

Nonlinear stability measures of synchronised states in a power-grid model

Tom S. Eaves^{1*}

^{1*}School of Science and Engineering, University of Dundee, Fulton Building, Dundee, DD1 4HN, UK.

Corresponding author(s). E-mail(s): teaves001@dundee.ac.uk;

Abstract

Nonlinear stability of desired dynamics in multi-stable systems (systems with more than one attracting state) depends on the shape and size of its basins of attraction. ‘Basin stability’ estimates the volume of a state’s basin of attraction and estimates the probability that a random initial condition evolves towards the state. If properties of the random initial conditions used are analysed, then basin stability can also provide estimates of the shape of the basin of attraction, but which are coarse-grained and lack details of small-scale features of its boundary. The closest approach of the basin boundary to the state can be computed via an optimisation procedure, providing minimum perturbation amplitudes to leave the desired region. Minimal disturbances are missed by basin stability estimations (by two orders of magnitude in perturbation energy for transition to fluid turbulence) and so offer a complimentary nonlinear stability measure to basin stability.

Minimal disturbances for desynchronisation are computed in the ‘swing equation’, a network of second-order Kuramoto oscillators which acts as a simple model for power-grid dynamics, in small four-node power-grids and a complex model UK power-grid. The amplitudes of minimal disturbances vary non-monotonically with the number of connections in the grid, depending on the details of the dynamical evolution of the perturbation across the grid. A comparison between the amplitude of minimal disturbances and basin stability for a range of nodal powers and dissipation rates shows that these nonlinear stability measures evolve independently, emphasising the need for both measures to be used in the design of nonlinear systems. The desynchronised dynamics of large power-grids are investigated in detail; an asymptotic expansion is developed to explain the grid-wide dynamics of single-node ‘dead-end’ desynchronisation events.

Keywords: nonlinear stability, optimisation, synchronisation, power-grids, networks

1 Introduction

Maintaining synchrony in a power-grid network is essential for its proper functioning; desynchronisation of one oscillator from the mean grid frequency can cause cascading desynchronisation across the grid and ultimately a blackout. Various models

have been used to study the large-scale dynamics of power-grids [1–7], and a particular focus has been placed on a simplified model known as the ‘swing equation’ (a second-order Kuramoto-like model) (e.g. [8]) which can be derived from considering near-synchronous power transmission between connected nodes of a idealised network [9]. In retaining some key aspects of power-grid

dynamics (interacting oscillators at each node with their own power output/usage and damping rate, linked via their phase differences), whilst removing details of the transmission lines, power generation, and higher order effects, the swing equation model has proven a fertile testing ground for analysing methodologies for deducing stability properties of these types of oscillators. Indeed, many approaches have been taken to quantify the robustness of its synchronised states, including via linear stability [10], the emergence of synchronisation [11], its response to noise [12], and transient dynamics [13, 14] and its associated ‘survivability’ [5, 15, 16].

Like many nonlinear systems, the swing equation admits multiple synchronised states [17] in addition to locally attracting desynchronised states and long-lived transient trajectories between them [18]. Analysis of this multi-stability requires a dynamical systems approach in which global features of the state space and its distinct basins of attraction are computed. The ‘basin stability’ [19] of synchronised states to single- and multi-node disturbances have been examined [20–23]; the trajectories of a very large number of initial conditions are computed, and this is used to estimate the volume of the basin of attraction of the synchronised states based upon the fraction of initial conditions which synchronise. Such analyses may also be constrained to certain sets of realistic disturbances [5, 24], or used to estimate the shape of the basin of attraction [25].

Basin stability studies provide an overall measure of the nonlinear stability of the synchronised state; the ratio of the estimated volume of the basin of attraction of synchronised states to the volume of all realistic perturbations provides an estimate of the synchronisation probability of the system, that is, the probability that a random (realistic) perturbation synchronises. For reasonably small networks, this volume estimation can be computed fairly rapidly. However as the network size becomes large, the number of random initial conditions needed to provide a reasonable estimate of the basin volume makes the method challenging to implement. For example, using the ‘brute force’ (non probabilistic) method of [22], if M different initial conditions are to be tested at each of N nodes, then the total number of

computations needed for multi-node basin stability is $(M + 1)^N - 1 = O(M^N)$, an exponential scaling in the number of nodes. Similar basin volume computations have been conducted in a fluid dynamics setting, where the fraction of initial conditions which laminarise compared to those which transition to turbulence are enumerated [25]. The extremely large number of initial conditions needed to provide a robust estimate of the laminarisation probability makes this a significant computational challenge, though it has recently been demonstrated that a substantial reduction can be made in the number of initial conditions needed by using a Bayesian method [26].

What these basin volume studies overlook, however, are details of the overall shape of the basin of attraction of the synchronised state. In a simple bi-stable system with two attracting states, the basins of attraction of the two states are separated by the stable manifold, or ‘edge’ manifold [27], of an unstable saddle point. The shape of this edge manifold, and particularly its closest approach to either of the attracting states, provides information about the most dangerous perturbations to the system. By computing the basin stability, fine details of the edge manifold will be missed, including relatively small intrusions of the edge manifold towards one of the attracting states. Indeed, in the fluid dynamics setting an initial condition which transitions to turbulence was found by Pershin *et al.* (2022) [26] with an amplitude two-orders-of-magnitude smaller than the amplitude at which 100% of initial conditions were predicted to re-laminarise. This initial condition, termed the ‘minimal seed for turbulence’ [28] is computed using a variational procedure to identify the initial condition which grows the most over large times. In Pershin *et al.* (2022) [26], it appears to be located at the end of a narrow intrusion of the turbulent attractor’s basin of attraction towards the laminar state and as such is missed by the coarse-graining involved when selecting a (large number of) random initial conditions. Although the probability of a random finite-amplitude perturbation taking the form of the minimal seed is small, it nevertheless represents a possible route away from the laminar state representing a ‘worst-case’ perturbation, and may be interpreted in terms of instantons in

large-deviation theory [29]. Minimal seeds have been computed in various shear flows [26, 30–33] in addition to problems in thermoacoustics [34], magnetohydrodynamics [35], and mixing [36, 37]. Whilst they are challenging to compute, the computational load is comparable to that of basin stability when the systems are turbulent (high dimensional) [26].

Halekotte & Feudel (2020) [38] translated the minimal seed methodology in fluid dynamics to transition in multi-stable networked dynamical systems (systems for which more than one attracting state exists at the same parameter values), including the swing equation. Instead of ‘minimal seed’, Halekotte & Feudel (2020) [38] called such initial conditions the ‘minimal shock’ which leads to desynchronisation; here we will use ‘minimal disturbance for desynchronisation’ as a more descriptively explicit title than either ‘seed’ or ‘shock’. Halekotte & Feudel (2020) [38] computed minimal disturbances in a model UK grid similar to that of Mitra *et al.* (2017) [22], and found a number of different small-amplitude perturbations to the synchronised state leading to localised desynchronisation events. These events typically occurred in ‘dead-end’ regions of the power-grid, which take a relatively simple form for single-node dead-ends [39] or evolve chaotically in small-group dead-ends [18].

The distribution of power generators and consumers in the model UK grid used by Halekotte & Feudel (2020) [38] was different to the multi-node basin stability computations of Mitra *et al.* (2017) [22], as was a key model parameter, making a direct comparison between the two stability measures challenging. Furthermore, the complexity of the model UK power-grid (with 120 nodes) precludes general conclusions about the behaviour of minimal disturbances in the swing equation being made. Halekotte & Feudel (2020) [38] also use a convergence measure for their optimisation algorithm which is known (in the turbulence transition context) to terminate the algorithm early and hence miss more optimal disturbances [31].

In this work, minimal disturbances for desynchronisation in the swing equation will be computed in the simple four-node grids of Ji & Kurths (2014) [20] and the large UK model of Mitra *et al.* (2017) [22], using the same parameters and node distributions as those basin stability computations, so that a direct comparison between

the two stability measures can be made. All the grids examined are multi-stable for the parameters under consideration. A detailed description of the optimisation procedure for finding such minimal disturbances will be given, along with a robust convergence criteria based upon gradients of the objective functional. Furthermore, an analytical asymptotic solution will be developed to fully explain the dynamics of single-node dead-end desynchronisation in the swing equation for large power-grids, rigorously demonstrating the validity of the estimates of Menck *et al.* (2014) [39] for such dead-ends.

A primary purpose of this work is to demonstrate that *both* basin stability and minimal disturbances are essential measures of the nonlinear robustness of synchronised states. They provide complimentary information about the state space; basin stability represents global information about the ‘likelihood’ of synchronising the system, but misses entirely the worst-case scenarios for desynchronisation that minimal disturbances identify. On the other hand, minimal disturbances provide very limited further information about the state space, owing to their extreme localisation. Together, however, they provide a good overview of the system’s nonlinear behaviour. Basin stability demonstrates how likely a *random* initial condition is to synchronise, whereas minimal disturbances highlight dangerous perturbations to a system that is *already* synchronised.

In section 2, the swing equation is introduced, along with the variational method for identifying minimal disturbances to cause desynchronisation and a detailed algorithm for its solution. Section 3 demonstrates the convergence of the method for a particular power-grid network, and section 4 shows how the minimal disturbances vary with network parameters and topology for the four-node grids of Ji & Kurths (2014) [20], comparing the minimal disturbances to their basin stability results. Section 5 considers the model UK grid of Mitra *et al.* (2017) [22], including a detailed analysis of its dead-end desynchronisation events. Conclusions are given in section 6.

2 The power-grid model

The swing equation models a power-grid as N coupled oscillators in a network. The i th oscillator has

phase $\phi_i(t) = \Omega t + \theta_i(t)$, where Ω is the synchronised grid frequency. The angular velocity of each oscillator is given by $\dot{\phi}_i(t) = \Omega + \dot{\theta}_i(t) \equiv \Omega + \omega_i(t)$, so that a synchronised state is represented by $\omega_i(t) = 0$ for $i = 1, \dots, N$. The oscillators are located on a graphed network (see Fig. 1) where each node (oscillator) on the graph is coupled to a number of other nodes along edges of the graph. These connections are represented by the graph's adjacency matrix A , such that $A_{ij} = 1$ if node i is connected to node j and $A_{ij} = 0$ otherwise. Nodes are not connected to themselves, so that $A_{ii} = 0$ for all i .

Power transmission between nodes is represented through a coupling strength K . For simplicity, and for later comparison with [20, 22, 38], K is taken to be the same constant for each connection in the graph, meaning that all connections in the grid transmit power equally efficiently (this simplification implies that all transmission lines are of the same length, and have identical properties, which is unrealistic for actual power grid dynamics, but is nonetheless a useful simplification for the analysis and comparison). Under the assumption of small deviation from the mean grid frequency, $|\omega_i| \ll \Omega$, and using an appropriate simplified description of power dynamics at each node, a Kuramoto-like model for the power-grid dynamics can be derived (see [9] for the details of the derivation and its simplification of nodal power dynamics), known as the swing equation:

$$\dot{\omega}_i = -\alpha\omega_i + P_i + K \sum_{j=1}^N A_{ij} \sin(\theta_j - \theta_i), \quad (1)$$

$$\dot{\theta}_i = \omega_i, \quad \text{for } i = 1, \dots, N. \quad (2)$$

In particular, power transmission between nodes depends on the phase difference between the nodes.

Each oscillator experiences a damping α which acts to synchronise the system by returning $\omega_i(t)$ to 0. Each oscillator represents either a generator or a consumer of power, and carries a (dimensionless) power P_i which is positive for generators and negative for consumers. The grid is taken to be in power balance, so that $\sum_{i=1}^N P_i = 0$. For simplicity, and again for later comparison with [20, 22, 38], we will require that exactly half of the nodes are consumers and exactly half of the nodes are generators (requiring that N is even),

and that all consumers have $P_i = -P$ and all generators have $P_i = +P$, where P is some constant (dimensionless) power.

The use of a single, constant parameter P to control the power generation and consumption (in addition single parameters K and α to control the coupling strength and nodal damping, respectively) represent a significant departure from realistic power-grid dynamics, beyond that already inherent in the derivation of the swing equation. However, it does allow for a straightforwardly designed parametric study of the swing equation's dynamics, and removes ambiguity about how to treat grids with different network topology (adjacency matrix A). Most presciently, it allows for a direct comparison between basin stability and minimal disturbances to be made, via the results of [20, 22, 38].

Synchronous states are found by setting $\omega_i = 0$. The phase differences between the nodes in the grid representing synchronised states therefore satisfy

$$K \sum_{j=1}^N A_{ij} \sin(\theta_j^s - \theta_i^s) = -P_i, \quad \text{for } i = 1, \dots, N. \quad (3)$$

These steady states may be linearly stable or unstable, and may coexist with stable asynchronous states depending upon the parameters P , α and K and the network topology and size (see e.g. [10, 20]). Equations (3) in fact define a number of families of synchronised states, since the equations are invariant under the transformation $\theta_i \mapsto \theta_i + \beta$ for any β , and inversion of A_{ij} (which can be singular) and sine introduce further distinct families of solutions even within the periodic restriction $-\pi \leq \theta_i < \pi$. Appendix A provides a description of the solution families for the graph in Fig. 1(a). To eliminate the translation symmetry in the synchronous phases, we shall instead refer only to the synchronous phase differences $\Delta\theta_{ij}^s = \theta_j^s - \theta_i^s$, where we note that $\Delta\theta_{ij}^s$ is invariant under a translation of the phases by β . Again for simplicity, for the remainder of this work the principle solutions to (3) using the arcsine function will be taken. Inversion of A_{ij} will take the Moore–Penrose pseudoinverse.

