

A Phase Transition Model for Cascading Network Failure

By Christopher L. DeMarco

As the scale of engineered systems such as electric power grids, communication networks, and the Internet expands and as society's dependence upon reliable operation of these networks increases, it is vital that system engineers seek a better understanding of how small-scale failures of individual elements may propagate to produce global, system-wide failures. Within this general class of problems, importance sampling methods have improved the efficiency of search techniques for identifying rare failures in probabilistic models for communication networks [1], [2]; similar progress has been made in characterizing the role of graph interconnection structure in propagation of failures in electric power networks [3], [4]. The work here is motivated largely by the electric power systems application. Efforts to date on

these problems for power networks have focused primarily on Markov chain-like representations, in which the state of the system is represented as discrete and the engineering constraints impacting failure transitions are those of the steady state. Our goal here will be somewhat different in flavor. In a deterministic model, we seek to capture the role of transient dynamic response following a specified initiating disturbance and examine subsequent ("cascading") element failures that are induced when operating thresholds for individual elements are exceeded along the state trajectory. In the formulation to be presented here, the potentially probabilistic nature of initiating events is not treated.

To allow tractable analysis of continuous dynamics in state-space models of high dimension, we will restrict attention to a specialized class of system representation, one in which energy storage associated with individual elements is

DeMarco (demarco@engr.wisc.edu) is with the Department of Electrical and Computer Engineering, University of Wisconsin–Madison, 1415 Engineering Drive, Madison, WI 53706, U.S.A.

closely tied to a structure of ordinary differential equations that describes the system dynamics. Hamiltonian systems will be members of this class, as well as Euler Lagrange systems with certain structure to their dissipation terms. This will include a wide class of $R-L-C$ models in electric circuits (including nonlinear inductive and capacitive elements) and classical models for electric power systems. Key to our formulation will be the fact that the threshold failure mechanisms modeled can be associated with energy-storing elements and that the threshold itself may be expressed in terms of energy stored in the branch. We believe that in many systems, this association of failure with excessive energy stored in an element is well supported by the underlying physics governing failure, even when details of the transition to the failed state are more crudely approximated in the macroscopic system model.

In the context of the control systems literature, systems displaying threshold failure mechanisms naturally suggest representation through the techniques of hybrid systems, in which continuous dynamics described by ordinary differential equations interact with discrete logic terms. In the scenario of interest, failure of individual elements would be governed by threshold logic, and the (assumed irreversible) failure of an element would induce a discontinuous change in the structure of the vector field, with a corresponding change in the equilibria, and a potentially significant transient in the state trajectory. From an analytic perspective, rigorous treatment of systems of differential equations with discontinuous right-hand sides raises a host of interesting and challenging issues, motivated by such practical problems as the dynamics and control of robotic mechanisms making contact with rigid surfaces [5]. Treatment of the analytic issues in differential equations with discontinuous right-hand sides is beyond the scope of this article. We will instead seek a smooth model from the outset, with the view that the element failures to be represented display behavior that can occur on a very fast time scale, but are fundamentally continuous (e.g., consider the transition to open circuit due to “burn-out” of a circuit element experiencing overcurrent). The key objective here will be to develop a model in which the tools of Lyapunov stability analysis can treat the interaction of multiple-element failure events. To judge whether the system maintains acceptable performance, one would like to determine whether or not the state returns to some acceptable stable equilibrium following the initiating event, and what elements have failures induced along the state trajectory leading to this equilibrium (if it exists). Hence a large piece of the work to be presented here can be interpreted as a stability analysis for a particular class of smoothed hybrid system.

