 8). This allocation is qualitatively similar to the result in [23]. Our interpretation is that virtual inertia is predominantly allocated at the generator nodes for two reasons: (i) only the frequency deviations at the generator nodes are penalized in the optimization, (ii) placing grid-following virtual inertia devices at buses without physical inertia, which acts as a low pass filter for the frequency measured by the virtual inertia devices, can result in large oscillations (see Section V-F).

For the purpose of comparing the efficacy of these different allocations vis-à-vis the original system, we compare in this subsection the performance metrics that have been introduce in Section III. In particular, we consider the total physical inertia, total virtual inertia, total virtual damping, damping ratio, and $\mathcal{H}_2$ norm. These metrics are independent of specific faults, and, as discussed in Section III, the $\mathcal{H}_2$ norm can effectively capture both system-wide performance as well as the need to use control energy efficiently. Additionally, we evaluate the response when the test system is subjected to a load increase of $100\mathrm{MW}$ at the load bus 12 in Figure 7 to test the resilience of these allocations. For this localized step disturbance we inspect the maximum RoCoF (rate of change of frequency), frequency nadir, peak power injection $\bar{p}_{\mathrm{v}}$, as well as the total energy imbalance $E_{\tau,\omega}$, and total virtual inertia and virtual damping control effort ($E_{\tau,\mathrm{m}}$ and $E_{\tau,\mathrm{d}}$) for $\tau = 30\mathrm{s}$.

A comparison of the aforementioned metrics across different allocations is given in Table II. Note that the constrained $\mathcal{H}_2$ allocation is optimized using only the generator bus frequencies as performance output (denoted by "$\mathcal{H}_2$ freq.") and the unconstrained $\mathcal{H}_2$ allocation is optimized for a performance output that contains both the generator bus frequencies and control inputs (denoted by "$\mathcal{H}_2$ sys.").

The following inferences can be drawn from these results:

- The RoCoF and frequency nadir are significantly increased by reducing the inertia and primary control gains of the original system. By allocating virtual inertia and virtual damping both the RoCoF and frequency nadir are significantly improved. Due to the delay in the response of the grid-following inverters the RoCoF is not reduced to the original level, even after restoring the total inertia of the system. The frequency nadir is improved compared to the original system because the virtual damping acts much faster than the turbine and speed drop control.
- The magnitude constraints derived from worst-case open-loop estimates ensure that the bounds on the peak power are satisfied, i.e. the peak virtual inertia power injection is below $100\mathrm{MW}$, but are very conservative. Removing the constraints and using the penalties r_{m} and r_{d} on the control effort results in an allocation with similar performance, i.e., slightly lower RoCoF and slightly higher frequency nadir, but more efficient use of control energy, i.e., lower control effort by the virtual inertia and damping devices. As discussed in Section III-C the main reason for these differences is that the penalty on the control effort is evaluated for the closed-loop system while the magnitude constraints are derived from worst-case open-loop estimates.
- At the same time, it can be seen that the intuition that the control effort is proportional to the budget $\bar{d}$ (see Section III-C) can be be misleading. In particular, the constrained and unconstrained $\mathcal{H}_2$ allocations utilize approximately the same amount of virtual damping but the control effort E_{d} used by the virtual damping devices in

TABLE II
COMPARISON OF PERFORMANCE METRICS FOR A LOAD STEP AT NODE 12

metric	original	low inertia	constr. $\mathcal{H}_2$	$\mathcal{H}_2$
$\sum_i m_i$	$219.76\,\mathrm{MW\,s^2}$	$85.98\,\mathrm{MW\,s^2}$	$85.98\,\mathrm{MW\,s^2}$	$85.98\,\mathrm{MW\,s^2}$
$\sum_i \tilde{m}_i$	-	-	$133.78\,\mathrm{MW\,s^2}$	$78.90\,\mathrm{MW\,s^2}$
$\sum_i \tilde{d}_i$	-	-	$37.50\,\mathrm{MW\,s}$	$36.52\,\mathrm{MW\,s}$
$\zeta_{\min}$	0.0048	0.0093	0.0311	0.0746
$\lvert\dot{\omega}\rvert_{\max}$ $^{\alpha}$	$0.219\,\frac{\mathrm{Hz}}{\mathrm{s}}$	$1.154\,\frac{\mathrm{Hz}}{\mathrm{s}}$	$0.368\,\frac{\mathrm{Hz}}{\mathrm{s}}$	$0.344\,\frac{\mathrm{Hz}}{\mathrm{s}}$
$\lvert\underline{\omega}\rvert$ $^{\alpha}$	$202\,\mathrm{mHz}$	$264\,\mathrm{mHz}$	$147\,\mathrm{mHz}$	$154\,\mathrm{mHz}$
$\bar{p}_{\mathrm{v}}$	-	-	$26.69\,\mathrm{MW}$	$26.76\,\mathrm{MW}$
$\mathcal{H}_2$ sys.$^{\beta}$	-	-	1.1257	0.0718
$\mathcal{H}_2$ freq.$^{\alpha}$	0.0734	0.1603	0.0405	0.0534
$E_{\tau,\omega}$ $^{\alpha,\gamma}$	3.279	5.689	2688	2782
$E_{\tau,\mathrm{m}}$ $^{\gamma}$	-	-	$7.30 \cdot 10^5$	$6.25 \cdot 10^5$
$E_{\tau,\mathrm{d}}$ $^{\gamma}$	-	-	$3.54 \cdot 10^6$	$2.70 \cdot 10^6$