For parameters which lead to bistability between synchronous and asynchronous states, we

search for the minimal disturbance to the synchronous state which is able to transition to an asynchronous one. For this purpose, a measure of the size of disturbance is required. We define two energy-like quantities $k(t)$ and $p(t)$ which respectively measure the average ‘potential’ and ‘kinetic’ energy of the disturbance and are given by

$$\begin{aligned} p(t) &\equiv \frac{K}{4N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} (\theta_j(t) - \theta_i(t) - \Delta\theta_{ij}^s)^2, \\ k(t) &\equiv \frac{1}{2N} \sum_{i=1}^N \omega_i(t)^2. \end{aligned} \quad (4)$$

They respectively measure the average phase difference between connected nodes and the average angular velocity of the nodes. An extra factor of $1/2$ in the definition of $p(t)$ accounts for the double-counting over both i and j of each connection in the grid given that $A_{ij} = A_{ji}$ and $\Delta\theta_{ij}^s = -\Delta\theta_{ji}^s$. It is clear why $k(t)$ can be interpreted as a ‘kinetic’ energy in the system. Labelling $p(t)$ as a ‘potential’ energy is in analogy to that of linear springs which seek to restore equilibrium. Additionally, an alternative measure of the size of a phase perturbation is given by Alvares & Banerjee (2023) [16], which is the work done by external torque in moving the phase at a single node to its new value, keeping the other phases fixed; for small phase perturbations at a single node, $p(t)$ is proportional to this work done. This definition of $p(t)$ is also invariant under a translation of the phases by β . The total energy of the system is defined as $E(t) = p(t) + k(t)$ and the initial energy of a disturbance is denoted $E_0 \equiv E(0)$. In the following figures, phases differences are taken in the range $-\pi \leq \theta_j - \theta_i < \pi$ so that the quantity $p(t)$ (and by extension $E(t)$) does not grow without bound when $\omega_i \neq 0$.

The minimal disturbance to cause desynchronisation is the set of initial conditions $I_0 = \{\theta_i(0), \omega_i(0)\}_{i=1}^N$ of smallest initial energy E_0 which leads to desynchronisation. Any synchronised state has $k = 0$, and so initial conditions are sought for which the kinetic energy k grows as $t \rightarrow \infty$, and the initial condition which does this with smallest initial energy E_0 is the minimal disturbance to cause desynchronisation. It should be noted that Halekotte & Feudel (2020) [38] only consider the kinetic energy, and only allow initial

conditions with frequency perturbations, setting $\theta_i(0) = 0$ for $i = 1, \dots, N$.

To solve for the minimal initial energy E_0 and associated initial conditions I_0 , initial conditions which maximise the kinetic energy k as $t \rightarrow \infty$ are found for *fixed* initial energy E_0 ,

$$\max_{I_0} \lim_{t \rightarrow \infty} k(t) \quad \text{s.t. } E(0) = E_0, \quad (5)$$

after which a further optimisation to minimise E_0 is performed until a critical energy E_c is reached for which all initial conditions with $E_0 < E_c$ re-synchronise. Due to the nature of the optimisation algorithm used below, initial conditions which grow the most at some value of E_0 may be used as initial guesses for initial conditions which grow the most at some smaller value of E_0 and this process is repeated until no further reductions in E_0 are possible whilst still identifying an initial condition which desynchronises. Instead of taking the limit $t \rightarrow \infty$, to solve (5) we find the initial conditions I_0 which maximise $k(T)$, where $T > 0$ is some large time which is able to distinguish between re-synchronising and desynchronising trajectories.

The optimisation is performed by constructing a constrained, extended Lagrangian:

$$\begin{aligned} \mathcal{L} = k(T) &- \frac{1}{N} \sum_{i=1}^N \int_0^T \left[\varphi_i (\dot{\theta}_i - \omega_i) + \right. \\ &\left. \eta_i \left(\dot{\omega}_i + \alpha \omega_i - P_i - K \sum_{j=1}^N A_{ij} \sin(\theta_j - \theta_i) \right) \right] dt \\ &- \frac{1}{N} \sum_{i=1}^N \varphi_i^0 (\theta_i(0) - \theta_i^0) - \frac{1}{N} \sum_{i=1}^N \eta_i^0 (\omega_i(0) - \omega_i^0) \\ &- c \left(\frac{1}{2N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} |\theta_j^0 - \theta_i^0 - \Delta\theta_{ij}^s|^2 \right. \\ &\quad \left. + \frac{1}{N} \sum_{i=1}^N |\omega_i^0|^2 - 2E_0 \right), \end{aligned} \quad (6)$$

where $\varphi_i(t)$, $\eta_i(t)$, φ_i^0 and η_i^0 for $i = 1, \dots, N$ and c are the Lagrange multipliers which enforce, respectively, the kinetic equations (2), the dynamic equations (1), initial conditions for the phases, initial conditions for the angular velocities, and the total energy of the initial conditions.

Taking variations with respect to θ_i and η_i and setting them to zero yields the ‘adjoint’ equations

for the Lagrange multipliers $\varphi_i(t)$ and $\eta_i(t)$:

$$\dot{\eta}_i = \alpha\eta_i - \varphi_i, \quad (7)$$

$$\dot{\varphi}_i = K \sum_{j=1}^N (\eta_i A_{ij} - \eta_j A_{ji}) \cos(\theta_j - \theta_i), \quad (8)$$

$$\varphi_i(T) = 0, \quad \eta_i(T) = \omega_i(T), \quad (9)$$

$$\varphi_i(0) = \varphi_i^0, \quad \eta_i(0) = \eta_i^0, \quad (10)$$

and taking variations with respect to the initial conditions leads to the gradients

$$\frac{\delta \mathcal{L}}{\delta \theta_i^0} = \frac{1}{N} \left[\varphi_i^0 + 2c \sum_{j=1}^N A_{ij} (\theta_j^0 - \theta_i^0 - \Delta \theta_{ij}^s) \right], \quad (11)$$

$$\frac{\delta \mathcal{L}}{\delta \omega_i^0} = \frac{1}{N} (\eta_i^0 - 2c\omega_i^0), \quad (12)$$

which may be used iteratively to find initial conditions I_0 which grow the most over a time T .

The Lagrange multiplier c is determined by the specific method used to search for better initial conditions. Steepest ascent reads

$$\begin{aligned} \theta_i^{0,n} = & \theta_i^{0,n-1} + \frac{\epsilon}{N} \left\{ \varphi_i^0 \right. \\ & \left. + 2c \sum_{j=1}^N A_{ij} \left[\theta_j^{0,n-1} - \theta_i^{0,n-1} - \Delta \theta_{ij}^s \right] \right\}, \end{aligned} \quad (13)$$

$$\omega_i^{0,n} = \omega_i^{0,n-1} + \frac{\epsilon}{N} \left[\eta_i^0 - 2c\omega_i^{0,n-1} \right], \quad (14)$$

where ϵ is a step-size between iterations, and n refers to the n th iterate in the optimisation procedure starting from some guess $\{\theta_i^{0,0}, \omega_i^{0,0}\}_{i=1}^N$ (which will typically arise from an optimal solution at a larger value of E_0). From these update relations, c is determined by ensuring that the new initial conditions have energy E_0 (which, in the case of steepest ascent, results in a quadratic equation for c which is given in Appendix B).

To summarise the algorithm for computing minimal disturbances which cause desynchronisation, Algorithm 1 shows a pseudocode outlining the procedure. The algorithm uses m to count the energy levels examined and n to count iterations at each energy level. To begin the algorithm at

$m = 0$, a random desynchronising initial condition with relatively large energy is identified with initial energy $E_0^m \gg E_c$. This initial condition is then uniformly re-scaled to a lower initial energy $E_0^{m+1} = rE_0^m$ where r is some re-scaling factor that is adjusted as the algorithm proceeds. This new initial condition is then integrated using the swing equation (2-1) until the target time $t = T$ is reached. The adjoint equations (7-8) are initialised at $t = T$ using (9) before being integrated backwards in time until $t = 0$. The Lagrange multipliers φ_i^0 and η_i^0 are set using (10) and then c is computed using (B11).

If this initial condition desynchronised the system, then m is incremented and the initial condition is immediately re-scaled in energy once more via $E_0^{m+1} = rE_0^m$. If the initial condition did not desynchronise the system, then n is incremented and a new initial condition at the same energy level is computed using (13-14). This process is repeated until a desynchronising initial condition is found at this energy level. The energy levels are continued to be rescaled until a synchronising initial condition has gradient residual \mathcal{R} , defined by

$$\mathcal{R} \equiv \sum_{i=1}^N \frac{1}{(\varphi_i^0)^2} \left(\frac{\delta \mathcal{L}}{\delta \theta_i} \right)^2 + \sum_{i=1}^N \frac{1}{(\eta_i^0)^2} \left(\frac{\delta \mathcal{L}}{\delta \omega_i^0} \right)^2, \quad (15)$$

smaller than some preset tolerance. The residual \mathcal{R} measures the magnitude of the gradient of \mathcal{L} , and division by φ_i^0 and η_i^0 is performed given that the adjoint equations (7-8) are linear in φ_i and η_i meaning that the amplitude of these gradients is somewhat arbitrary (albeit set by $\eta_i(T) = \omega_i(T)$). In contrast, Halekotte & Feudel (2020) [38] identify convergence when improvements in the objective functional (in this case $k(T)$) between consecutive iterations become sufficiently small. Though this weaker convergence criterion has some precedence in the fluid mechanics turbulence transition problem (see e.g. [40]), it is known that minimal changes in the objective functional may occur for many hundreds of iterations while a gradient residual equivalent to \mathcal{R} remains large, before sudden increases in the objective functional are found, accompanied by an eventual decay in \mathcal{R} [31].

When a synchronising initial condition with a small residual is found, if the difference between

Algorithm 1 Finding minimal disturbances

```
1:  $m \leftarrow 0$ 
2: Find desynchronising IC  $\{\theta_i^0, \omega_i^0\}_{i=1}^N$ 
3:  $m \leftarrow m + 1$ 
4:  $E_0^m = rE_0^{m-1}$   $\triangleright$  Reduce  $E_0$ 
5:  $\theta_i^0 \leftarrow \theta_i^s + \sqrt{\frac{E_0^m}{E_0^{m-1}}}(\theta_i^0 - \theta_i^s)$ 
6:  $\omega_i^0 \leftarrow \sqrt{\frac{E_0^m}{E_0^{m-1}}}\omega_i^0$ 
7:  $n \leftarrow 0$ 
8:  $\theta_i(0) \leftarrow \theta_i^0$  and  $\omega_i(0) \leftarrow \omega_i^0$ 
9: Solve swing equation (1–2) for  $0 \leq t \leq T$ 
10:  $k_n(T) \leftarrow \frac{1}{2N} \sum_{i=1}^N \omega_i(T)^2$ 
11: if  $k_n(T) < k_{n-1}(T)$  then  $\triangleright$  No improvement
12:   Reduce step-size  $\epsilon$ 
13:   Recompute  $c$  via (B11)
14:   Recompute gradients  $\frac{\delta \mathcal{L}}{\delta \theta_i^0}, \frac{\delta \mathcal{L}}{\delta \omega_i^0}$  via (11–12)
15:   Create new ICs via (13–14)
16:   go to 8  $\triangleright$  Restart from  $t = 0$ 
17: else if  $k_n(T)$  shows desynchronisation then
18:   go to 3  $\triangleright$  Reduce  $E_0$ 
19: end if
20:  $\varphi_i(T) \leftarrow 0$  and  $\eta_i(T) \leftarrow \omega_i(T)$ 
21: Solve adjoint equation (7–8) for  $T \geq t \geq 0$ 
22:  $\varphi_i^0 \leftarrow \varphi_i(0)$  and  $\eta_i^0 \leftarrow \eta_i(0)$ 
23: Compute  $c$  via (B11)
24: Compute gradients  $\frac{\delta \mathcal{L}}{\delta \theta_i^0}, \frac{\delta \mathcal{L}}{\delta \omega_i^0}$  via (11–12)
25: if Consecutive gradients aligned then
26:   Increase step-size  $\epsilon$ 
27:   Recompute  $c$  via (B11)
28:   Recompute gradients via (11–12)
29: else if Gradients misaligned then
30:   Decrease step-size  $\epsilon$ 
31:   Recompute  $c$  via (B11)
32:   Recompute gradients via (11–12)
33: end if
34: Compute residual  $\mathcal{R}$  via (15)
35: if  $\mathcal{R}$  does not indicate convergence then
36:    $n \leftarrow n + 1$ 
37:   Create new ICs via (13–14)
38:   go to 8  $\triangleright$  Continue iterating at this  $E_0$ 
39: end if  $\triangleright$  Converged residual
40: if  $E_0^{m-1} - E_0^m > \epsilon_{E_c}$  then  $\triangleright$  Large energy gap
41:    $E_0^m \leftarrow \frac{E_0^m + E_0^{m-1}}{2}$   $\triangleright$  Increase  $E_0$ 
42:   Decrease  $r$ 
43:   go to 5  $\triangleright$  New ICs with larger  $E_0$ 
44: end if
45: Converged:  $E_0^m < E_c < E_0^{m-1}$ 
```

its initial energy and the previous initial energy level, $E_0^{m-1} - E_0^m = (1 - r)E_0^{m-1}$, is below some threshold ϵ_{E_c} (taken here to be 10^{-4}), then optimal initial conditions either side of the edge have been found and (a local result for) the critical minimal energy E_c for desynchronisation satisfies $E_0^m < E_c < E_0^{m-1}$. However, if the current difference between consecutive energy levels is too large, then the algorithm may still be a long way from converging to E_c . In this case, the energy level E_0^m is replaced by $(E_0^m + E_0^{m-1})/2$ and the re-scaling factor r is reduced accordingly for all further iterations of the algorithm. These re-scaling operations continue until the edge is sandwiched by converging initial conditions separated in energy by no more than ϵ_{E_c} .

