

Model reduction for analysis of cascading failures in power systems

Pablo A. Parrilo* Sanjay Lall* Fernando Paganini†
George C. Verghese‡ Bernard C. Lesieutre‡ Jerrold E. Marsden*

Abstract

In this paper, we apply a principal-orthogonal decomposition based method to the model reduction of a hybrid, nonlinear model of a power network. The results demonstrate that the sequence of fault events can be evaluated and predicted without necessarily simulating the whole system.

Keywords: Cascading failures, power systems, swing equations, Karhunen-Loève decomposition, nonlinear model reduction.

1 Introduction

Power systems are among the largest and most intricate engineered systems in operation today. The sheer complexity of these large-scale networks makes the planning, management and operation a difficult task, especially in the presence of multiple decision agents. These distributed agents perform both command and regulatory actions, and usually do so in an independent but cooperative fashion.

Upcoming changes in the power industry will force companies to adopt even more complicated software and management techniques to satisfy increasing demands and cooperation requirements. Obviously, there is a strong need to guarantee that these changes will not have undesirable consequences, such as network-wide power oscillations, instabilities or major cascading failures. Additionally, it would be advantageous to have guidelines about failure dynamics, and the possible mechanisms for stopping a fault cascade, or even reversing it once it started. As is becoming increasingly clear, these second-order effects are not easy (or even possible) to predict. The heterogeneity of the dynamics, and the lack of a centralized authority have been long-standing obstacles in the formulation of a unified theory.

In this paper we present some initial investigations on the possibility of applying recently developed model reduction tools to the study of cascading failures in power networks (see Figure 1).

The focus of this paper is on the construction of low-order models which closely approximate the global behavior of the hybrid nonlinear power network. There is a growing

*Control & Dynamical Systems, California Institute of Technology.

†Electrical Engineering, University of California - Los Angeles.

‡Electrical Engineering & Computer Science, Massachusetts Institute of Technology.

recognition of the strong need for rapid and reliable computation of the system dynamics. We will demonstrate in this paper that recent developments in model reduction techniques can contribute to these goals.

The paper is organized as follows: in Section 2 we present the swing equation model used throughout the paper. Section 3 demonstrates the cascading failures obtained in the simulations of the system model. In Section 4, an overview of the model reduction procedure introduced in [5] is presented, and the following section shows the promising results obtained from the application of this methodology to the power network model used in Section 3. Finally, the last section presents some conclusions and discussion of future work.

2 Modeling

2.1 Swing equations

Our starting point will be the study of so-called *swing dynamics*, representing the coupling between real power flows and frequency variations across the transmission network of the power system. This is the most natural starting point for a study of geographically extensive dynamic phenomena on the power system. Swing dynamics, which occur on the time scale of seconds, can span the system even in a small-signal setting, whereas voltage dynamics tends to extend across the system only when significantly mediated by nonlinear effects. Swing dynamics are also potentially able to interact with particular protection and control mechanisms in a way that could plausibly lead to cascading events.

A typical swing-equation model [1, 4, 8] involves a second-order differential equation (essentially Newton's law in rotational form, for the dynamics of the rotor) associated with each generator node or bus, and an algebraic equation associated with each load bus. Assuming that all bus voltages are well regulated, we can write the following equation for the i th bus:

$$M_i \ddot{\delta}_i + D_i \dot{\delta}_i = P_{mi} - P_{gi}, \quad (i = 1, \dots, n) \quad (1)$$

where δ_i are the generator rotor angle deviations with respect to a synchronously rotating reference frame, P_{mi} is the mechanical power input, P_{gi} is the electrical power output, and n is the number of generators. The parameters M_i and D_i are the i th-generator's normalized inertia and damping coefficient, respectively. The instantaneous frequency deviation from synchronous frequency at bus i is given by $\dot{\delta}_i(t)$. The supplied mechanical power P_{mi} is assumed constant in this simple model, but more generally is adjusted by a speed-dependent governor. The damping coefficient D_i reflects electromechanical damping in the generator itself, as well as the effective damping introduced by other components such as power system stabilizers, which are not directly represented in the swing model.