The recent perspective provided in [6] gives an overview of a variety of approaches to stability analysis within hybrid systems. As noted there, a considerable literature exists employing Lyapunov-based techniques for determining stabil-

ity in such systems, with many approaches resting upon a partitioning of the state space, and the “stitching together” of individual Lyapunov functions appropriate to dynamics within each partition. Such approaches are often employed for the design of a logic-based switching controller, and one seeks conditions to guarantee acceptable performance. This contrasts with our goal of characterizing possible paths to system failure. The difference is important, because many of the techniques reviewed in [6] presume a single, fixed equilibrium point of interest about which stability is to be established. Our analysis will be strongly focused on multiple points, and we will seek to estimate the attractive regions associated with locally stable equilibria and to characterize the transitions between them, the “phase transitions” alluded to in the title. Despite these differences, the approach to be presented here does have strong connections to the ideas underlying the use of partitioned Lyapunov functions in hybrid systems. While we will directly construct a single, global Lyapunov function, our modeling of element failure introduces a natural partitioning of the state space in which the Lyapunov function varies its behavior across failure boundaries. This will be clearly illustrated in the circuit example to be presented below.

To facilitate the proposed Lyapunov analysis in a system of fixed structure, we will smooth the discontinuity of an ele-

A large piece of the work to be presented here can be interpreted as a stability analysis for a particular class of smoothed hybrid system.

ment failure event. The ideal binary indicator variable describing $\{element\ in\ service\}$ versus $\{element\ out\ of\ service\}$ is treated as a continuously varying state in a smooth differential equation. The construction of the state equation for the indicator variable will be such that it exhibits strongly bistable behavior, maintaining its value in a small neighborhood of 1 so long as the element is operational, and rapidly forcing its value to a small neighborhood of 0 during the failure event. It is important to emphasize that this approach is adopted to facilitate Lyapunov-based estimates of stability and of attractors associated with multiple equilibria that arise in this model. It is not advocated as the most effective means for numerical simulation. While the example to follow will employ simulated trajectories to verify the qualitative behavior of the model, use of smoothed representations of element failure is typically not efficient in simulations. As illustrated in our circuit example, the smoothed model is constructed so that the failure transition, when triggered, occurs much more rapidly than the dynamic evolution of states between failures. This inherently introduces widely separated time scales, creating a stiff set of differential equations.

Given the fairly specialized failure model this work will develop, a key goal of the exposition to follow will be to convince the reader that the class of applicable systems is wide enough to be of importance. Our motivation stems from a class of systems that contains classical Hamiltonian dynamics as a special case and applies to a variety of nonlinear cir-

While we will directly construct a single, global Lyapunov function, our modeling of element failure introduces a natural partitioning of the state space.

cuit models and mechanical system models for which energy storage plays an intimate role in determining the structure of the state equations. The recent review in [7] argued strongly for the importance of energy storage and energy-shaping techniques in control analysis and design; the work to be presented here will seek to further reinforce that view.

Hamiltonian-Like System Model Structure: Failure-Free Case

As noted above, the class of systems to which our element failure approach can be applied includes Hamiltonian systems as a special case; it is useful to briefly review the Hamiltonian structure as motivation. Recall that the state space of a Hamiltonian system has a special symplectic structure; in the standard coordinate system of classical mechanics, one has an even number of system states, which partition into a vector of positions, \mathbf{z} , and an associated vector of momenta, \mathbf{p} . The Hamiltonian function is then a differentiable real-valued function of these states, $H(\mathbf{z}, \mathbf{p})$, with the state equations taking the form

$$\begin{bmatrix} \dot{\mathbf{z}} \\ \dot{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix} \nabla_{\mathbf{z}, \mathbf{p}} H(\mathbf{z}, \mathbf{p}).$$

The nondissipative property of the system is then reflected in the fact that the derivative of H along trajectories is guaranteed to be zero; i.e.,

$$\frac{dH(\mathbf{z}(t), \mathbf{p}(t))}{dt} = [\nabla H(\mathbf{z}, \mathbf{p})]^T \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix} \nabla H(\mathbf{z}, \mathbf{p}) = 0.$$