During iterations within each energy level, two methods are used to adjust the step-size ϵ used in the steepest ascent updates (13–14). Firstly, if at time $t = T$ the current iteration has a kinetic energy $k(T)$ which is smaller than that of the previous iteration, then the gradient ascent update (13–14) has overshot a critical point and the step-size ϵ is too large. A reduction in ϵ is therefore made, which then requires the Lagrange multiplier c and the associated gradients from the previous iteration to be recomputed, given that they depend on ϵ , before a new initial condition is produced. Secondly, a normalised (between -1 and 1) dot-product between the $2N$ -dimensional gradient vectors at the current and the previous iteration is taken to see the alignment between the two consecutive gradient directions. If the dot-product is greater than 0.95, then the two consecutive gradients are closely aligned and it is possible to use a larger step-size ϵ as both gradients are indicating that the critical point is in the same direction. On the other hand, if the dot-product is less than -0.5 then the consecutive gradient vectors are misaligned and the algorithm is zigzagging towards the critical point, meaning that ϵ must be reduced. This dot-product is taken at most once per iteration to avoid over-adjusting the step-size ϵ . Any adjustments to ϵ made in one energy level are carried forward to the next.

On occasion, the algorithm fails to find desynchronising initial conditions or improvements to $k(T)$ whilst the residual \mathcal{R} remains large. In such instances, some manual fine-tuning may be required, with options including reducing the energy re-scaling factor r anyway and restarting

the algorithm from there, or increasing the target time T . The latter approach is particularly important as E_c is approached since trajectories spend larger amounts of time on the edge manifold before either synchronising or desynchronising as the distance to the edge decreases.

A similar methodology has been used in fluid dynamics to identify minimal seeds for turbulence [26, 30, 32, 33] in addition to approaching E_c from below [31, 32] (see Kerswell (2018) [28] for a review).

3 Identification of minimal disturbances

The basin stability computations of Ji & Kurths (2014) [20] considered four simple graphs, labelled $G1$ through $G4$ and shown in Fig. 1. Two nodes (1 and 4) were taken as consumers with $P_i = -P$ and two nodes (2 and 3) as generators with $P_i = P$. The adjacency matrices for these graphs are given by

$$\mathbf{A}_1 = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}, \quad \mathbf{A}_2 = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix},$$

$$\mathbf{A}_3 = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}, \quad \mathbf{A}_4 = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}. \quad (16)$$

All four graphs exhibit multi-stability in the swing equations (2-1), when $K = 1$ and α and P are varied. In particular, when $\alpha = 0.5$ and $P = 0.8$, a random initial condition has around 30% to 40% chance of desynchronising each of the networks (see [20] for exact results). These parameter values will be used throughout this section when finding minimal disturbances for each graph; other parameter values will be examined in section 4. It can be seen that the dynamics of graphs $G2$ and $G3$ are equivalent under the transformation $\theta_i \mapsto -\theta_i$ and $P_i \mapsto -P_i$, and so only graphs $G1$, $G2$ and $G4$ will be considered further here.

Fig. 2(a) shows the time-evolution of the total energy $E(t)$ for optimal disturbances at a selection of initial energies E_0^m , starting the algorithm at $E_0^0 = 2$. For clarity of the figure, not all values of m used during the computation are shown.

As the initial energy E_0 is decreased, the optimal solutions spend longer and longer near a plateau of intermediate energy $E \approx 1.44$ which is associated with a saddle point in state space separating the attracting synchronised and desynchronised dynamics. The smallest initial perturbation which desynchronises is plotted in thick blue, and has initial energy $E_0 = 0.6656$. The thin red line shows the trajectory of the optimal disturbance with $E_0 = 0.6655$, which ultimately energetically decays and returns to the synchronised state; no initial conditions which cause desynchronisation are found by the optimisation procedure at this lower energy. We can therefore conclude that the minimum disturbance to cause desynchronisation (at least in the sense of a local minimum, see later discussion) has critical energy E_c in the range $0.6655 < E_c < 0.6656$.

The dynamics of this minimal disturbance are shown in figures 2(b,c) which plot the evolution of its angular velocities and phases (plotted restricted to $[-\pi, \pi)$) respectively. The trajectory admits a ready interpretation. Some kinetic energy is initially located in two of the nodes (node 1, a consumer and node 3, a generator) whilst the other two nodes begin with negligible kinetic energy. Meanwhile, nodes 1 and 3 also begin with a moderate phase perturbation, as do nodes 2 and 4 albeit at a smaller amplitude. The initial kinetic energy in nodes 1 and 3 drives the phase of these oscillators whilst leaving the other two nodes largely undisturbed from their initial configuration. As the phase of node 1 and 3 are driven by their angular velocity, they lose kinetic energy until $t \approx 3$ at which point all nodes have negligible kinetic energy and their phases remain roughly constant for $3 \lesssim t \lesssim 8$. The set of phases sustained during this quasi-steady period is in fact a solution of the synchronised equations (3), also plotted in Fig. 2(c). However, this particular synchronised solution is (linearly) unstable and so the trajectory eventually exponentially diverges from this saddle-point state along its unstable manifold; the minimal disturbance itself lies on its stable manifold. Once the trajectory leaves the vicinity of this saddle point, it breaks down into desynchronised dynamics.

To demonstrate that the solution plotted in figures 2(b,c) is a converged solution of the optimisation problem, Fig. 3 plots the residual \mathcal{R} defined




Fig. 1 The four simple graphs, $G1$, $G2$, $G3$ and $G4$. Generators are shown with circles and consumers are shown with squares.

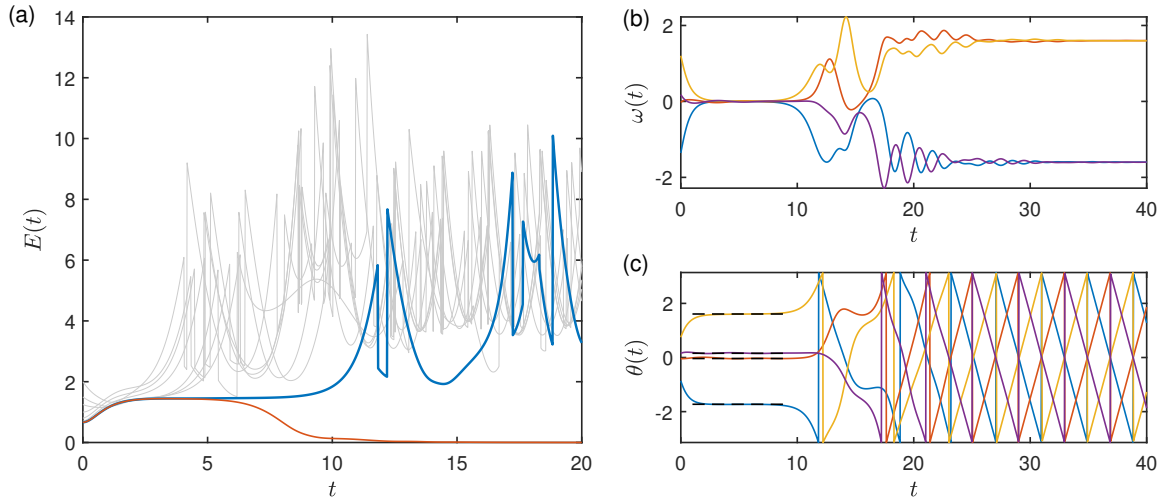


Fig. 2 Dynamics of optimal disturbances in graph $G1$. (a) Time-evolution of the total energy $E(t)$ for optimal disturbances with initial energies $E_0 = 2, 1.5, 1.25, 1, 0.8, 0.7,$ and 0.675 (grey) and 0.6656 (blue, thick) and 0.6655 (red). (b,c) Time-evolution of the angular velocities and phases respectively of each node for the minimal disturbance with initial energy $E_0 = 0.6656$, with node 1 (blue), node 2 (red), node 3 (yellow) and node 4 (purple). In (c), a steady state solution for the phases (detailed in Appendix A) is shown (black, dashed).




Fig. 3 Residual \mathcal{R} defined by (15) (left-hand axis, blue) and initial energy E_0 (right-hand axis, red) against iteration number n .

in (15). Convergence of the solution requires that $\mathcal{R} \rightarrow 0$. Fig. 3 shows that whilst \mathcal{R} remains large for initial energies E_0 substantially larger than the minimum disturbance energy E_c , convergence of \mathcal{R} occurs as the initial energy E_0 is reduced. This is because, far from E_c there are many initial conditions which desynchronise, and it is impossible to define an ‘optimal’ such initial condition; nearby to E_c there are far fewer initial conditions which do so, and a clear optimal emerges. The convergence of \mathcal{R} is somewhat chaotic, with small values interspersed with large values in the latter stages of the algorithm; this feature is in common with minimal seed computations in fluid turbulence [32] and other optimisation problems exhibiting ‘chaotic gradient descent’ (see e.g. [41]).

The nature of the optimisation algorithm means that any solutions found are inherently local optima which do not rule out the possibility of other, less energetic disturbances which desynchronise the system. To investigate the potential of a better local optimal solution, a number of alternative strategies were employed to initialise and run the optimisation algorithm. The results in Fig. 2 were initialised at $E_0 = 2$ with a random initial condition (each θ_i and ω_i sampled from a normal distribution with mean 0 and variance 1, then scaled uniformly to have energy $E_0 = 2$). Another random initialisation chosen from the same distribution was gradually reduced in energy by individually reducing $k_0 \equiv k(0)$ and $p_0 \equiv p(0)$ instead of E_0 , which led to the same optimal solution. Furthermore, two alternative initialisations were trialled (followed by reducing E_0) which (1) disturbed only node 1, and (2) disturbed pairs of

nodes 1 and 4, and 2 and 3 identically owing to the symmetry $1 \leftrightarrow 4$ and $2 \leftrightarrow 3$. Both strategies ultimately led to the same optimal disturbance as E_0 was lowered. Although these further examples do not demonstrate that the solution presented is in fact a global optimum, they provide evidence that the solution may be such, and at least demonstrate that the solution found by the original algorithm represents a dynamically significant trajectory.

For the remainder of this paper, only solutions found using the original algorithm will be presented, with the caveat that they may only be local optimal solutions. One could imagine searching for other local minima by extending the Lagrangian (6) with a further term which penalises initial conditions close to the already identified minima, though it is unclear how the new optimisation problem would behave, and such an investigation is beyond the scope of this work. However, as section 4 will demonstrate, finding *any* minimal disturbance at all represents an improvement over basin stability computations (when the purpose is identifying worst-case scenarios), which miss such perturbations entirely. Indeed, the two stability measures are designed to measure fundamentally different aspects of the state space and its trajectories, and as such each supplies complementary information inaccessible to its counterpart, however robust that information may be.

The dynamics of minimum disturbances to cause desynchronisation found in graphs G_2 and G_4 are shown in figures 4(a,c) and 4(b,d) respectively (with the same parameters $K = 1$, $\alpha = 0.5$ and $P = 0.8$). The critical energy for G_2 satisfies $1.40054 < E_c < 1.40055$ and for G_4

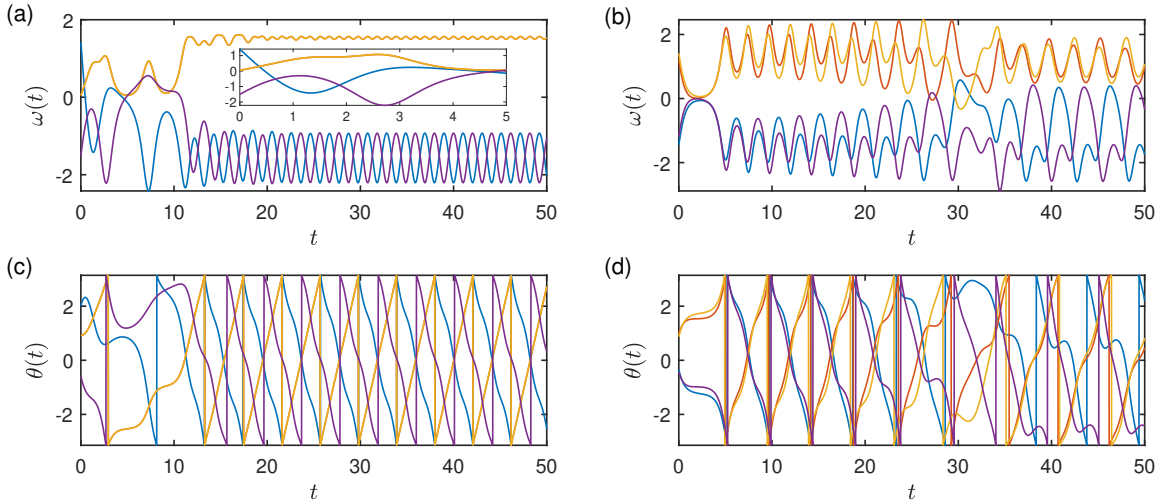


Fig. 4 Dynamics of the minimal disturbances in graphs $G2$ (a,c) and $G4$ (b,d). Time-evolution of the angular velocities (a,b) and phases (c,d). Line styles as in figures 2(b,c). For graph $G2$ in (a,c), the dynamics of nodes 2 and 3 are identical. Insert in (a) shows early-time dynamics.

satisfies $1.26725 < E_c < 1.26750$. Both critical energies are substantially larger than that found for $G1$, indicating that the more connected grids are substantially more (nonlinearly) stable to perturbations than $G1$, despite having similar sized basins of attractions at these parameters according to Ji & Kurths (2014) [20]. This underlines one of the central motivations for this work: local properties of the basin of attraction, in particular the closest approach of the edge manifold to the synchronised state, significantly affect the robustness of the synchronised state and may be missed by a global basin-size approach.