$^{\alpha}$generator frequencies.
$^{\beta}$generator frequencies and virtual inertia and damping control.
$^{\gamma}$evaluated for $\tau = 30\,\mathrm{s}$, $r_{\mathrm{m},i} = 1$, $r_{\mathrm{d},i} = 1$, and $q_i = 1$.

the constrained $\mathcal{H}_2$ allocation is much higher.

- In comparison to the original system and to the low inertia system without virtual inertia or damping, the two optimized allocations reduce the $\mathcal{H}_2$ norm from the disturbance inputs to the frequency deviations at the generator buses. However, this choice of performance output does not account for efficient use of control energy. Evaluating the $\mathcal{H}_2$ norm with a performance output that contains the generator frequencies as well as the power injection by virtual inertia and damping results in a much larger $\mathcal{H}_2$ norm for the allocation that was optimized using the budget and magnitude constraints.
- It can be seen that the damping ratio does not correlate well with the RoCoF or the frequency nadir. This underlines that the spectrum itself is not very insightful.

The unconstrained $\mathcal{H}_2$ norm optimization results in an allocation that is similar to the original inertia profile of the grid before we reduced the inertia for this test case. Moreover, virtual damping devices are placed at the buses where we reduced the gain of speed droop control. Given that the test system was designed around the original inertia and damping profile, it is reasonable to expect that the optimal allocations are similar and cannot substantially outperform this allocation when limited to a comparable output power and energy capacity. Nonetheless, as discussed above, the comparably fast response time of the virtual damping results in an allocation that decreases the frequency nadir.

C. Tuning of the weights in the $\mathcal{H}_2$ optimization

The simulations in Section V-B showed that including the budget and magnitude constraints derived as in Section III-C (i.e., based on an open-loop analysis of the system), is not an effective tool to account for the use of control energy. Instead, it is much more effective to consider a system $\mathcal{H}_2$ norm that directly includes penalties on the control effort.

In this section, we show how these penalty terms, and specifically the weights r_m and r_d, can be tuned in order to achieve the desired peak power injection or energy storage capacity of the virtual inertia devices. To illustrate this fact, we solved the unconstrained $\mathcal{H}_2$ norm optimization for

- $r_\mathrm{d} = 0.08$ and r_m ranging from 0.004 to 0.02,
- $r_\mathrm{m} = 0.0083$ and r_d ranging from 0.04 to 0.3.

The resulting peak power, RoCoF, frequency nadir and virtual damping control effort are shown in Figure 9. We observe that increasing the penalty r_m reduces the peak power $\bar{p}_\mathrm{v}$ but increases the RoCoF $|\dot{\omega}|_{\max}$, at the same time the frequency nadir $|\underline{\omega}|$ does not change. In contrast, increasing the penalty r_d reduces the virtual damping control effort $\int_{t=0}^{\tau} p_\mathrm{d}^2\, dt$ and increases the frequency nadir $|\underline{\omega}|$, but the RoCoF $|\dot{\omega}|_{\max}$ does not change. In other words, the penalty r_m can be used to limit the short term peak power injections due to virtual inertia and the penalty r_d can be used to limit the energy use due to virtual damping on longer time scales.

D. Simulation of a load step

In this subsection, we present a simulation which enables a better understanding of the post-fault system behavior.