The remaining model equations are given by the load and network characteristics. The expression for the electrical power output is given by:

$$P_{gi} = \text{Re}[V_i^* I_i] = \text{Re}[V_i^* \sum_j Y_{ij} V_j] = - \sum_j |V_i| |V_j| b_{ij} \sin(\delta_i(t) - \delta_j(t)), \quad (i = 1, \dots, n)$$

where $V_i = |V_i| e^{j\delta_i}$, and $Y = G + jB$ is the admittance matrix representing the network connections. We assume throughout that the voltage magnitudes do not change, and the transmission lines are lossless, i.e. Y is purely imaginary ($G = 0$). In the particular case




Figure 1: Sequence of failures in simulation of swing-equation models. The black circles represent working generators, and gray lines represent failures.

presented in this paper, the loads are modeled as synchronous motors, with small inertia and small damping constants.

The quantity b_{ij} is the susceptance of the line between buses i and j , and is therefore 0 if these two buses are not connected (and is also 0 for $j = i$); the summation on the right side of the above equation is thus effectively over all buses that are connected to bus i . The term $|V_i||V_j|b_{ij} \sin(\delta_i(t) - \delta_j(t))$ is the (real) power transmitted through the line from bus i to bus j . Thus the *maximum power* that can be transferred through the line, in the absence of other limits, is given by $|V_i||V_j|b_{ij}$. However, thermal dissipation constraints on the line (to limit thermal expansion of the line, and associated line sag) may limit the transferred power to values lower than this theoretical limit.

A key point to note is that the susceptance matrix $B = [b_{ij}] = \text{Im}(Y)$ provides a direct representation of the *network structure*; apart from its diagonal, it is essentially a *weighted incidence matrix* for the network graph. If the network topology evolves, as transmission lines are taken out in response to failures, for example, then the pattern of zero and nonzero entries in B (or Y) evolves in a corresponding manner.

In the limit when the motor constants are zero, we obtain a differential-algebraic equation (DAE) system, which can be interpreted as the singular perturbation limit of the model presented. In this case, we can write equations (1) in the form:

$$(1 - \gamma_i)P_{Li} + \gamma_i[M_i\ddot{\delta}_i(t) + D_i\dot{\delta}_i(t) - P_{mi}] = - \sum_j |V_i||V_j|b_{ij} \sin(\delta_i(t) - \delta_j(t)) \quad (2)$$

where γ_i is 1 if i is a generator bus and 0 if i is a load bus, P_{Li} is the power consumed at this bus by the load if i is a load bus, M_i is the normalized inertia of the generator if i is a generator bus. The DAE model is very useful from a conceptual viewpoint, since the loads are represented as *constant power* sinks. However, it also has some analytical difficulties, such as the nonexistence of solutions for some parameter values. For these practical reasons, we only use equations (1) in the present work.

2.2 A mechanical analogue

The above swing equations can also be interpreted as a model for a nonlinear spring-mass system, with a planar network of nonlinear springs interconnected at a set of nodes. The nonlinear characteristics of the springs are given by the sinusoidal terms in equation (2).

At node i , we have either a mass M_i oscillating normally to the planar network (rising above it by an amount δ_i), driven up by an external normal force P_{mi} , or we have simply an external normal force P_{Li} acting down normally to the plane. This mental picture is useful in developing some intuition for the motion of the power system, and also makes clear that what is known about the dynamics of spring-mass systems is directly applicable to the model. There is a very rich body of theoretical results about the analysis and control of nonlinear mechanical systems (in the precise sense defined in [7]).

2.3 Failure mechanisms

The network model employed in the simulations has multiple safety mechanisms, to represent real-world operation limits and security safeguards. As we will see, the interaction between the swing dynamics and the discrete topological changes induced by these overload protections gives rise to a succession of fault events, known as *cascading failures*.