The trajectory of the minimal disturbance of graph $G2$ plotted in figures 4(a,c) shows that the dynamics of nodes 2 and 3 (which are newly connected in comparison to $G1$) are identical for the minimal disturbance despite the initial guess at $E_0 = 2$ not having this restriction. Additionally, no steady state saddle-point is visited by the trajectory. Unlike $G1$, for $G2$ the initial kinetic energy is stored in each consumer (nodes 1 and 4) which are the least connected nodes in the graph. This (relative) isolation in the grid appears to allow the initial kinetic energy to drive the phases of these two nodes in a way that minimises losses to the rest of the grid. In turn, this drives the other two nodes in the grid. The efficiency of storing energy in the more isolated nodes at the outset, and then driving the rest of the grid, appears to be the main

factor behind the minimal disturbance in this particular grid. Its dynamics are thus distinct from both $G1$ and $G4$, for which all nodes are equally well connected.

The trajectory of the minimal disturbance of graph $G4$ plotted in figures 4(b,d) does visit an unstable steady state briefly around $t \approx 2$ though not for any substantial period of time. The initial kinetic energy is distributed in all four nodes, in contrast to both $G1$ and $G2$. Unlike $G2$, no pair of nodes shares an identical evolution, although the dynamics of the pair of nodes 1 and 4, and the pair of nodes 2 and 3 are similar throughout. This is in contrast to $G1$ where this grouping is evident only after the steady state is visited; for $G4$, the connections between these pairs of nodes appears to force them to be grouped from the offset, therefore requiring a larger initial energy for the minimal disturbance. Indeed, the critical energy E_c for $G4$ is almost exactly twice that of $G1$, suggesting that the initial dynamics of $G1$ which are largely restricted to only half of the grid, are mirrored across the other half of the grid in $G4$. At around $t \approx 30$, the (already desynchronised) dynamics transition to an alternative behaviour in which nodes 2 and 3 are in-phase with one another whilst nodes 1 and 4 are out-of-phase with one another. The structure of the underlying dynamical systems appear to become more complicated as the graphs become more connected; indeed, the desynchronised state in $G1$ has steady angular

velocities, but this is not true of $G2$ and $G4$ where the desynchronised angular velocities vary periodically, and $G4$ undergoes two distinct regimes of desynchronisation.

Of particular note is that, although these more connected grids are more nonlinearly stable than $G1$, this stability does not increase monotonically with the number of links; in $G4$ all possible connections between nodes are realised, but its critical energy is lower than that for $G2$. It is unclear why this ordering arises, though the minimal disturbance in $G2$ does not visit a steady state ‘stepping stone’ on route to desynchronisation, which may result in a large initial energy being required for transition. However, any further interpretation is made challenging due the significant complexity of steady states in $G2$ and $G4$ (see Appendix A).

4 Parametric dependence

The dependence of the minimal disturbance energy E_c on the parameters α and P will be investigated for graph $G1$ with $K = 1$. Ji & Kurths (2014) [20] compute the probability p_{sync} of a random initial condition synchronising for $0.1 \leq P \leq 1.9$ and $0.1 \leq \alpha \leq 2.9$, each in steps of 0.1. Comparing these probabilities with the (upper bound of) the critical energy E_c will indicate to what extent the variation of global state space structures aligns with the nonlinear robustness of a particular synchronised state. For simplicity of discussion, the upper bound for E_c will itself be referred to as E_c in the remainder of this work.

Two cases will be examined: $0.5 \leq P \leq 2$ with $\alpha = 0.5$ and $0.2 \leq \alpha \leq 0.9$ with $P = 0.8$. The synchronisation probabilities for these two cases are shown in Fig. 5(c,d), and it can be seen that in the former case, Ji & Kurths (2014) [20] estimate that $p_{sync} \approx 1$ for $P \leq 0.5$ and $p_{sync} \approx 0$ for $P \geq 1.3$, and in the latter case, $p_{sync} \approx 1$ for $\alpha \geq 1$ and $p_{sync} \approx 0$ for $\alpha \leq 0.1$.

Fig. 5(a) shows the critical energy E_c for varying P with fixed $\alpha = 0.5$. It is seen that E_c varies gradually from around 1 when $P = 0.5$ to 0 when $P = 2$. When $P > 2$ the synchronised state does not exist (see Appendix A). In contrast to the critical energy E_c , the synchronisation probability for this case plotted in Fig. 5(c) decreases sharply from $p_{sync} \approx 1$ at $P = 0.5$ to $p_{sync} \approx 0.04$ at $P = 1$. For $P > 1$, the synchronisation probability remains small and is essentially recorded to be

zero for $P \gtrsim 1.3$. The critical energy E_c places no such importance on or around the value $P = 1$.

These results clearly demonstrate how E_c and its associated initial condition, local properties of the nonlinear structure of the state space, evolve almost entirely independently from its global structure in terms of basins of attraction of synchronised states. Indeed, were the system *already* in this particular synchronised state, its robustness to a finite-amplitude perturbation is not reduced nearly as rapidly as the synchronisation probabilities would suggest as P increases. Whilst a *random* initial condition with $P = 1$ has a very small probability (4%) of synchronising, disturbances about this particular synchronised state need only be around 50% smaller energetically than at $P = 0.5$ (where the probability of synchronisation to *some* synchronised state is nearly 100%) to ensure that the system remains synchronised. Furthermore, when $P = 0.5$, there still exists a minimal disturbance for desynchronisation, despite the fact that $p_{sync} \approx 1$.

Fig. 5(b) shows the variation of E_c with α for fixed $P = 0.8$. We see that there is a gradual increase in E_c as α increases from 0.1 to 0.7, followed by a sudden increase in E_c at $\alpha = 0.75$ and erratic behaviour around this higher value of E_c for $0.75 \leq \alpha \leq 0.9$. Meanwhile in Fig. 5(d) the synchronisation probability $p_{sync} \approx 0$ when $\alpha = 0.1$ and remains small for $\alpha \lesssim 0.4$. For $\alpha > 0.4$ the synchronisation probability gradually increases so that $p_{sync} \approx 1$ when $\alpha = 1$. As with the variation in P in figures 5(a,c), there are features in the variation of p_{sync} with α that are not replicated in the variation of E_c with α and *vice versa*. Whilst there is a sudden increase in p_{sync} around $\alpha \approx 0.4$, the critical energy E_c varies smoothly through this region. Similarly, the erratic behaviour of E_c for $\alpha \geq 0.75$ is not present in the synchronisation probability.

It is unclear what structural changes occur in the system which lead to the sudden increase in E_c at $\alpha = 0.75$ in Fig. 5(b). However, it is interesting to note that the steady edge state visited by minimal disturbance trajectories for $\alpha < 0.75$ has two unstable directions, whereas the state visited for $\alpha \geq 0.75$ has three unstable directions. It is unclear why the minimal disturbance swaps which desynchronisation route it takes; the original edge state still exists for $\alpha \geq 0.75$ but is apparently




Fig. 5 Variation of the critical energy E_c against (a) P with $\alpha = 0.5$ and (b) α with $P = 0.8$ for graph $G1$. Probability of a random initial condition synchronising, p_{sync} , against (c) P with $\alpha = 0.5$ and (d) α with $P = 0.8$ for graph $G1$ taken from Ji & Kurths (2014) [20].

no longer part of an optimal path to desynchronisation. The erratic behaviour for $\alpha \geq 0.75$ may be due to difficulties in the convergence of the optimisation algorithm; as the number of unstable directions of the edge state increases, the convergence performance of the algorithm decreases [26, 42].

The computational time required to find the minimal disturbances varies substantially across the parameters considered here. When varying P with fixed α , the results with $P \geq 0.7$ took around 1 to 2 hours (on a single, relatively slow 1.60 GHz processor running alongside other low-intensity day-to-day tasks). For $P < 0.7$, the total computational time increased dramatically, with 15 hours for $P = 0.6$ and 76 hours for $P = 0.5$ (using the same hardware). When varying α with fixed P , the results with $0.1 \leq \alpha \leq 0.3$ took 1 to 4 hours (with $\alpha = 0.2$ taking the least time), and the time varied non-monotonically for all values of α , generally increasing as α increased. For example, $\alpha = 0.4$ took 16.5 hours, $\alpha = 0.7$ took 42.5 hours, and $\alpha = 0.9$ took 360 hours (15 days). Surprisingly, between these increasingly long computations, $\alpha = 0.6$ took only 4 hours and $\alpha = 0.8$ took 5 hours.

These times to convergence cannot be known *a priori* by the method. However, in general, the computational time appears to increase as the *basin stability* measure p_{sync} approaches 1. At first glance, this appears to be in contrast to the algorithm's behaviour when used to study transition to turbulence [26, 33] and other chaotic systems [42] in which convergence difficulties have

been associated with the complexity of the dividing saddle-point visited by the minimal trajectory. However, $p_{sync} \rightarrow 1$ means that the basin of attraction of the desynchronised state is shrinking, and hence it is likely colliding with the saddle point in a bifurcation, thus potentially complicating the dynamics in the vicinity of the saddle point, leading to longer times to converge in a similar manner to that understood in those other contexts.

Although the convergence time becomes excessively long for some parameter values, the time to reach an initial energy E_0 within 10% of the final value of E_c is less than 4 hours in most cases, with the few outliers (typically the longest cases overall) reaching this 10% threshold within the first third of their total computational time. The vast majority of the computational time is spent finessing initial conditions near the final answer in order to converge (reach a small residual \mathcal{R} and small energy gap ϵ_{E_c}); identification of reasonable approximations to worst-case scenarios can be achieved much more rapidly.

This puts the computation of minimal disturbances on a favourable footing alongside multi-node basin stability; the latter takes $\mathcal{O}(M^N)$ computations of the swing equation (where M is the number of initial conditions examined at each node) to estimate global properties of the basin of attraction, which scales exponentially with the number of nodes N . In doing so it is very unlikely to identify any worst-case trajectories, even approximately. On the other hand, the minimal disturbance algorithm uses steepest

ascent, the convergence rate of which is not explicitly tied to the system size (although the dynamics near the dividing saddle-point are likely to become more complex as the system size increases), but an adjusted algorithm is needed to search for further local optimal solutions. To provide rigorous (converged) results within a reasonable computational time, algorithms for both basin stability and minimal disturbances require improvement, for example using Bayesian methods for basin stability [26] or a penalisation method for (minimal) trajectory optimisation [42].

5 A UK grid

Fig. 6(a) shows a ‘United Kingdom’ power network graph with $N = 120$ nodes used by Mitra *et al.* (2017) [14, 22] to study basin stability and perturbation recovery in the swing equation. Mitra *et al.* (2017) [22] noted that (in agreement with [39]) network dead-ends ‘undermine’ power-grid stability. In particular, they find that disturbances at nodes adjacent to dead-end nodes are the most likely to lead to a desynchronised state. Halekotte & Feudel (2020) [38] also studied this United Kingdom grid using the swing equation, computing minimal disturbances for desynchronisation (which they termed ‘minimal shocks’). However, though they used the same graph topology (the adjacency matrix A), the distribution of consumers and generators used by Halekotte & Feudel (2020) [38] was different to that of Mitra *et al.* (2017) [14, 22], as was the choice of the parameter K , thus ruling out a direct comparison between the works.

Following Mitra *et al.* (2017) [22], the parameters are chosen to be $P = 1$, $\alpha = 0.1$ and $K = 8$, and the distribution of consumers and generators is that of Mitra *et al.* (2017) [22] shown in Fig. 6(a). With these parameters, the dynamics of the swing equation admit several localised desynchronisation events, wherein the perturbation is concentrated in a small number of nodes, but the entire grid has a persistent desynchronisation with $\omega_i \neq 0$ given the coupling between the nodes. One such localised desynchronisation is shown in Fig. 6(b) which was found by sampling a few randomly chosen initial conditions. It has initial energy $E_0 = 8.73$ with all nodes active across a spread of initial angular velocities $|\omega_i(0)| \lesssim 10$. As time progresses, the perturbations at $N - 2$

nodes decay, and only two nodes remain ‘active’ with a substantial ongoing perturbation. These two nodes are indicated in Fig. 6(a), which shows that one of these nodes is a dead-end and the other is the dead end’s only connecting node. The dead-end node oscillates around an average value of $\langle \omega_i \rangle = P_i/\alpha = -10$ with an oscillation frequency of $|P_i|/\alpha = 10$, as suggested by Menck *et al.* (2014) [39] and discussed below. Meanwhile, its only neighbour oscillates around an average value of $\langle \omega_i \rangle \approx 0.11$ with the same oscillation frequency.

The randomly-generated desynchronisation event of Fig. 6(b) is used as an initial guess for the optimisation procedure to find a minimal disturbance. Given that the UK-grid admits several isolated desynchronisation events, there is a significant chance that the optimisation will converge to a local solution only, focussed around one particular desynchronisation event; further random initialisations of the algorithm could be used to study other desynchronisation events. Nevertheless, the converged solution shown in Fig. 6(c) represents an entirely different desynchronised solution from the initial guess, located in a different part of the graph. This new desynchronisation once more consists of a single dead-end node which oscillates around an average value of $\langle \omega_i \rangle = P_i/\alpha = 10$ whilst its only neighbour oscillates around $\langle \omega_i \rangle \approx -0.11$ (the sign change from Fig. 6(b) is due to a swap of generator and consumer for the dead end and its only neighbour). The critical energy of this converged solution is $E_c = 2.69$. This substantial reduction in initial energy is reflected in the tighter spread of initial frequencies in the grid, with $|\omega_i(0)| \lesssim 5$.