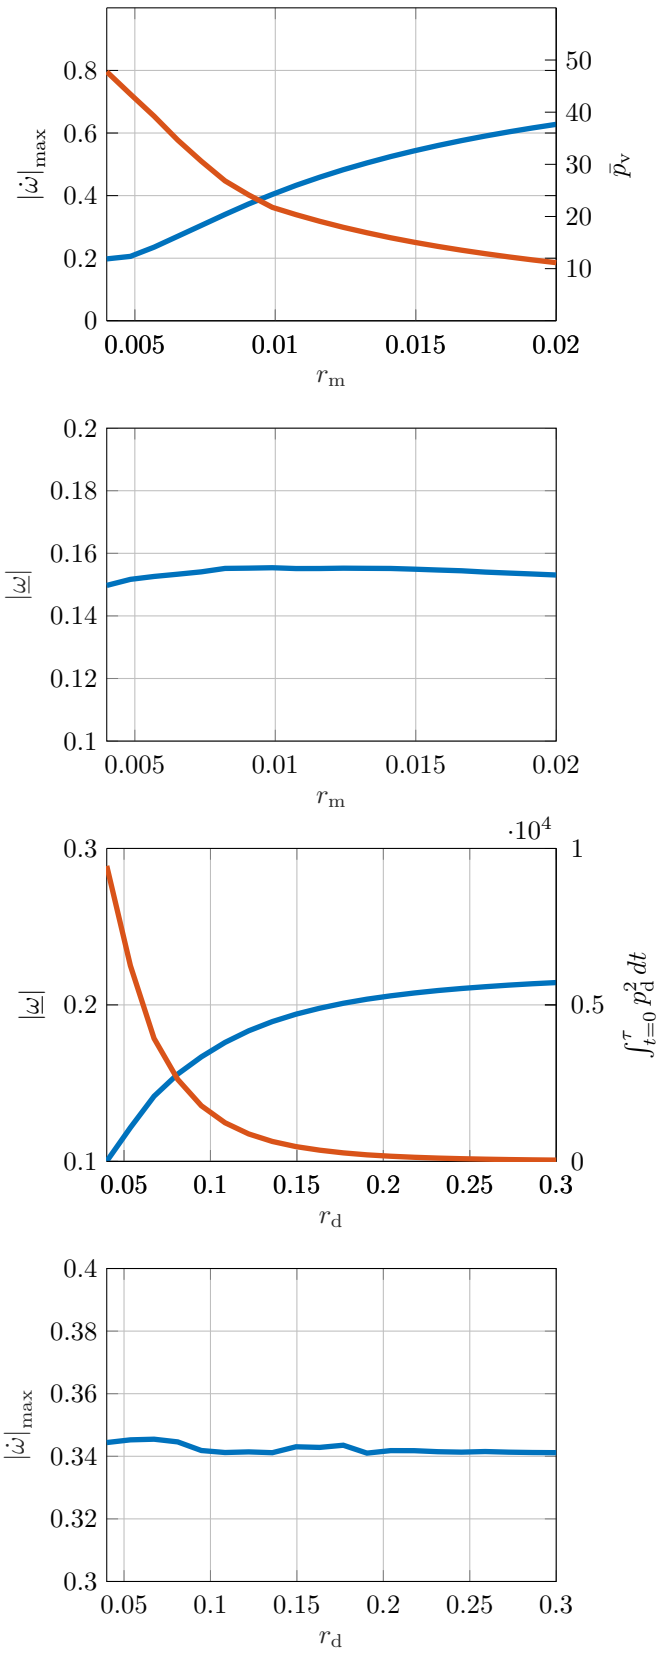

Fig. 9. Trade-off between frequency nadir, peak virtual inertia power injection, virtual damping control effort, and RoCoF, for different penalties r_m and r_d.

Specifically, we consider a load increase (step) of 100 MW at bus 12. The time-domain plots in Figure 10 suggest that all generators experience a significantly improved transient frequency response behavior for the unconstrained $\mathcal{H}_2$ allocation. The control effort displayed in Figure 10 shows a relatively large peak power due to the virtual inertia, in contrast the peak power due to virtual damping is smaller but the power injection is sustained over a longer time horizon. We observe that the virtual inertia control dominates the initial response to a disturbance and, in practice, will be limited by the peak power capabilities of the virtual inertia device. In contrast, the virtual damping dominates the response for $t > 5\,\mathrm{s}$ and, in practice, would mostly be limited by the energy storage capacity of the virtual inertia device.