The first protection mechanism is related to line overloading. If the sine of difference of the angles between two connected nodes $|\delta_i - \delta_j|$ exceeds a given threshold value, then the corresponding line is removed. In other words, we set the corresponding elements in the matrix Y (or B) to zero. As seen previously, the rationale behind this kind of action is basically a power flow limit protection.

The generators in our model are also susceptible to fault events. If the frequency deviation δ_i with respect to the synchronous reference exceeds some preassigned limit, the corresponding generator is tripped. This simulates the protective relay action present in most generators. In the model, the consequences of this are approximately captured by representing the tripped generator as a load bus with $P_{Li} = 0$. The topology of the network is not directly affected by generator failures, but only through line failures.

An important fact that needs to be taken into account is that of the loss of generating power in the event of a generator failure. In this case, the necessary conditions for existence of an equilibrium are no longer satisfied, since the system is consuming more power than the available supply. In our simulation, this is addressed by increasing in equal amounts the mechanical power input of the remaining generators. This is a somewhat crude representation of the actual mechanisms, which include governors that sense a drop in the local frequency, and correspondingly increase the mechanical power input to the generators. Future work will incorporate governor models for this adjustment. It should be noted, however, that in the cases where after the failure the network has many connected components, an explicit mechanism is needed to ensure that the power balance condition is satisfied for *every* connected subnetwork. Additional necessary conditions for the existence of an equilibrium can be obtained by exploiting the *max-flow/min-cut* theorem [2].

The complete model is a hybrid nonlinear system of coupled differential equations. The hybrid nature is given by the discrete topology changes induced by the relay actions. As we will see in the next section, this model has the potential for generating complex fault events, as a consequence of the complicated interaction between the dynamics and the overloading protection mechanisms.

3 Observed cascading failures

The specific network model used in the simulation is a 10×10 square grid, with four generators distributed in a regular pattern (see Figure 1). The resulting model is a nonlinear dynamical system, with 200 state variables, interconnected with the discrete relay components to form a hybrid system. The choice of network size was made mainly based on the prospect of reduced computation time for the full order model. Future implementations will include larger networks, as well as more complicated topologies.

The simulation is started from an equilibrium state near the origin, computed using a Gauss-Newton method. A line failure is simulated by setting two elements of the admittance matrix Y to zero ($Y_{ij} = Y_{ji} = 0$), since this implies no direct connection between the nodes i and j . The initial event is given by the opening of one branch of the network, chosen at random. This initial condition modifies the power flow equilibrium in the network. Depending on the initial event, and the network load, the system might converge to a new equilibrium, without violating the power flow constraints, or go through a succession of fault events, taking down additional lines and/or generators in the process. The simulation is run until the system reaches a new equilibrium, or all the generators trip.

4 Model reduction

The models we have used so far in this paper are hybrid, nonlinear interconnected systems, with 200 states used to describe the interaction of four generators with an idealized load distribution. At this scale, the model is readily amenable to simulation. However, we would like to be able to simulate significantly larger systems, both for the purpose of obtaining reliable statistics and for simulating a section of the physical power grid. Larger models quickly become unwieldy for simulation, and so we would like a method of constructing reduced-order models.

Critical difficulties in model reduction of this class of systems are presented by both their hybrid nature and their nonlinearity. Well-known methods, such as balanced-truncation, are not applicable other than to the linearization of the system, and do not take account of either of these crucial global system properties. In particular, the presence of the relays is responsible for the discrete chain of cascading failures, and we would like the reduced-order system to accurately predict these events.

An alternative model reduction procedure makes use of the Karhunen-Loève decomposition of the state space. This procedure was introduced for the analysis of turbulence by Lumley [6], and is described in detail in [3].

For high-dimensional systems whose state space is \mathbb{R}^n , a directly computable problem is to find the smallest dimensional subspace which contains observed points on the trajectories of the system; the Karhunen-Loève decomposition provides a method for doing this [5, 3]. A further refinement of this method is to relax the constraint that the subspace contain this set, but instead require that the subspace *approximately* contain it. Once such a subspace is found, a Galerkin projection can be applied to project the dynamics onto it, so that the high-dimensional system is approximated by a small number of nonlinear ordinary differential equations.