Interestingly, a grid-wide disturbance is still required to trigger this desynchronisation. This is in contrast to the basin stability computations of Mitra *et al.* (2017) [22], who indicate small synchronisation probabilities ($p_{sync} \approx 0.2$) for disturbances at a *single* node for nodes which neighbour a dead-end. However, Mitra *et al.* (2017) [22] consider initial frequencies up to $|\omega_i(0)| = \omega_{max} = 100$, which as a single-node perturbation leads to an initial kinetic energy of $k(0)_{max} = \omega_{max}^2/2N \approx 41.7$. As a naive estimate (assuming a uniform probability distribution for synchronisation), we may expect single-node disturbances to lead to desynchronisation for energies above $p_{sync}k(0)_{max} \approx 8.33$ which is substantially larger

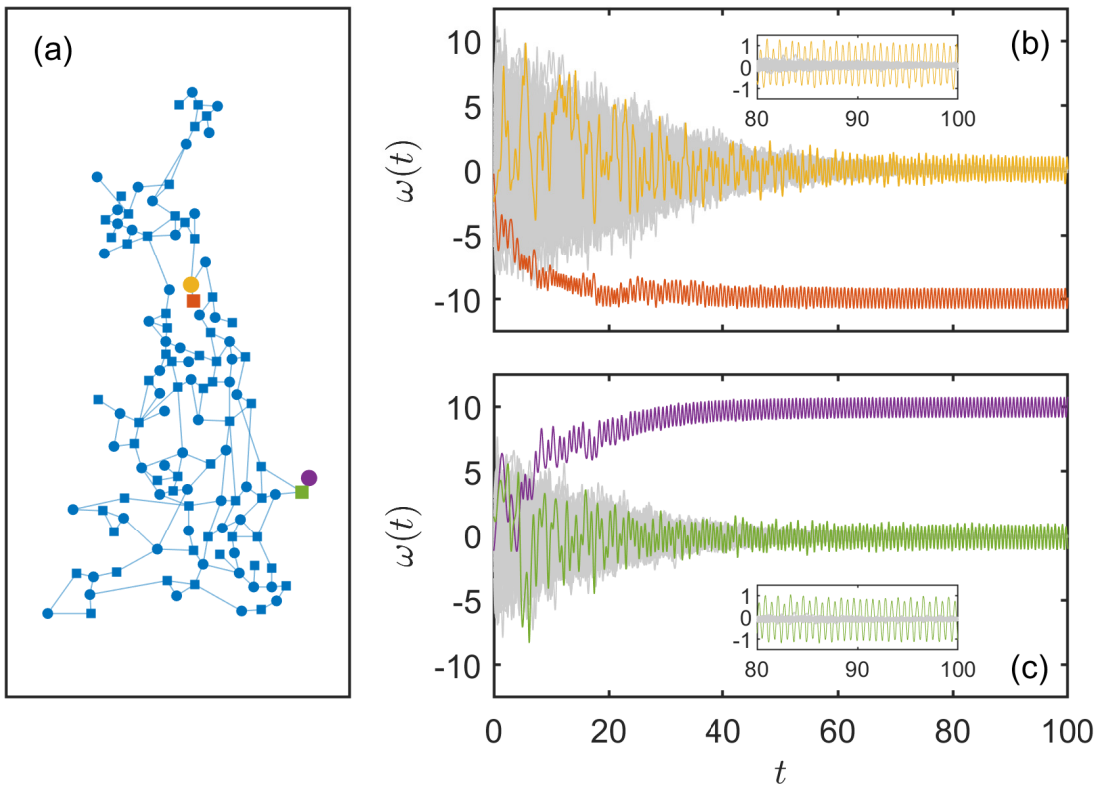


Fig. 6 (a) A United Kingdom-like grid with $N = 120$ nodes split equally into generators (circles) and consumers (squares). Desynchronising trajectories $\omega(t)$ for (b) a random initial condition with $E_0 = 8.73$ used as a starting point for minimal disturbance optimisation, and (c) the associated converged minimal disturbance with $E_c = 2.69$. In (b,c) 118 nodes are shown in light grey, and 2 nodes are shown in colour, as indicated in the grid in (a). Results shown for $P = 1$, $\alpha = 0.1$ and $K = 8$, as in Mitra *et al.* (2017) [22].

than the critical energy of the minimal disturbance, $E_c = 2.69$. It is reasonable to hypothesise therefore that the minimal disturbance requires the entire grid to be disturbed and in doing so accesses much lower initial energies. Mitra *et al.* (2017) [22] show that the average synchronisation probability for *multi-node* disturbances decreases rapidly as the number of disturbed nodes is increased, saturating at $p_{sync} \approx 0.0005$. Performing a similar estimate in which the entire grid is perturbed leads to $k(0)_{max} = (\omega_{max}^2/2N)N = 5000$ and $p_{sync}k(0)_{max} = 2.5$, which is comparable to E_c . However, the reliability of this estimate is difficult to interpret given a lack of energetic information (in particular, its synchronisation probability distribution) provided by Mitra *et al.* (2017) [22].

The dynamics of the minimal disturbance found by Halekotte & Feudel (2020) [38] are somewhat different to that computed here. Most of the initial frequencies in their minimal disturbance are zero, and those in a dead-end group of 7 nodes are all negative whilst a nearby ‘counterbalancing’ cluster within the grid are all positive. The frequencies remain nearby their initial values for a while, before they all decay to near zero. After a short period with near-zero frequencies, there is a sudden burst of kinetic energy, with all frequencies across the grid evolving chaotically at elevated values. This high activity continues for a substantial period, before the frequency of all but the dead-end nodes collapse to near zero and the remaining active nodes oscillate around either zero or $\pm P/\alpha$. The latter stages of these dynamics (chaotic frequencies across the grid, collapsing

onto well-defined dead-end desynchronisation) are similar to those found here. There are a few possible reasons for the differences in the initial dynamics, and why the initial condition and initial dynamics of Halekotte & Feudel (2020) [38] are ‘cleaner’ and more readily interpreted. Most significantly, [38] only allow initial conditions with frequency perturbations, without phase perturbations. This means that there must be an initial period in which the frequencies drive the phases away from equilibrium in order to activate the nonlinear coupling term. This appears to allow a more ‘controlled’ behaviour which drives the phases towards an unstable saddle point equilibrium (as in the four-node graph $G1$ in Fig. 2(b,c)) which then breaks down to chaotic motion. An additional factor is that the coupling constant K is set to 5 by Halekotte & Feudel (2020) [38] instead of 8 used here (and by Mitra *et al.* (2017) [22]), meaning that nonlinear interactions are weaker in the results of Halekotte & Feudel (2020) [38]. Finally, the distribution of generators and consumers was chosen randomly by Halekotte & Feudel (2020) [38], whereas the current work uses the same distribution as Mitra *et al.* (2017) [22]. This means that the final dead-end desynchronisation of Halekotte & Feudel (2020) [38] does not exist in the current work; their final dead-end consists of two nodes which are both consumers, whereas the same two nodes in this work (the two uppermost nodes in Fig. 6(a)) consist of one generator and one consumer.

5.1 Dead-end desynchronisation dynamics

While Menck *et al.* (2014) [39] provide a rationale for why a desynchronised dead-end node should oscillate around an average value of $\langle \omega_i \rangle = P_i/\alpha$, they do not provide a full solution for the entire grid which would permit these dynamics, instead deriving this result for a single oscillating node attached to a fixed, synchronised grid. Here, we derive an asymptotic solution for such dead-end desynchronisation events, including the reaction of the rest of the grid. Label the isolated node as $i = 1$ and its only neighbour as $i = 2$, which is then connected in some way to the rest of the grid, $i = 3, \dots, N$. To make analytical progress, introduce the small parameter $\varepsilon = \alpha^2/P \ll 1$ (this parameter is $\varepsilon = 10^{-2}$ in the results above)

and scale the variables as follows: $\omega_i = P\hat{\omega}_i/\alpha$, $t = \alpha\hat{t}/P$, $P_i = P\hat{P}_i$ and $\theta_i = \hat{\theta}_i$, where the hat-decoration denotes the new, scaled variables. For ease of presentation, this hat-decoration will be dropped and all variables are the scaled variables in the remainder of this section unless specifically contradicted.

The dynamics in Fig. 6 exhibit multiple time-scales, the oscillation frequency of ω_i itself (for $i = 1$ and 2), along with the evolution of θ_i which is largely governed by the average of ω_i . A long time $\tau = \varepsilon^2 t$ is introduced to account for these distinct time-scales and an asymptotic analysis is performed on the resulting equations:

$$\theta_{i;t} + \varepsilon^2 \theta_{i;\tau} = \omega_i, \quad (17)$$

$$\begin{aligned} \omega_{i;t} + \varepsilon^2 \omega_{i;\tau} = & -\varepsilon \omega_i + \varepsilon P_i \\ & + \varepsilon \frac{K}{P} \sum_{j=1}^N A_{ij} \sin(\theta_j - \theta_i), \end{aligned} \quad (18)$$

where a semi-colon denotes partial differentiation and $\theta_i = \theta_i(t, \tau)$ and $\omega_i = \omega_i(t, \tau)$ in which the two time-scales t and τ are treated as independent variables.

First, consider the dead-end node $i = 1$. Expanding the angular frequency as $\omega_1 = \omega_{1,0} + \varepsilon \omega_{1,1} + O(\varepsilon^2)$ gives

$$O(\varepsilon^0): \quad \omega_{1,0;t} = 0, \quad (19)$$

$$O(\varepsilon^1): \quad \omega_{1,1;t} = -\omega_{1,0} + P_1 + \frac{K}{P} \sin(\theta_2 - \theta_1). \quad (20)$$

The first equation indicates that $\omega_{1,0} = \Omega_1(\tau)$ and $\theta_1 = \Omega_1(\tau)t + \Theta_1(\tau) + O(\varepsilon)$. The solution to (20) must remain bounded for asymptotic consistency; it will be shown shortly that $\theta_2 - \theta_1$ depends linearly on t (it is not a constant) and so the final term in (20) is t -periodic and leads to a t -periodic component in $\omega_{1,1}$. However, the other terms contributing to $\omega_{1,1;t}$ are constants (with respect to t) and so for asymptotic consistency, we require that $\omega_{1,0} = \Omega_1(\tau) = P_1$.

Now consider the main bulk of the grid, $i = 3, \dots, N$. This part of the grid is largely unaffected by the desynchronised dead-end node. As such, a solution in which $\theta_j - \theta_i = \Delta\theta_{ij}^s + \varepsilon D_{ij} + O(\varepsilon^2)$ is sought, so that the dynamics in the bulk of the

grid represent a small perturbation from the synchronised state. This leads to an expansion of the form $\omega_i = \varepsilon\omega_{i,1} + \varepsilon^2\omega_{i,2} + O(\varepsilon^3)$. Furthermore, in order that the nodes attached to the dead-end neighbour $i = 2$ remain close to their synchronised values, the phase at $i = 2$ must also remain close to its synchronised value. These expansions of $\theta_j - \theta_i$ (and hence also ω_i) must therefore be valid for $i, j = 2, 3, \dots, N$.

The swing equations for the main bulk of the grid, $i = 3, \dots, N$, therefore become

$$\begin{aligned} \omega_{i,t} + \varepsilon^2\omega_{i,\tau} &= -\varepsilon\omega_i + \varepsilon P_i \\ &+ \varepsilon \frac{K}{P} \sum_{j=2}^N A_{ij} \sin(\Delta\theta_{ij}^s) \\ &+ \varepsilon^2 \frac{K}{P} \sum_{j=2}^N A_{ij} D_{ij} \cos(\Delta\theta_{ij}^s) + O(\varepsilon^3) \quad (21) \\ &= -\varepsilon\omega_i + \varepsilon^2 \frac{K}{P} \sum_{j=2}^N A_{ij} D_{ij} \cos(\Delta\theta_{ij}^s) + O(\varepsilon^3), \end{aligned} \quad (22)$$

where the second equality arises from the relationship between P_i and the steady-state phases (3) given that θ_1 only contributes to the equation of ω_2 and so the fact that $j = 1$ is missing from the sums in these equations for $i = 3, \dots, N$ is irrelevant. Inserting the expansion for ω_i yields

$$\begin{aligned} O(\varepsilon^1) : \quad \omega_{i,1;t} &= 0, \quad (23) \\ O(\varepsilon^2) : \quad \omega_{i,2;t} &= -\omega_{i,1} + \frac{K}{P} \sum_{j=2}^N A_{ij} D_{ij} \cos(\Delta\theta_{ij}^s). \end{aligned} \quad (24)$$

The first equation results in $\omega_{i,1} = \Omega_i(\tau)$ and so $\theta_i = \theta_i^s + \varepsilon(\Omega_i(\tau)t + \Theta_i(\tau))$. The deviation of the phase-differences from their steady values is then $D_{ij} = (\Omega_j - \Omega_i)t + \Theta_j - \Theta_i$ (it will be shown that this is also the case for $j = 2$). For asymptotic consistency, D_{ij} must remain bounded in t and so $\Omega_i = \Omega_j \equiv \Upsilon$ for $i, j = (2), 3, \dots, N$. From this, the two terms on the right-hand-side of (24) are constants, and so for $\omega_{i,2}$ to remain bounded in t ,

there is a consistency condition relating Υ to D_{ij} :

$$\Upsilon = \frac{K}{P} \sum_{j=2}^N A_{ij} D_{ij} \cos(\Delta\theta_{ij}^s), \quad \text{for } i = 3, \dots, N. \quad (25)$$