E. Load model comparison

As discussed in Section II our model uses a small inertia constant, i.e. $m_i = 10^{-3}$, at the load buses. To better understand the consequences of this choice we simulated a load increase (step) of 100 MW at bus 12 of the low-inertia system without virtual inertia or virtual damping, i.e. $\tilde{m}_i = 0$ and $\tilde{d}_i = 0$, and load bus inertia constants $m_i = 10^{-3}$ as well as $m_i = 0$. The time-domain plots in Figure 11 show the resulting frequency estimates by the PLLs, i.e. the signals that would be used to implement virtual inertia and damping. Both time-domain simulations show a prominent dip in the estimated frequency of the load bus affected by the disturbance for $t < 0.05\,\mathrm{s}$. The simulations confirm the expectation that the differences between the models are negligible. Finally, we solved the optimization problem (25) without constraint $\mathcal{C}$ and with input penalties as in the previous subsections for the case $m_i = 0$. The resulting allocation is almost identical to the one obtained with $m_i = 10^{-3}$ and time-domain simulations of faults, i.e. Figure 10, show nearly identical results. Overall, as predicted by singular perturbation analysis [31], this confirms that using $m_i = 0$ in place of the small inertia constant $m_i = 10^{-3}$ results in a small change in the system response.

F. Limitations of grid-following virtual inertia devices

We observe that no virtual inertia is allocated at the load buses by the unconstrained $\mathcal{H}_2$ norm optimization while the constrained $\mathcal{H}_2$ norm allocates some virtual inertia at load buses. However, the latter allocation results in a large $\mathcal{H}_2$ norm once the control effort is explicitly considered in the performance outputs. This suggests that allocating virtual inertia at the load buses of the three-area test system may not increase resilience. In particular, when simulating impulsive disturbances (e.g. line faults) we observe large oscillations of the power infeed of virtual inertia devices at load buses.

We conjecture that the root cause for these oscillations is that the virtual inertia devices in this work are grid-following, i.e., heavily rely on measuring the local frequency at their bus. Directly after the fault a large but very brief frequency transient occurs at the load bus and the virtual inertia device reacts with a large but delayed power injection. After the initial frequency transient has decayed the frequency at the load bus is dominated by this large power injection of the virtual inertia device. In other words, at this point the virtual

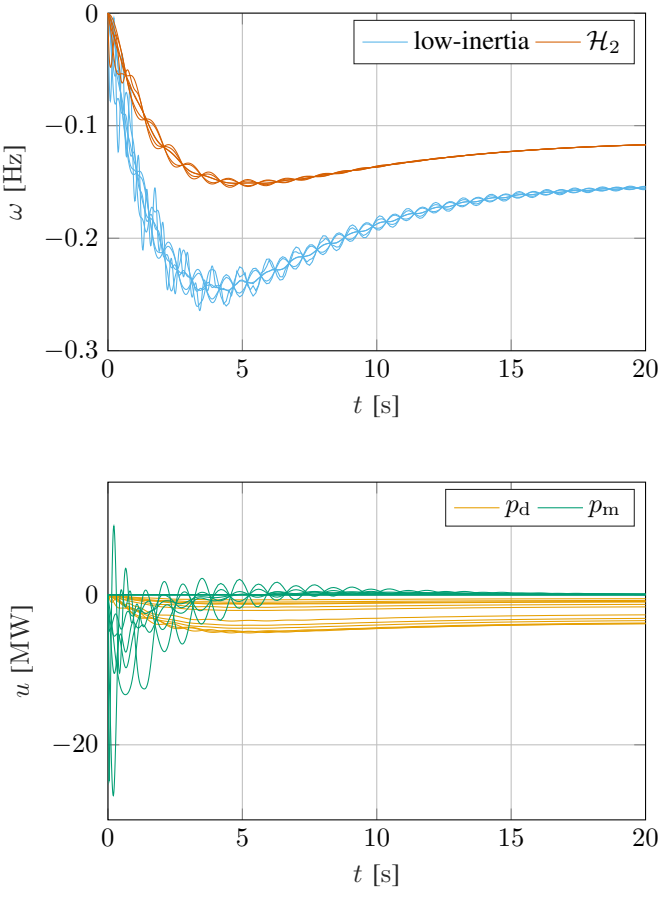

Fig. 10. Time domain simulation for the frequency variation at the generator nodes and virtual inertia / damping infeed after a load step at load bus 12 and the low-inertia system as well as the allocation obtained via unconstrained $\mathcal{H}_2$ optimization.