Suppose the state of the system is described by $x(t) \in \mathbb{R}^n$. In practical application, the method makes essential use of *empirical data*, taken either from experiments or from

numerical simulation, consisting of sampled measurements $\{x^{(1)}, \dots, x^{(N)}\}$ of $x(t)$. The next step is to perform a principal component analysis of this data, to find how well it may be approximated by projection onto a k -dimensional subspace of the original n -dimensional state space.

The Karhunen-Loève decomposition. The Karhunen-Loève decomposition provides a method for finding this best approximating subspace. The method is known in the literature by several names, including principal component analysis, factor analysis, and total-least-squares estimation. The method has been extensively analyzed in the literature, although the original concept goes back to Karl Pearson [9]. We give here a brief outline of the method; for details see [3].

We can characterize the subspace $S \subset \mathbb{R}^n$ by the projection operator Q mapping \mathbb{R}^n onto S . We would like to find Q to minimize

$$H(Q) = \sum_{i=1}^N \|x^{(i)} - Qx^{(i)}\|_2^2$$

the total squared perpendicular distance of the points from the k -plane. The following result is standard.

Theorem 1. (Total Least Squares) *Let R be the correlation matrix of the data, defined by*

$$R := \sum_{i=1}^N x^{(i)} x^{(i)*}$$

and let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the ordered eigenvalues of R . Then

$$\min_Q H(Q) = \sum_{i=n-k+1}^n \lambda_i$$

where the minimum is over all rank k projections Q .

In general, R may not have rank n , if the given data lies within a strict subspace of \mathbb{R}^n . Let $s = \text{rank } R$, and let $\phi_1, \phi_2, \dots, \phi_s$ be orthonormal eigenvectors of R , corresponding to the nonzero λ_i . Each $x^{(i)}$ can be written as

$$x^{(i)} = \sum_{j=1}^s a_{ij} \phi_j$$

where $a_{ij} = \langle x^{(i)}, \phi_j \rangle$, and $\langle \phi_i, \phi_j \rangle = \delta_{ij}$. The optimal k -dimensional subspace approximant is given by

$$\hat{x}^{(i)} = \sum_{j=1}^k a_{ij} \phi_j$$

Denote by P the $k \times n$ matrix whose rows are ϕ_1, \dots, ϕ_k , so that $PP^* = I$. The projected approximant to x is given by $P^*Px \in S$, and $y = Px$ is a representation in terms of the new

coordinates ϕ_i on S . This subspace approximant is then optimal, in the sense that the total energy (2-norm) in the subspace is given by

$$\sum_{i=1}^N \|Px^{(i)}\|_2^2 = \sum_{j=1}^k \lambda_j$$

and this is the maximum achievable by any k -plane.

Model reduction. Since in general the data collected may not have zero mean, approximation by an affine subspace would give better results. In statistical terms, we are more interested in the covariance of the data than in its correlation. Hence, we construct the correlation matrix as

$$R_{ij} = \sum_{p=1}^N (x_i^{(p)} - \bar{x}_i)(x_j^{(p)} - \bar{x}_j)$$

where $\bar{x} = \frac{1}{N} \sum_{i=1}^N x^{(i)}$ is the mean of the data. The eigenvalues of R now provide us with information as to how close an approximation of the data is provided by a k -dimensional subspace; the goal is to choose k such that the fraction of the total energy in the subspace

$$\sum_{i=1}^k \lambda_i / \sum_{i=1}^n \lambda_i$$

is close to one, yet k is sufficiently small. Clearly this will not always be possible, with models which are better approximated by low-dimensional systems having relatively few large eigenvalues.

We now subtract the mean \bar{x} from the dynamic equation, defining $e = x - \bar{x}$ to arrive at $\dot{e}(t) = f(e(t) + \bar{x})$. We would like to approximate $e(t)$ in this equation by $P^*Pe(t) \in S$.