Next, consider the dead-end neighbour node $i = 2$. Using the fact that, for $j = 3, \dots, N$, we have $\theta_j - \theta_2 = \Delta\theta_{2j}^s + \varepsilon D_{2j} + O(\varepsilon^2)$, the equation for ω_2 is

$$\begin{aligned} \omega_{2;t} + \varepsilon^2\omega_{2,\tau} &= -\varepsilon\omega_2 + \varepsilon P_2 + \varepsilon \frac{K}{P} \sin(\theta_1 - \theta_2) \\ &+ \varepsilon \frac{K}{P} \sum_{j=2}^N A_{2j} \sin(\Delta\theta_{2j}^s) \\ &+ \varepsilon^2 \frac{K}{P} \sum_{j=2}^N A_{2j} D_{2j} \cos(\Delta\theta_{2j}^s) + O(\varepsilon^3). \end{aligned} \quad (26)$$

Before making further progress, note that the swing equation (1) admits a global relation by adding all N equations together:

$$\frac{d}{dt} \left(\sum_{i=1}^N \omega_i \right) = - \sum_{i=1}^N \omega_i, \quad (27)$$

from which it is seen that the total angular frequency, $\sum_{i=1}^N \omega_i$, exponentially decreases to zero for all initial conditions. We will take this total angular frequency to be identically zero, given that the analysis is for a stable dead-end desynchronisation. In terms of the individual nodes of interest, this reads

$$\omega_1 + \omega_2 + \sum_{i=3}^N \omega_i = 0, \quad (28)$$

or, using the previous results,

$$\begin{aligned} \omega_2 &= -\omega_1 - \varepsilon \sum_{i=3}^N \Upsilon + O(\varepsilon^2) \quad (29) \\ &= -\omega_1 - \varepsilon \frac{K}{P} \sum_{i=3}^N \sum_{j=2}^N A_{ij} D_{ij} \cos(\Delta\theta_{ij}^s) + O(\varepsilon^2). \end{aligned} \quad (30)$$

For large N , this first expression is asymptotically consistent with $\omega_1 = O(\varepsilon^0)$ and $\omega_2 = O(\varepsilon^1)$ since $\varepsilon \sum_{i=3}^N \Upsilon = \varepsilon(N-2)\Upsilon$ may be $O(\varepsilon^0)$.

The combination $A_{ij}D_{ij} \cos(\Delta\theta_{ij}^s)$ is antisymmetric in i and j and so almost all of the sum in (30) cancels except for the terms with $j = 2$. Therefore,

$$\omega_2 = -\omega_1 - \varepsilon \frac{K}{P} \sum_{i=3}^N A_{i2} D_{i2} \cos(\Delta\theta_{i2}^s) + O(\varepsilon^2) \quad (31)$$

$$= -\omega_1 + \varepsilon \frac{K}{P} \sum_{i=3}^N A_{2i} D_{2i} \cos(\Delta\theta_{2i}^s) + O(\varepsilon^2). \quad (32)$$

The sum in this final expression is the same as that in the final term of (26) after replacing the dummy index i with j and noting that $A_{22} = 0$ (and multiplying by ε). Eliminating this sum in preference of ω_1 and ω_2 in (26) gives

$$\begin{aligned} \omega_{2;t} + \varepsilon^2 \omega_{2;\tau} &= \varepsilon \omega_1 + \varepsilon P_2 + \varepsilon \frac{K}{P} \sin(\theta_1 - \theta_2) \\ &+ \varepsilon \frac{K}{P} \sum_{j=2}^N A_{2j} \sin(\Delta\theta_{2j}^s) \end{aligned} \quad (33)$$

$$= \varepsilon \omega_1 + \varepsilon \frac{K}{P} \sin(\theta_1 - \theta_2) - \varepsilon \frac{K}{P} \sin(\Delta\theta_{21}^s), \quad (34)$$

where the second equality arises from the equation for P_2 in terms of the steady phases (3).

The expansion $\omega_2 = \varepsilon \omega_{2,1} + \varepsilon^2 \omega_{2,2} + O(\varepsilon^3)$ is substituted into (34) along with $\omega_1 = \omega_{1,0} + \varepsilon \omega_{1,1} + O(\varepsilon^2)$ and $\omega_{1,0} = \Omega_1 = P_1$ to give

$$\begin{aligned} O(\varepsilon^1) : \quad \omega_{2,1;t} &= P_1 + \frac{K}{P} \sin(\theta_1 - \theta_2) \\ &- \frac{K}{P} \sin(\Delta\theta_{21}^s) = \frac{K}{P} \sin(\theta_1 - \theta_2), \end{aligned} \quad (35)$$

$$\begin{aligned} \omega_{1,1;t} &= -\omega_{1,0} + P_1 + \frac{K}{P} \sin(\theta_2 - \theta_1) \\ &= \frac{K}{P} \sin(\theta_2 - \theta_1), \end{aligned} \quad (36)$$

where the simplification in the first expression arises from the equation for P_1 in terms of the steady phase $\Delta\theta_{12}$ (3).

Finally, recall that $\theta_1 = P_1 t + \Theta_1 + O(\varepsilon)$. Furthermore, since $\theta_j = \theta_j^s + \varepsilon(\Upsilon t + \Theta_j) + O(\varepsilon^2)$ for $j = 3, \dots, N$, we have that $D_{2j} = \Upsilon t + \Theta_j - \varepsilon^{-1}(\theta_2 - \theta_2^s) + O(\varepsilon)$ and so in order for D_{2j} to remain bounded in t and $O(\varepsilon^0)$ we must

have $\theta_2 = \theta_2^s + \varepsilon(\Upsilon t + \Theta_2(\tau) + F(t)) + O(\varepsilon^2)$ for some $F(t) = O(\varepsilon^0)$ uniformly in t . This leads to the leading-order form of ω_2 to be given by $\omega_{2,1} = \Upsilon + f(t)$ with $f(t) = F'(t)$, indicating that $\omega_2 = O(\varepsilon)$ varies around the same value as the rest of the grid, $\varepsilon \Upsilon$.

For times $t < O(\varepsilon^{-1})$, we may write $\theta_2 = \theta_2^s + O(\varepsilon)$, so that

$$O(\varepsilon^1) : \quad \omega_{2,1;t} = \frac{K}{P} \sin(P_1 t + \Theta_1 - \theta_2^s), \quad (37)$$

$$\omega_{1,1;t} = -\frac{K}{P} \sin(P_1 t + \Theta_1 - \theta_2^s), \quad (38)$$

which may be trivially integrated with respect to t .

Combining the results together, and defining $\Phi = \Theta_1(\tau) - \theta_2^s$, we arrive at

$$\omega_1 = P_1 + \varepsilon \frac{K}{P P_1} \cos(P_1 t + \Phi) + O(\varepsilon^2) \quad (39)$$

$$\omega_2 = \varepsilon \Upsilon - \varepsilon \frac{K}{P P_1} \cos(P_1 t + \Phi) + O(\varepsilon^2) \quad (40)$$

$$\omega_i = \varepsilon \Upsilon + O(\varepsilon^2), \quad \text{for } i = 3, \dots, N. \quad (41)$$

All that remains is to determine Υ , which may be obtained from the global constraint $\omega_1 + \omega_2 + \sum_{i=3}^N \omega_i = 0$, giving

$$\Upsilon = -\frac{P_1}{\varepsilon(N-1)}, \quad (42)$$

which may be $O(\varepsilon^0)$ when N is large (the present case with $N = 120$ and $\varepsilon = 10^{-2}$ gives $\Upsilon \approx -0.84 P_1$ with $|P_1| = 1$). The results for the un-scaled variables read

$$\omega_1 = \frac{P_1}{\alpha} + \frac{\alpha K}{P_1} \cos\left(\frac{P_1}{\alpha} t + \Phi\right) \quad (43)$$

$$\omega_2 = -\frac{P_1}{\alpha(N-1)} - \frac{\alpha K}{P_1} \cos\left(\frac{P_1}{\alpha} t + \Phi\right) \quad (44)$$

$$\omega_i = -\frac{P_1}{\alpha(N-1)}, \quad \text{for } i = 3, \dots, N. \quad (45)$$

These solutions are plotted along with the numerical results in Fig. 7(a). The phaseshift Φ is undetermined by the analysis and depends on how this solution is approached transiently; in Fig. 7(a), $\Phi = 0.689$. Furthermore, this value drifts at an almost constant rate on a slower timescale εt (Fig. 7(b)), as anticipated after neglecting $\varepsilon \Upsilon t$ in




Fig. 7 The dead-end desynchronisation converged to by the dynamics of Fig. 6(c) with (a) angular frequency ω_i for nodes $i = 3, \dots, N$ (grey), node $i = 1$ (blue) and node $i = 2$ (red). The asymptotic solutions represented by (43–45) are plotted with black dashed lines, using $\Phi = 0.689$. All variables are the original, un-scaled versions. (b) The undetermined phase Φ , computed by finding the best fit between (43) and the numerical data locally over a single period of oscillation, against the intermediate timescale εt .

the leading order phase θ_2 for times $t < O(\varepsilon^{-1})$. This drift could presumably be accounted for with a more sophisticated asymptotic expansion (e.g. using elements of WKB theory) but is of little importance here. Satisfactory agreement is seen with the numerical results, particularly for ω_1 . The amplitude of the oscillation of ω_2 is not quite right, and this is likely because its two other (non dead-end) neighbours also exhibit a small oscillation in ω_i which was not accounted for in the analysis. The analysis also implicitly assumes that $K/P = O(\varepsilon^0)$, and here $K/P = 8$ which is perhaps a little large, though $\varepsilon K/P = 0.08$ is (somewhat) small.

Interestingly, (43) is identical to the result for a single oscillator coupled to a fixed, synchronised grid derived by Menck *et al.* (2014) [39]. The derivation presented here shows how the same dynamics arise asymptotically from a dead-end node attached to a grid which responds to the dead-end, and also shows how the rest of the grid behaves. The rest of the grid shifts its mean

frequency by a uniform amount in order to preserve the total angular frequency $\sum_{i=1}^N \omega_i = 0$ in an average sense by compensating for the elevated mean frequency of the dead-end node, and the dead-end node’s neighbour oscillates with the same amplitude and frequency as the dead-end node, with a phase shift of π between them, to ensure that the total angular frequency is instantaneously preserved.

A similar analysis can be performed for dead-ends formed of a small group of nodes, and is presented in Appendix C. However, the resulting equations do not admit an analytic solution, and suggest a rich variety of such dead-end dynamics, as found by Menck *et al.* (2014) [39] and Halekotte *et al.* (2021) [18].

6 Conclusions

Nonlinear measures of stability are required when examining the robustness of a desired behaviour in multi-stable systems. The size and shape of the basin of attraction of desired and undesired states

must be computed in order to provide such measures. Basin stability [19] provides estimates of the size of these basins of attraction (and can be interpreted in terms of the probability of a random initial condition evolving towards the desired state), and can give some indication of the shape of the basin of attraction if initial conditions are sorted in some state-space norm [25]. However, these coarse-grained estimates miss small-scale features of the boundary between different basins of attraction and the closest approach of this boundary to the desired dynamics can have no relation to the overall basin volume [26]. Initial conditions at this closest approach of the basin boundary to the desired dynamics represent the minimal disturbances to the desired behaviour that can trigger transition to undesired dynamics, and can be computed using a variational method [28] and has recently been explored in some networked dynamical systems, including the ‘swing equation’ for (simplified) power-grid synchronisation [38].

This current work provided a direct comparison between basin stability and the amplitude of minimal disturbances to cause desynchronisation in the swing equation for a variety of grid layouts, namely three simple four-node grids ($G1$, $G2$ and $G4$) [20] and a complex model UK grid [22]. The algorithm for computing minimal disturbances was presented in detail, including an improved convergence criteria in comparison to Halekotte & Feudel (2020) [38]. For two of the simple four-node grids ($G1$ and $G4$), the minimal disturbance initial condition drives the system towards a steady, synchronised, unstable saddle point on the basin boundary separating stable synchronised and desynchronised states, before transitioning to long-term desynchronisation. Each node in $G1$ is connected to exactly two out of the three other nodes, whereas in $G4$ each node is connected to all three of the other nodes. This leads to the minimal disturbance in $G1$ being largely confined to half the grid, propagating across the grid as the system evolves, whereas the minimal disturbance in $G4$ takes the isolated initial condition of $G1$ and mirrors it across the other half of the grid, exciting the entire grid at once and leading to a minimal disturbance amplitude almost exactly twice that of $G1$.

The dynamics of $G2$ are substantially different. Exactly two of the four nodes are connected

to all three other nodes, and this asymmetric connection in the grid causes the dynamics of these two nodes to be completely identical throughout the evolution of the minimal disturbance. This also results in the minimal disturbance amplitude varying non-monotonically with the number of connections in the network, with grid $G2$ having the largest such amplitude, and emphasises the impact of details of the network topology on the system’s nonlinear dynamics.