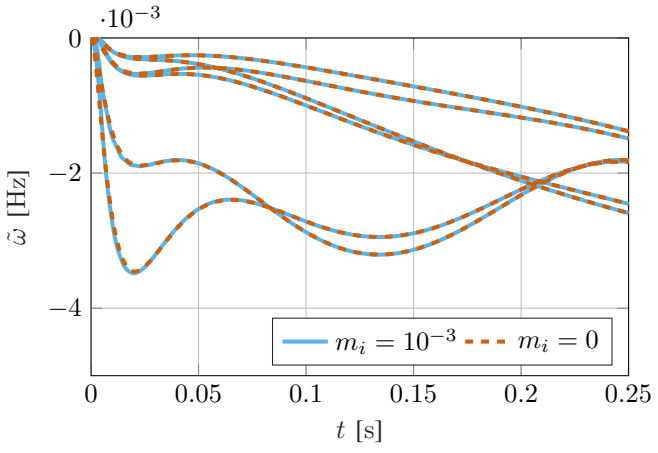

Fig. 11. Time domain simulation for the frequency estimates by the PLLs at the load nodes after a load step at load bus 12 for the low-inertia system and load bus inertia constant $m_i = 10^{-3}$ (solid blue lines) as well as load bus inertia constant $m_i = 0$ (dashed red lines).

inertia device computes its current power injection based on a slightly delayed measurement of its power injection. This poorly damped feedback interconnection via the delayed device response and fast bus dynamics can result in large frequency swings and power oscillations. This points towards the need to investigate the use of grid-forming inverters to provide virtual inertia and damping.

VI. Summary and Conclusions

In this paper we considered a linear multi-machine power system model, discussed several performance metrics for low-inertia power systems, and studied the optimal placement and tuning of control devices providing virtual inertia and damping. We provided an example that showed that certain metrics, such as the total inertia and the damping ratio, which are commonly associated with the frequency stability of a power system may not be suitable to characterize the system resilience. Instead, we formalized the problem of "how and where to optimally place virtual inertia" based on the control-theoretic notion of an $\mathcal{H}_2$ system norm which directly characterizes the gain from a disturbance input to a performance output. Based on this performance metric, we proposed a computational approach to optimize the inertia allocation and used a low-inertia three-area test case to illustrate the results and demonstrate how to adjust parameters of the optimization procedure to trade-off control performance and use of control effort.

VII. Acknowledgements

The authors wish to thank Theodor Borsche for providing the initial model for the three area case study.

References

[1] P. Tielens and D. V. Hertem, "The relevance of inertia in power systems," *Renewable and Sustainable Energy Reviews*, vol. 55, pp. 999–1009, 2016.

[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 and Energy Magazine*, vol. 13, no. 1, pp. 60–74, 2015.

[3] J. Slootweg and W. Kling, "Impacts of distributed generation on power system transient stability," in *Proc. IEEE Power Engineering Society Summer Meeting*, 2002, pp. 862–867.

[4] G. Lalor, J. Ritchie, S. Rourke, D. Flynn, and M. O'Malley, "Dynamic frequency control with increasing wind generation," in *IEEE Power Engineering Society General Meeting*, 2004, pp. 1715–1720.

[5] A. Ulbig, T. S. Borsche, and G. Andersson, "Impact of low rotational inertia on power system stability and operation," in *Proc. 19th IFAC World Congress*, 2014, pp. 7290–7297.

[6] RG-CE System Protection & Dynamics Sub Group, "Frequency stability evaluation criteria for the synchronous zone of continental europe," ENTSO-E, Tech. Rep., 2016.

[7] Svenska kraftnät, Statnett, Fingrid and Energinet.dk, "Challenges and opportunities for the nordic power system," Tech. Rep., 2016.

[8] AEMO, "Update Report - Black System Event in South Australia on 28 September 2016," Tech. Rep., 2016.

[9] Q.-C. Zhong and G. Weiss, "Synchronverters: Inverters that mimic synchronous generators," *IEEE Trans. on Industrial Electronics*, vol. 58, no. 4, pp. 1259–1267, 2011.

[10] T. Jouini, C. Arghir, and F. Dörfler, "Grid-friendly matching of synchronous machines by tapping into the dc storage," in *Proc. 6th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, 2016, pp. 192–197.

[11] N. Soni, S. Doolla, and M. C. Chandorkar, "Improvement of transient response in microgrids using virtual inertia," *IEEE Trans. on Power Delivery*, vol. 28, no. 3, pp. 1830–1838, 2013.

[12] H. Bevrani, T. Ise, and Y. Miura, "Virtual synchronous generators: A survey and new perspectives," *Intl. Journal of Electrical Power & Energy Systems*, vol. 54, pp. 244–254, 2014.