To generalize the class of systems of interest, consider a state space of arbitrary dimension, $\mathbf{x} \in \mathbf{R}^n$, with an associated differentiable function, $\vartheta(\mathbf{x})$, and a full-rank, negative semidefinite matrix $\mathbf{A} \in \mathbf{R}^{n \times n}$. The systems to be treated here, prior to the addition of the element failure model, will be required to take the form

$$\dot{\mathbf{x}} = \mathbf{A} \nabla \vartheta(\mathbf{x}). \quad (1)$$

This structure appears in certain Euler-Lagrange systems, provided the Rayleigh dissipation can be captured in a term closely related to the energy storage. The reader is encouraged to consult [8, chap. 4], in which a class of generalized Hamiltonian systems also closely related to (1) is reviewed; there \mathbf{A} is allowed to be a smooth function of state, with a requirement of skew symmetry to yield losslessness. A number of physical examples developed in [8] argue for the usefulness of that class; specifically, it is observed that LC circuits possess the generalized Hamiltonian structure, and that “the inclusion of resistive elements will affect the Hamiltonian form of the equations, but not the passivity of the system.” For the circuit example to be presented below, the addition of linear resistive terms simply requires relaxing the skew symmetric requirement on \mathbf{A} , by addition to \mathbf{A} of a symmetric (often diagonal) negative semidefinite matrix that captures the dissipative effects of the resistors.

Note that the assumption of \mathbf{A} full rank ensures that equilibria for (1) will occur only at those points \mathbf{x}^e for which $\nabla \vartheta(\mathbf{x}^e) = 0$; \mathbf{A} negative semidefinite ensures that $\vartheta(\mathbf{x}(t))$ is nonincreasing along any $\mathbf{x}(t)$ that is a solution trajectory of (1). Given $\vartheta(\mathbf{x}(t))$ nonincreasing, sets of the form $\{\mathbf{x} | \vartheta(\mathbf{x}) < \text{constant}\}$ are invariant. With certain additional assumptions (to ensure that $d\vartheta/dt$ will not remain identically zero along nontrivial trajectories), $\vartheta(\mathbf{x})$ serves as a natural, global Lyapunov function for (1), and $\nabla^2 \vartheta(\mathbf{x}^e) > 0$ provides a necessary and sufficient condition for an equilibrium \mathbf{x}^e to be asymptotically stable. In this framework, a geometric picture emerges suggestive of that describing phase transitions in materials; that is, stable equilibria are local minima of “potential wells” with respect to $\vartheta(\mathbf{x})$, and the ease of transition between two equilibria is governed by the height of the potential barrier that must be overcome in the transition. The lowest energy path of transition between two stable equilibria will typically pass through a “saddle exit” point, which is itself a critical point at which $\nabla \vartheta(\mathbf{x}) = 0$. In the framework we will describe, each element failure creates the possibility of a new stable equilibrium. For very severe failures, an equilibrium may fail to exist. Our goal will be to exploit the geometric insights regarding the behavior of $\vartheta(\mathbf{x})$ to supplement simulation techniques in evaluating the vulnerability of a system to transitions in which large numbers of elements fail or to a case in which the state never reaches an acceptable stable equilibrium.

To allow the element failure representation to be added, we will place additional assumptions on the form of $\vartheta(\mathbf{x})$. In particular, for each element subject to possible failure, $\vartheta(\mathbf{x})$ must contain a separable term that identifies energy storage in that element. As noted in the introduction, we require that the element failure threshold be expressible as a limiting

value on stored energy in the element. In many cases there will be an equivalence between a more intuitive form of the failure threshold and this threshold expressed in terms of stored energy. To illustrate, consider an interconnected spring-mass-damper model, as might arise in a finite element representation of a flexible mechanical element. In such a framework, individual spring elements represent the connective bonds. The reader will recall that the behavior of a linear spring, with spring constant K , is simply represented as $\{\text{restorative force}, f\} = K \times \{\text{displacement}, x\}$, and its stored energy as $(1/2)Kx^2$. To a good first approximation, failure of such a spring would be triggered by the spring force magnitude exceeding some specified threshold; upon failure, the effective spring constant K goes to zero. These simple relations then provide the motivation for our failure model. One first observes that in the operative state (K at its “normal” operational value), there is a one-to-one relationship between spring force magnitude and stored energy in the spring; hence, the failure threshold may be expressed in terms of either quantity. Once the failure threshold is exceeded, the system dynamics must drive the spring constant K to zero.