Simple grid $G1$ was investigated in further detail, comparing the minimal disturbance amplitude with the synchronisation probabilities of Ji & Kurths (2014) [20] over a range of nodal powers and damping rates. Parameter values for which near-zero and near-certain synchronisation probabilities are computed by Ji & Kurths (2014) [20] carry no meaning for the minimal disturbance amplitude, which varies smoothly through these points. This observation demonstrates that, just as in the turbulence transition problem [26], the coarse-graining associated with basin stability computations overlooks small-scale features of basin boundaries and provides an incomplete picture of the nonlinear stability of a desired behaviour. While the presented algorithm can only identify *local* optimal initial conditions, finding *any* such ‘worst-case’ perturbation represents an improvement over basin stability computations if identifying worst-case scenarios is critical. Though it is unlikely that a random initial condition matches the minimal disturbance, it is nevertheless a possible route away from the desired dynamics and must be accounted for when maintaining the desired behaviour is of utmost importance for the application in mind. Indeed, these two nonlinear measures provide complementary information and should be used together when making relevant design decisions.

In isolation, the computational time required to find a minimal disturbance cannot be known *a priori*, and this computational time can be quite short, or become unmanageably long for some parameter values if the saddle-point dividing the two basins of attraction becomes overly complex. However, the comparison with basin stability results demonstrated that there is a potential link between this increased complexity (or at least with the increased convergence time of the minimal disturbance algorithm) with the basin stability’s synchronisation probability approaching one.

A minimal disturbance in the complex model UK grid with parameters corresponding to the basin stability computations of Mitra *et al.* (2017) [22] was computed. The minimal disturbance requires the entire grid to be perturbed, before the dynamics settle down onto an isolated ‘dead-end’ desynchronisation in which a node with a single connection to the rest of the grid oscillates around an elevated frequency. The dynamics of the dead-end desynchronisation were investigated in further detail, and an asymptotic expansion was developed to formalise the estimate of Menck *et al.* (2014) [39]. The resulting solution shows how the zero total angular frequency of the grid is maintained by a small frequency perturbation across the rest of the grid which balances the elevated mean frequency of the dead-end node, and the dead-end’s single neighbouring node has an angular frequency which oscillates exactly out-of-phase with the dead-end node to instantaneously ensure zero total angular frequency. A rough estimate of the minimum size of whole-grid perturbations required to desynchronise the grid using the results of Mitra *et al.* (2017) [22] is comparable to the minimal disturbance amplitude, though this correspondence should not be over-interpreted, owing to the assumption made of a uniform synchronisation probability distribution in perturbation energy, which is unlikely to be realistic [25].

In summary, both basin stability and minimal disturbance amplitudes provide important, and often complementary information about the non-linear stability of desired dynamics in multi-stable systems, and this work advocates for the use of *both* measures when analysing nonlinear systems. In the case of power-grid dynamics specifically, minimal disturbances in models which more realistically account for the power grid dynamics, such as the normal form model [6, 7] or the third-order model [1] should be investigated, in addition to the swing equation itself with the added complexities of non-equal properties (nodal power, damping, line admittance) across the grid. The underlying methodologies translate readily across applications, and improved algorithms using a Bayesian approach for basin stability [26] and a penalised optimisation procedure [42] that may be incorporated into the computation of minimal disturbances to improve its convergence rate have recently been put forward to improve convergence.

Both methods present computational challenges, particularly as the system size (number of nodes) grows and if the state space becomes complex. For the case of minimal disturbances, improvements to the algorithm are required to find other local optimal solutions, perhaps through penalisation of trajectories that move close to already discovered optimal initial conditions. It is evident that further investigation of the interplay (or lack thereof) between these nonlinear measures is warranted, along with further improvements to the associated algorithms.

Supplementary information. The code used to compute minimal disturbances is deposited here: https://github.com/TomEaves/DAL_swing.

Acknowledgements. Preliminary computations of some minimal disturbances using an alternative algorithm for reducing initial perturbation energy and an alternative (non-invariant) definition of potential energy were performed by Mr. Saad Kashaf (BEng Dundee, 2021) as part of their Bachelor’s thesis. Two anonymous reviewers are thanked for their constructive comments that helped improve the manuscript.

Declarations

Not applicable.

Appendix A Steady solutions and linear dynamics

The graph $G1$ has adjacency matrix given by

$$\mathbf{A} = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}. \quad (\text{A1})$$

This form of adjacency matrix, along with the relation $S_{ij} \equiv \sin(\Delta\theta_{ij}^s) = -\sin(\Delta\theta_{ji}^s)$ means that the steady state equations

$$K \sum_{j=1}^N A_{ij} \sin(\Delta\theta_{ij}^s) = -P_i \quad \text{for } i = 1, \dots, N \quad (\text{A2})$$

depend on only 4 independent synchronous phase differences $\Delta\theta_{12}^s$, $\Delta\theta_{13}^s$, $\Delta\theta_{24}^s$ and $\Delta\theta_{34}^s$. The sine of these phase differences then satisfy

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \\ 0 & 0 & -1 & -1 \end{pmatrix} \begin{pmatrix} S_{12} \\ S_{13} \\ S_{24} \\ S_{34} \end{pmatrix} = \begin{pmatrix} R \\ -R \\ -R \\ R \end{pmatrix}, \quad (\text{A3})$$

where $R = P/K$. These equations are underdetermined. Their solution is given by

$$\begin{pmatrix} S_{12} \\ S_{13} \\ S_{24} \\ S_{34} \end{pmatrix} = \frac{R}{2} \begin{pmatrix} 1 \\ 1 \\ -1 \\ -1 \end{pmatrix} + \lambda \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \end{pmatrix} \quad (\text{A4})$$

where λ is a constant. To ensure that $|S_{ij}| \leq 1$, we must have

$$-1 + \frac{R}{2} \leq \lambda \leq 1 - \frac{R}{2}, \quad (\text{A5})$$

which also demonstrates that no synchronous solutions exist if $R > 2$.

Inverting sin adds further solutions. In particular, if $\Delta\theta_{ij}^s \in [-\pi, \pi)$, and using the convention that $\sin^{-1}(x) \in [-\pi/2, \pi/2]$, we introduce primary and secondary solutions for each phase difference:

$$\begin{aligned} \Delta\theta_{ij}^{s,1} &= \sin^{-1}(S_{ij}), \\ \Delta\theta_{ij}^{s,2} &= \text{sgn}(\sin^{-1}(S_{ij}))\pi - \sin^{-1}(S_{ij}). \end{aligned} \quad (\text{A6})$$

Finally, the phases themselves are found by inverting

$$\mathbf{M} \begin{pmatrix} \theta_1^s \\ \theta_2^s \\ \theta_3^s \\ \theta_4^s \end{pmatrix} \equiv \begin{pmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \\ 0 & 0 & -1 & 1 \end{pmatrix} \begin{pmatrix} \theta_1^s \\ \theta_2^s \\ \theta_3^s \\ \theta_4^s \end{pmatrix} = \begin{pmatrix} \Delta\theta_{12}^s \\ \Delta\theta_{13}^s \\ \Delta\theta_{24}^s \\ \Delta\theta_{34}^s \end{pmatrix}. \quad (\text{A7})$$

The matrix \mathbf{M} is singular, and these equations are also underdetermined. Their solution may be written as

$$\begin{pmatrix} \theta_1^s \\ \theta_2^s \\ \theta_3^s \\ \theta_4^s \end{pmatrix} = \mathbf{M}^+ \begin{pmatrix} \Delta\theta_{12}^s \\ \Delta\theta_{13}^s \\ \Delta\theta_{24}^s \\ \Delta\theta_{34}^s \end{pmatrix} + \beta \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad (\text{A8})$$

where \mathbf{M}^+ is the Moore–Penrose pseudoinverse of the matrix \mathbf{M} and β is an arbitrary constant that accounts for the translational symmetry $\theta_i \mapsto \theta_i + \beta$. This solution is valid provided that $(\Delta\theta_{12}^s, \Delta\theta_{13}^s, \Delta\theta_{24}^s, \Delta\theta_{34}^s)$ is in the image of \mathbf{M} , which is the case here for all 16 combinations of solution branches (i.e. for all combinations of primary and secondary solutions for each of the four steady phase differences).

The unstable synchronous state plotted in Fig. 2(c), with $R = 0.8$ so that $-0.6 \leq \lambda \leq 0.6$, has $\lambda = 0.5928$, branches $\Delta\theta_{12}^{s,2}$, $\Delta\theta_{13}^{s,2}$, $\Delta\theta_{24}^{s,1}$ and $\Delta\theta_{34}^{s,1}$, and $\beta = 0$. The linearly stable synchronous state about which the energy $E(t)$ is measured has $\lambda = \beta = 0$ and primary solutions for each phase difference.

The dynamics seen in Fig. 2(b) as the trajectory leaves the unstable steady state can be described in more detail by observing that, after linearisation, the equations of motion are

$$\begin{aligned} \ddot{\theta}_i + \alpha\dot{\theta}_i + \left[K \sum_{j=1}^N A_{ij} \cos \Delta\theta_{ij}^s \right] \theta_i \\ = K \sum_{j=1}^N A_{ij} \theta_j \cos(\Delta\theta_{ij}^s), \end{aligned} \quad (\text{A9})$$

which take the form of damped harmonic oscillation of each node, forced by neighbouring nodes. However, the ‘restoring’ force $\left[K \sum_{j=1}^N A_{ij} \cos \Delta\theta_{ij}^s \right] \theta_i \equiv \mathcal{K}_i \theta_i$ is not necessarily restoring since $\mathcal{K}_i < 0$ is possible. Indeed, the two branches of solutions satisfy $\cos(\Delta\theta_{ij}^{s,1}) > 0$ and $\cos(\Delta\theta_{ij}^{s,2}) < 0$. For the steady state plotted in Fig. 2(c), we have $\mathcal{K}_1 < \mathcal{K}_3 < 0 < \mathcal{K}_2 < \mathcal{K}_4$. This ordering is reflected in Fig. 2(b); the first node to depart from the steady state is mode 1, followed by node 3, then node 2, and finally node 4, in the order of largest destabilising force $\mathcal{K}_1\theta_1$ through to largest restoring force $\mathcal{K}_4\theta_4$. Given that the nodes are coupled together, any slight initial disturbance in node 1 will grow exponentially, and will eventually cause growth in all other nodes even if they were initially completely stationary.

Whilst in theory, similar exercises may be carried out for graphs $G2$ and $G4$, the equivalent expressions to (A3) involve non-square matrices, introducing further dimensions to the solution

families. The details of these solutions will not be computed here.

Appendix B Computing the Lagrange multiplier c .

To determine the Lagrange multiplier c , the steepest ascent update (13–14) must lead to a new set of initial conditions $\{\theta_i^{0,n}, \omega_i^{0,n}\}_{i=1}^N$ with energy E_0 :

$$E_0 = \frac{1}{2N} \sum_{i=1}^n |\omega_i^{0,n}|^2 + \frac{1}{4N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} |\theta_j^{0,n} - \theta_i^{0,n} - \Delta\theta_{ij}^s|^2. \quad (\text{B10})$$

Inserting the steepest ascent update (13–14) and using the fact that the previous set of initial conditions $\{\theta_i^{0,n-1}, \omega_i^{0,n-1}\}_{i=1}^N \equiv \{\theta_i^0, \omega_i^0\}_{i=1}^N$ also have energy E_0 , rearranging (B10) gives a quadratic equation for c :

$$\begin{aligned} & \frac{4\epsilon}{N} \left[2 \sum_{i=1}^N |\omega_i^0|^2 + K \sum_{i=1}^N \sum_{j=1}^N A_{ij} C_{ij}^2 \right] c^2 \\ & + 4 \left[K \sum_{i=1}^N \sum_{j=1}^N A_{ij} (\theta_j^0 - \theta_i^0 - \Delta\theta_{ij}^s) C_{ij} \right. \\ & + \frac{\epsilon K}{N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} (\varphi_j^0 - \varphi_i^0) C_{ij} - 2 \sum_{i=1}^N |\omega_i^0|^2 \\ & \left. - \frac{2\epsilon}{N} \sum_{i=1}^N \eta_i^0 \omega_i^0 \right] c + 4 \sum_{i=1}^N \eta_i^0 \omega_i^0 \frac{2\epsilon}{N} \sum_{i=1}^N |\eta_i^0|^2 \\ & + 2K \sum_{i=1}^N \sum_{j=1}^N A_{ij} (\theta_j^0 - \theta_i^0 - \Delta\theta_{ij}^s) (\varphi_j^0 - \varphi_i^0) \\ & + \frac{\epsilon K}{N} \sum_{i=1}^N \sum_{j=1}^N A_{ij} (\varphi_j^0 - \varphi_i^0)^2 = 0, \quad (\text{B11}) \end{aligned}$$

where

$$C_{ij} = \sum_{k=1}^N (A_{jk} (\theta_k^0 - \theta_j^0 - \Delta\theta_{jk}^s) - A_{ik} (\theta_k^0 - \theta_i^0 - \Delta\theta_{ik}^s)). \quad (\text{B12})$$

In general, this quadratic equation for c may have up to two real solutions. If there are no real

solutions, then a reduction in ϵ can be made to guarantee that a real solution exists. When two real solutions are available, the one with smallest magnitude is sought, as this is the solution leading to new initial conditions closest to the previous set of initial conditions, as expected from an iterative gradient-based optimisation method (see [43] for a geometric interpretation of this choice).