[13] S. D'Arco and J. A. Suul, "Virtual synchronous machines – classification of implementations and analysis of equivalence to droop controllers for microgrids," in *Proc. IEEE Powertech*, 2013.

[14] T. Vu Van, K. Visscher, J. Diaz, V. Karapanos, A. Woyte, M. Albu, J. Bozelie, T. Loix, and D. Federenciuc, "Virtual synchronous generator: An element of future grids," in *IEEE PES Innovative Smart Grid Technologies Conf. Europe*, 2010.

[15] J. Morren, S. de Haan, W. Kling, and J. Ferreira, "Wind turbines emulating inertia and supporting primary frequency control," *IEEE Trans. on Power Systems*, vol. 21, no. 1, pp. 433–434, 2006.

[16] M. Torres and L. A. Lopes, "Virtual synchronous generator control in autonomous wind-diesel power systems," in *Proc. IEEE Electrical Power & Energy Conference*, 2009.

[17] F. Hughes, O. Anaya-Lara, N. Jenkins, and G. Strbac, "Control of DFIG-based wind generation for power network support," *IEEE Trans. on Power Systems*, vol. 20, no. 4, pp. 1958–1966, 2005.

[18] O. Anaya-Lara, F. Hughes, N. Jenkins, and G. Strbac, "Contribution of DFIG-based wind farms to power system short-term frequency regulation," *IEE Proceedings - Generation, Transmission and Distribution*, vol. 153, no. 2, pp. 164–170, 2006.

[19] J. Ekanayake, L. Holdsworth, and N. Jenkins, "Control of DFIG wind turbines," *Power Engineer*, vol. 17, no. 1, pp. 28–32, 2003.

[20] J. Ekanayake and N. Jenkins, "Comparison of the response of doubly fed and fixed-speed induction generator wind turbines to changes in network frequency," *IEEE Trans. on Energy Conversion*, vol. 19, no. 4, pp. 800–802, 2004.

[21] S. Wang, J. Hu, and X. Yuan, "Virtual synchronous control for grid-connected DFIG-based wind turbines," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 3, no. 4, pp. 932–944, 2015.

[22] E. Rakhshani, D. Remon, A. M. Cantarellas, and P. Rodriguez, "Analysis of derivative control based virtual inertia in multi-area high-voltage direct current interconnected power systems," *IET Generation, Transmission & Distribution*, vol. 10, no. 6, pp. 1458–1469, 2016.

[23] T. S. Borsche, T. Liu, and D. J. Hill, "Effects of rotational inertia on power system damping and frequency transients," in *54th IEEE Conference on Decision and Control*, 2015, pp. 5940–5946.

[24] A. Mešanović, U. Münz, and C. Heyde, "Comparison of H-∞, H-2, and pole optimization for power system oscillation damping with remote renewable generation," in *IFAC Workshop on Control of Transmission and Distribution Smart Grids*, 2016, pp. 103–108.

[25] B. K. Poolla, S. Bolognani, and F. Dörfler, "Placing rotational inertia in power grids," *arXiv preprint arXiv:1510.01497*, 2015.

[26] M. Pirani, E. Hashemi, B. Fidan, and J. W. Simpson-Porco, "H-∞ robustness in mechanical and power networks," *Submitted.*, 2016.

[27] P. W. Sauer and M. Pai, *Power system dynamics and stability*. Prentice Hall, 1997.

[28] J. Machowski, J. Bialek, and J. Bumby, *Power system dynamics: stability and control*. John Wiley & Sons, 2011.

[29] A. R. Bergen and D. J. Hill, "A structure preserving model for power system stability analysis," *IEEE Transactions on Power Apparatus and Systems*, vol. 100, no. 1, pp. 25–35, 1981.

[30] F. Milano and A. Ortega, "Frequency divider," *IEEE Transactions on Power Systems*, vol. 32, no. 2, pp. 1493–1501, 2017.

[31] H. K. Khalil, *Nonlinear systems*, 3rd ed. Prentice hall New Jersey, 2002.

[32] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice-Hall, 1996.

[33] F. Dörfler and F. Bullo, "Kron reduction of graphs with applications to electrical networks," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, no. 1, pp. 150–163, 2013.

[34] T. Rautert and E. W. Sachs, "Computational design of optimal output feedback controllers," *SIAM Journal on Optimization*, vol. 7, no. 3, pp. 837–852, 1997.