Galerkin projection. The Galerkin projection has been used extensively to construct numerical solutions to partial differential equations. It can also be used to construct lower order *mathematical models* of a given dynamical system. The idea is to replace the given dynamics by an associated dynamics on a k -dimensional subspace $S \subset \mathbb{R}^n$ of the original state space.

Given the ordinary differential equation

$$\dot{x}(t) = f(x(t))$$

we would like to approximate this equation with a lower order description, of the form

$$\dot{y}(t) = g(y(t))$$

for $y(t) \in \mathbb{R}^k$. The Galerkin projection method is to write $x(t)$ in the form

$$x(t) = P^*y(t) + r(t)$$

for $y(t) \in \mathbb{R}^k$. Substituting this expression into the original ODE, we have

$$P^*\dot{y}(t) + \dot{r}(t) = f(P^*y(t) + r(t))$$

Here $\dot{r}(t)$ is known as the *residual*. The method is then to force the residual to be orthogonal to the k -dimensional subspace onto which the dynamics are projected, so that $Pr(t) = 0$. This results in

$$\dot{y}(t) = Pf(P^*y(t) + r(t))$$

So far this equation is exact; the approximation is then to assume that we have chosen the projection P so that $r(t)$ is minimized in some sense, so that the reduced-order system is given by

$$\dot{y}(t) = Pf(P^*y(t))$$

The initial conditions of this reduced-order system are given by $y(0) = Px(0)$.

Up until now we have neglected the mean in this projection; applying this method to the mean-adjusted equation $\dot{e}(t) = f(e(t) + \bar{x})$ leads to the reduced order approximation

$$\dot{y}(t) = Pf(P^*y(t) + \bar{x})$$

The Karhunen-Loève method therefore projects the dynamics onto the subspace containing most of the energy of the system. If we keep all of the eigenvectors corresponding to non-zero eigenvalues, then this subspace will contain the attractor (a set to which most of the trajectories converge to), and hence will have dimension greater than or equal to the Hausdorff dimension of the attractor. In general, we would expect that the more eigenvectors we keep, the better approximation we will obtain.

So far, we have discussed application of the Karhunen-Loève method to systems which can be described by autonomous differential equations. For systems which have a hybrid component, essentially the same approach applies, with the differential equation replaced by a mixed discrete/continuous system. Numerically, such a system is integrated in discrete timesteps, which can be represented as

$$x(t_{k+1}) = h(x(t_k)).$$

The Galerkin projection then leads to the discrete iteration

$$y(t_{k+1}) = Ph(P^*y(t_k) + \bar{x})$$

which can be directly implemented for nonlinear hybrid systems, such as the power system considered in this paper.

5 Initial test of the model reduction approach

In this section we describe an application of the above procedure to the 200 state swing-model. We consider the four-generator example described above, and initiate a failure close to the generator in the top-left corner.

In order to try to capture this behavior with a reduced order model, a reasonable strategy might be to model reduce only *part* of the system. In this case, the initial failure spreads in the top left corner of the network, and so we might expect that the dynamics of the bottom half of the network do not play a significant role.




Figure 2: First mode of the Karhunen-Loève decomposition

The Karhunen-Loève method can be directly applied to only part of the system; the states of the system are partitioned into (x_1, x_2) , where x_1 are the states to be kept and x_2 are to be replaced by a lower-order approximation. A projection P is constructed from measurements of $x_2(t)$ collected from simulation. This is extended to a projection \hat{P} on the whole state space, defined by $\hat{P}^T = [I \ P^T]$, which is then used to perform the Galerkin projection.

Data was collected from the simulation of the full-order system, and the correlation and projection operators were constructed from the 100 states corresponding to the bottom half of the network. Figure 4 shows the eigenvalues of the correlation matrix, scaled so that they sum to one. As can be seen from the plot, most of the energy is concentrated in the first few modes. Figure 2 shows the mode shape corresponding to the first singular value.

The energy distribution of the first modes suggests that a low order model might explain adequately the dynamics of this part of the system. A ten state model was chosen since these ten singular values sum to 0.998, and hence account for more than 99% of the ‘energy’ in the system.