Appendix C Multiple-node dead-end desynchronisation

The analysis of section 5 can be extended to a small dead-end group of nodes in a similar manner. In particular, if $k \ll N$ nodes form a dead-end, linked via a single node $k+1$ to the rest of the grid, then (in terms of the scaled variables)

$$\omega_i = \Omega_i(\tau) + \epsilon\omega_{i,1} + O(\epsilon^2) \quad \text{for } i = 1, \dots, k \quad (\text{C13})$$

$$\omega_{k+1} = \epsilon\Upsilon + \epsilon\omega_{k+1,1} + O(\epsilon^2), \quad (\text{C14})$$

$$\omega_i = \epsilon\Upsilon + O(\epsilon^2) \quad \text{for } i = k+2, \dots, N. \quad (\text{C15})$$

The main grid frequency offset is

$$\Upsilon = -\frac{\sum_{i=1}^k \Omega_i}{\epsilon(N-k)}, \quad (\text{C16})$$

which requires that $\epsilon(N-k) = O(\epsilon^0)$, which is achievable when $1 \leq k \ll N$. No clear statement about each individual $\Omega_i(\tau)$ can be made, as they arise from the solution to the equations for $\omega_{i,1}$, for $i = 1, \dots, k$ along with $\omega_{k+1,1}$ (given below), which are somewhat more complex than the equivalent equations when $k = 1$ and appear (numerically) to admit several different solutions depending on their initial conditions.

The structure of the result for a k -node dead-end has a similar form to the single-node dead-end. The bulk of the grid (nodes $i = k+2, \dots, N$) have their frequency shifted by an amount $\epsilon\Upsilon$ which is related to balancing the total net power associated with the k dead-end nodes. The single node which connects the dead-end group to the rest of the grid (node $k+1$) has the same mean frequency of $\epsilon\Upsilon$, but oscillates around this frequency.

For the single-node dead-end, this oscillation is due to the activity of the dead-end node, whereas for a multi-node dead-end, this oscillation depends on the dynamics of the dead-end group.

The dynamics within the dead-end and its single neighbour are complex. Each dead-end node has a frequency varying on the slow time-scale τ at leading order ($\Omega_i(\tau)$), but this frequency does not straightforwardly relate to a given node's power P_i . The $O(\varepsilon)$ variation around this leading order frequency is given by

$$\begin{aligned} \omega_{i,1;t} = & -\Omega_i(\tau) + P_i \\ & + \frac{K}{P} \sum_{j=1}^k A_{ij} \sin((\Omega_j - \Omega_i)t + \Theta_j - \Theta_i) \\ & - \frac{K}{P} A_{i,k+1} \sin(\Omega_i t + \Theta_i - \theta_{k+1}^s), \end{aligned} \quad (\text{C17})$$

and the $O(\varepsilon)$ variation around $\varepsilon\Upsilon$ in its single neighbour is given by

$$\begin{aligned} \omega_{k+1;t} = & \sum_{i=1}^k \Omega_i - \frac{K}{P} \sum_{j=1}^k A_{k+1,j} \sin(\Delta\theta_{k+1,j}^s) \\ & + \frac{K}{P} \sum_{j=1}^k A_{k+1,j} \sin(\Omega_j t + \Theta_j - \theta_{k+1}^s). \end{aligned} \quad (\text{C18})$$

The dynamics of the dead-end are self-contained, and this drives the oscillation at node $k+1$. Consistency conditions are not possible to apply in general. Numerical experiments reveal multiple solutions in the dead-end (for example the cluster at the top-right of the UK grid in Fig. 6(a)) in which connected nodes may have the same or different values of Ω_i , and these values are typically observed to be $\pm P_i$ or 0. These observations mean that any of the sine terms in (C17–C18) which at first glance appear periodic in t (and hence not part of the consistency condition) may in fact be constant, and so they may or may not be included in the consistency condition depending on the exact dead-end solution under consideration.

If none of the leading order dead-end frequencies are zero, then it can be deduced that $\sum_{i=1}^k \Omega_i = \sum_{i=1}^k P_i$, by adding all k equations (C17) and applying the consistency condition. Halekotte & Feudel (2020) [38] present a particular multi-node dead-end desynchronisation for

which *all* nodes in the dead-end have the same non-zero leading order frequency. Using $\Omega_i(\tau) \equiv \bar{\omega}$ and the steady state phase relation (3) leads to

$$\bar{\omega} = \frac{K}{kP} \sum_{j=1}^k A_{k+1,j} \sin(\Delta\theta_{k+1,j}^s) \quad (\text{C19})$$

which determines the dead-end frequency in terms of the synchronised steady phase differences.

The desynchronised leading-order phase differences are then given by the asymptotic consistency of (C17) given that the sine connections within the dead-end all become constants:

$$0 = -\bar{\omega} + P_i + \frac{K}{P} \sum_{i=1}^k \sin(\Theta_j - \Theta_i). \quad (\text{C20})$$

The leading order phase differences in the dead-end satisfy a similar relation to the phase differences in the synchronised state (3), except that the connection to the rest of the grid via node $k+1$ plays no role, and the powers $P_i \mapsto P_i - \bar{\omega}$. The full solution for the entire grid in this case is

$$\begin{aligned} \omega_i = & \bar{\omega} + \varepsilon \frac{K}{P\bar{\omega}} A_{i,k+1} \cos(\bar{\omega}t + \Phi_i) + O(\varepsilon^2) \\ & \text{for } i = 1, \dots, k \end{aligned} \quad (\text{C21})$$

$$\begin{aligned} \omega_{k+1} = & \varepsilon\Upsilon - \varepsilon \frac{K}{P\bar{\omega}} \sum_{j=1}^k A_{k+1,j} \cos(\bar{\omega}t + \Phi_j) \\ & + O(\varepsilon^2), \end{aligned} \quad (\text{C22})$$

$$\omega_i = \varepsilon\Upsilon + O(\varepsilon^2) \quad \text{for } i = k+2, \dots, N, \quad (\text{C23})$$

with

$$\Upsilon = -\frac{k\bar{\omega}}{\varepsilon(N-k)}, \quad (\text{C24})$$

and $\Phi_j = \Theta_j(\tau) - \theta_{k+1}^s$, which has a similar form to the single-node dead end solution (39–41).

References

- [1] Machowski, J., Bialek, J. W. & Bumby, J. *Power system dynamics: stability and control* (Wiley, 2011).
- [2] Nishikawa, T. & Motter, A. E. Comparative analysis of existing models for power-grid synchronization. *New Journal of Physics* **17**, 015012 (2015).

- [3] Auer, S., Kleis, K., Schultz, P., Kurths, J. & Hellmann, F. The impact of model detail on power grid resilience measures. *The European Physical Journal Special Topics* **225**, 609–625 (2016).
- [4] Hellmann, F. *et al.* Network-induced multistability through lossy coupling and exotic solitary states. *Nature Communications* **11**, 1–9 (2020).
- [5] Büttner, A., Kurths, J. & Hellmann, F. Ambient forcing: sampling local perturbations in constrained phase spaces. *New Journal of Physics* **24**, 053019 (2022).
- [6] Kogler, R., Plietzsch, A., Schultz, P. & Hellmann, F. Normal form for grid-forming power grid actors. *PRX Energy* **1**, 013008 (2022).
- [7] Büttner, A. & Hellmann, F. Complex couplings - a universal, adaptive, and bilinear formulation of power grid dynamics. *PRX Energy* **3**, 013005 (2024).
- [8] Nnoli, K. P. & Kettemann, S. Spreading of disturbances in realistic models of transmission grids in dependence on topology, inertia and heterogeneity. *Scientific Reports* **11**, 23742 (2021).
- [9] Filatrella, G., Nielsen, A. H. & Pederson, N. F. Analysis of a power grid using a Kuramoto-like model. *The European Physical Journal B* **61**, 485–491 (2008).
- [10] Bhatta, K., Hayat, M. M. & Sorrentino, F. Modal decomposition of the linear swing equation in networks with symmetries. *IEEE Transactions on Network Science and Engineering* **8**, 2482–2494 (2021).
- [11] Motter, A. E., Myers, S. A., Anghel, M. & Nishikawa, T. Spontaneous synchrony in power-grid networks. *Nature Physics* **9**, 191–197 (2013).
- [12] Ji, P., Lu, W. & Kurths, J. Stochastic basin stability in complex networks. *EPL* **122**, 40003 (2018).
- [13] Roberts, L. G. W., Champneys, A. R., Bewll, K. R. W. & di Bernardo, M. Analytical approximations of critical clearing time for parametric analysis of power system transient stability. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* **5**, 465–476 (2015).
- [14] Mitra, C., Kittel, T., Choudhary, A., Kurths, J. & Donner, R. V. Recovery time after localized perturbations in complex dynamical networks. *New Journal of Physics* **19**, 103004 (2017).
- [15] Hellmann, F., Schultz, P., Grabow, C., Heitzig, J. & Kurths, J. Survivability of deterministic dynamical systems. *Scientific Reports* **6**, 29654 (2016).
- [16] Alvares, C. & Banerjee, S. How network topology affects the strength of dangerous power grid perturbations. *arXiv 2401.00552* (2023).
- [17] Manik, D., Timme, M. & Witthaut, D. Cycle flows and multistability in oscillatory networks. *Chaos* **27**, 083123 (2017).
- [18] Halekotte, L., Vanselow, A. & Feudel, U. Transient chaos enforces uncertainty in the British power grid. *Journal of Physics: Complexity* **2**, 035015 (2021).
- [19] Menck, P. J., Heitzig, J., Marwan, N. & Kurths, J. How basin stability complements the linear-stability paradigm. *Nature Physics* **9**, 89–92 (2013).
- [20] Ji, P. & Kurths, J. Basin stability of the Kuramoto-like model in small networks. *The European Physical Journal Special Topics* **223**, 2483–2491 (2014).
- [21] Schultz, P., Heitzig, J. & Kurths, J. Detours around basin stability in power networks. *New Journal of Physics* **16**, 125001 (2014).
- [22] Mitra, C., Choudhary, A., Sinha, S., Kurths, J. & Donner, R. V. Multiple-node basin stability in complex dynamical systems. *Physical Review E* **95**, 032317 (2017).

- [23] Nauck, C. *et al.* Predicting basin stability of power grids using graph neural networks. *New Journal of Physics* **24**, 043041 (2022).
- [24] van Kan, A., Jegminat, J., Donges, J. F. & Kurths, J. Constrained basin stability for studying transient phenomena in dynamical systems. *Physical Review E* **93**, 042205 (2016).
- [25] Pershin, A., Beaume, C. & Tobias, S. M. A probabilistic protocol for the assessment of transition and control. *Journal of Fluid Mechanics* **895**, A16 (2020).
- [26] Pershin, A., Beaume, C., Eaves, T. S. & Tobias, S. M. Optimizing the control of transition to turbulence using a bayesian method. *Journal of Fluid Mechanics* **941**, A25 (2022).
- [27] Skufca, J. D., Yorke, J. A. & Eckhardt, B. Edge of chaos in a parallel shear flow. *Physical Review Letters* **96**, 174101 (2006).
- [28] Kerswell, R. R. Nonlinear nonmodal stability theory. *Annual Review of Fluid Mechanics* **50**, 319–345 (2018).
- [29] Lecoanet, D. & Kerswell, R. R. Connection between nonlinear energy optimization and instantons. *Physical Review E* **97**, 012212 (2018).
- [30] Monokrousos, A., Bottaro, A., Brandt, L., Di Vita, A. & Henningson, D. S. Nonequilibrium thermodynamics and the optimal path to turbulence in shear flows. *Physical Review Letters* **106**, 134502 (2011).
- [31] Pringle, C. C. T., Willis, A. P. & Kerswell, R. R. Minimal seeds for shear flow turbulence: using nonlinear transient growth to touch the edge of chaos. *Journal of Fluid Mechanics* **702**, 415–443 (2012).
- [32] Rabin, S. M. E., Caulfield, C. P. & Kerswell, R. R. Triggering turbulence efficiently in plane Couette flow. *Journal of Fluid Mechanics* **712**, 244–272 (2012).
- [33] Eaves, T. S. & Caulfield, C. P. Disruption of SSP/VWI states by a stable stratification. *Journal of Fluid Mechanics* **784**, 548–564 (2015).
- [34] Juniper, M. P. Transient growth and triggering in the horizontal Rijke tube. *International Journal of Spray Combustion Dynamics* **3**, 209–224 (2011).
- [35] Willis, A. P. On the optimised magnetic dynamo. *Physical Review Letters* **109**, 251101 (2012).
- [36] Foures, D. P. G., Caulfield, C. P. & Schmid, P. J. Optimal mixing in two-dimensional plane Poiseuille flow at high Péclet number. *Journal of Fluid Mechanics* **748**, 241–277 (2014).
- [37] Marcotte, F. & Caulfield, C. P. Optimal mixing in two-dimensional stratified plane Poiseuille flow at finite Péclet and Richardson numbers. *Journal of Fluid Mechanics* **853**, 539–385 (2018).
- [38] Halekotte, L. & Feudel, U. Minimal fatal shocks in multistable complex networks. *Scientific Reports* **10**, 11783 (2020).
- [39] Menck, P. J., Heitzig, J., Kurths, J. & Schellnhuber, H. J. How dead ends undermine power grid stability. *Nature Communications* **5**, 3969 (2014).
- [40] Cherubini, S., De Palma, P., Robinet, J.-C. & Bottaro, A. Rapid path to transition via nonlinear localized optimal perturbations in a boundary-layer flow. *Physical Review E* **82**, 066302 (2010).
- [41] van den Doel, K. & Ascher, U. The chaotic nature of faster gradient descent methods. *Journal of Scientific Computing* **51**, 560–581 (2012).
- [42] Chung, S. W. & Freund, J. B. An optimization method for chaotic turbulent flow. *Journal of Computational Physics* 111077 (2022).
- [43] Foures, D. P. G., Caulfield, C. P. & Schmid, P. J. Localization of flow structures using infinity-norm optimization. *Journal of Fluid*

Mechanics **729**, 672–701 (2013).