Extending this perspective to circuit applications, the analogous element failure would be triggered by the magnitude of a branch current i exceeding a specified threshold. In a circuit context, a common representation of such a failure would be a “fuse” element, modeled as a nonlinear resistor with a transition from a low resistance value to a high (possibly infinite) resistance value when the triggering threshold is exceeded. In contrast, our approach will require that branches subject to failure have series inductance, L . The inductor, analogous to the spring in the mechanical context, becomes the focal point for representing failure in the branch, and the branch flux, λ , becomes the state variable analogous to displacement in the mechanical context. Given an inductor current magnitude as a threshold for failure, there exists an equivalent failure threshold expressed as a limit of stored inductor energy $(1/2L)\lambda^2$. Once the threshold is triggered, the appropriate response to represent an infinite impedance (failed) branch will be to drive $1/L$ to zero. To capture this behavior in a smooth state space model of the form (1), our approach requires that the “original” $\vartheta(\mathbf{x})$ for the circuit (before modification for the failure representation) contains the separable term $(1/2L)\lambda^2$. The following section will demonstrate how the state equations to represent the threshold-driven failure can be specially structured to create a new model that maintains the form of (1), with only a very simple modification to $\vartheta(\mathbf{x})$. This observation is at the heart of our approach.

A Bistable Branch Threshold Failure Model

To represent failure in a branch, the state space of our model will be augmented to associate an approximately binary indicator with each branch that is subject to the possi-

bility of failure. A dynamic structure is imposed that yields two (potentially) stable equilibria for each indicator variable. As will be illustrated below, by construction these equilibria will be placed in a neighborhood of 1, corresponding to the operational state of the branch, and in a neighborhood of 0, indicating a failed state. While representations of branch recovery are possible within this general framework, our exposition here will focus on a “one-way” model; (i.e., a branch that undergoes failure will have no possibility of returning to the operational state). We let γ_k denote the failure indicator state for branch k .

As motivated in the previous section, the mechanism of failure to be represented removes a branch from service when a specified threshold of energy storage in the branch is exceeded. In our illustrative R-L-C circuit model to follow, the energy function $\vartheta(\mathbf{x})$ contains terms associated with magnetic energy storage in inductive branches, of the form

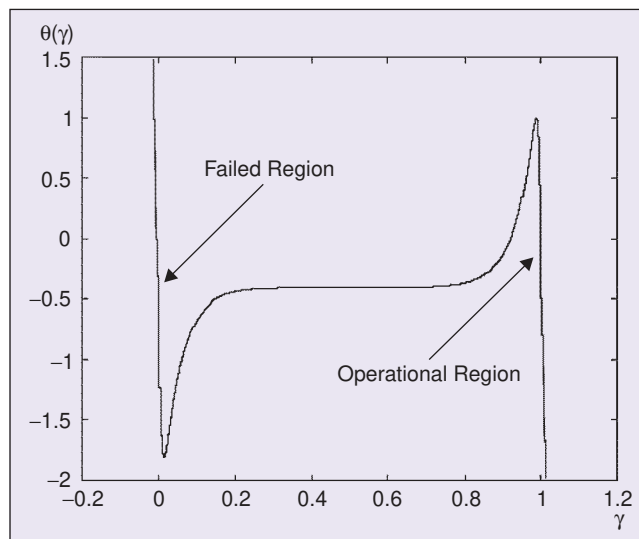


Figure 1. Plot of $\theta(\gamma)$ versus γ .




Figure 2. Plot of $\Theta(\gamma)$ versus γ .