After realizing the corresponding reduction, the reduced order model is interconnected with the unreduced remaining subnetwork, and a new simulation is performed. See Figure 3 for a comparison of the response of the two different models to the same failure.

Visually, the agreement between the simulations of the two models is extremely good, despite the fact that the reduced order part has only 10 states compared with the original 100 states. More importantly, the sequence of failures obtained is exactly the same as that of the full-order system. The sequence of failures obtained is detailed in Table 1. The initial failure occurs at time $t = 0.25$, and it causes the tripping of all the generators. Both the fault progression and the final equilibrium of the system are exactly the same for both models

This is a promising result; the Karhunen-Loève decomposition is motivated by the desire to reproduce the *large scale behavior* of the system, and in this case we have obtained reliable predictions of the sequence of failures. Naturally, it would *not* be reasonable to expect the methodology to work in every case, and for arbitrarily long intervals. One reason for this is that the system might present chaotic behavior, and therefore arbitrarily small perturbations in the dynamics or the initial conditions will cause an exponentially growing divergence in the state.

This initial test of this methodology indicates that a *multi-resolutional* approach might be successful for larger scale power networks, using models which are more detailed at parts of




Figure 3: Original and reduced order model simulations.

Fault type	Time (Full)	Time (Reduced)
Initial line	0.25	0.25
Line	1.764	1.773
Line and Generator 1	1.845	1.848
Generators 3 & 4	2.570	2.574
Generator 2	2.582	2.594

Table 1: Failure times for the full and reduced order models.

the network which are dynamically more active. It should also allow for successive refinement, so that the system can be modeled in increasing detail in those areas which are susceptible to failures. This type of modeling process has been described as *hierarchical modeling*; a tree-structure of models is used to describe the behavior of the system, and high resolution (in this case, a large number of states) is only used where needed.

6 Conclusions

The complexity of modern-day power networks makes necessary the careful evaluation and analysis of global events, such as cascading failures. Additional demands such as decentralization and the need for obtaining solutions in reasonable computation times make necessary a reevaluation of the suitability of current techniques in dealing with these new challenges.

A possible solution is the use of reduced order models for analysis. The methodology presented in this paper, based on the Karhunen-Loève decomposition, deals directly with the problem of finding a subspace where “most” of the dynamics occur, and uses this information in the construction of a reduced order model. The results shown in Section 3 demonstrate




Figure 4: Normalized singular values associated with the reduction procedure.

that the sequence of fault events can be evaluated and predicted (at least, in a reasonable horizon length) without necessarily simulating the full-order model.

The initial investigations presented here are very promising, and give a hint of the potential value of the methodology. Future work will explore important related issues, such as a possible tradeoff between approximation error and prediction horizon.

References

- [1] A. Bergen. *Power Systems Analysis*. Prentice-Hall, 1986.
- [2] H. Frank and I. Frisch. *Communication, Transmission, and Transportation Networks*. Addison-Wesley, 1971.
- [3] P. Holmes, J. L. Lumley, and G. Berkooz. *Turbulence, coherent structures, dynamical systems, and symmetry*. Cambridge University Press, 1996.
- [4] P. Kundur. *Power System Stability and Control*. McGraw-Hill, 1994. EPRI Power System Engineering Series.
- [5] S. Lall, J. E. Marsden, and S. Glavaški. Empirical model reduction of controlled nonlinear systems. Caltech CDS technical report 98-008, 1998. Available at http://www.cds.caltech.edu/~sanj/home/publications/ifac99_empmr.
- [6] J. L. Lumley. The structure of inhomogeneous turbulence. In A. M. Yaglom and V. I. Tatarski, editors, *Atmospheric turbulence and wave propagation*, pages 166–178. Nauka, Moscow, 1967.
- [7] J. E. Marsden. *Lectures on Mechanics*. Cambridge University Press, 1992.
- [8] M. A. Pai. *Power System Stability*. North-Holland Publishing Company, 1981.
- [9] K. Pearson. On lines and planes of closest fit to points in space. *Philosophical magazine*, 2:609–629, 1901.