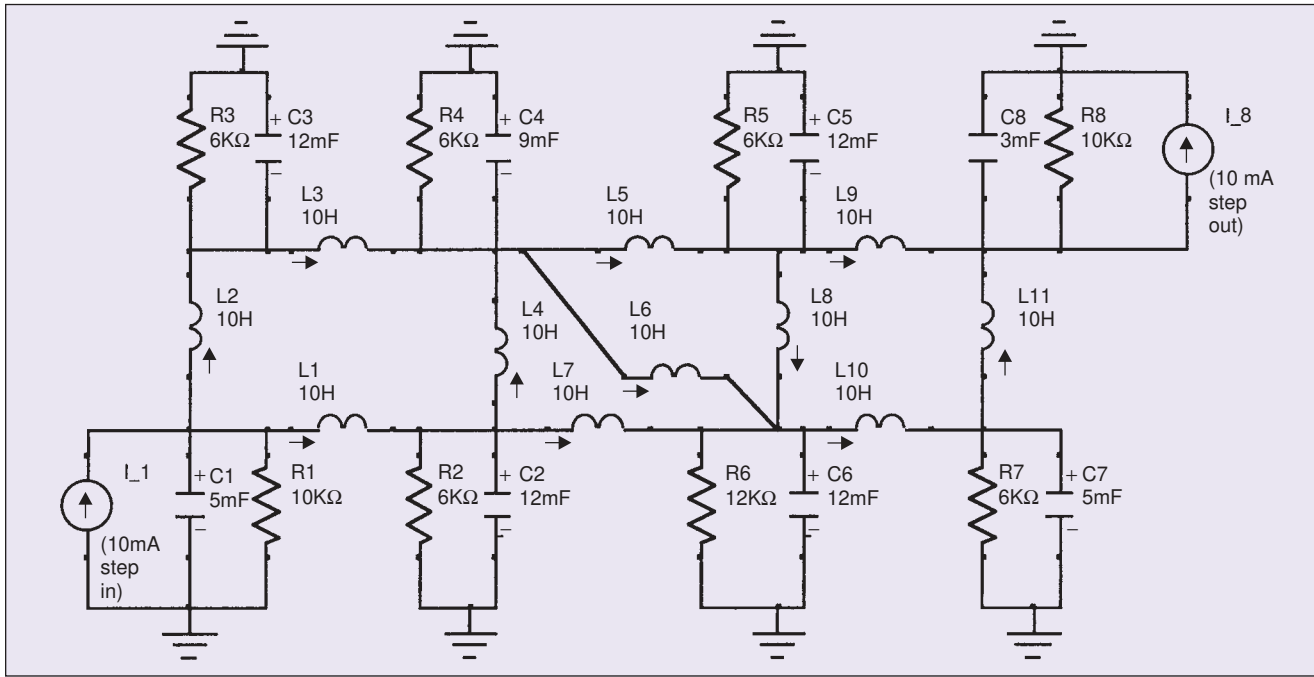


Figure 3. Eight-node R-L-C circuit example.

$(1/2L)\lambda^2$. We seek to modify $\vartheta(\mathbf{x})$ to introduce dependence on the branch failure states, γ . In particular, the new $\vartheta(\mathbf{x}, \gamma)$ should have the property that $\partial\vartheta/\partial\gamma$ recovers a term suitable for “driving” the time derivative of γ .

A construct to achieve this behavior within the dynamics of γ proves quite simple. As a first step, the term describing energy stored in branch k is simply pre-multiplied by γ_k ; for example, in the circuit model, $(1/2L)\lambda_k^2$ in $\vartheta(\mathbf{x})$ is replaced by $(1/2L)\gamma_k\lambda_k^2$. Moreover, for each branch, we add another term to $\vartheta(\mathbf{x})$ that is dependent only on the γ_k for that branch; this term characterizes the threshold that must be overcome to induce a failure. Denote the term to be added to $\vartheta(\mathbf{x})$ as $\Theta(\gamma_k)$; its associated derivative, with sign reversed, is denoted in lower case, i.e., $\theta(\gamma_k) := -(\partial\Theta/\partial\gamma_k)$. A representative example of a $\theta(\gamma)$ reflecting a unity failure threshold is displayed below (the subscript k on γ is dropped to minimize notational clutter):

$$\theta(\gamma) := 2 \times [-\exp(-20\gamma) + \exp(-200\gamma) + \exp(20(\gamma-1)) - \exp(200(\gamma-1)) - 0.2].$$

A plot of this unity threshold $\theta(\gamma)$ is illustrated in Fig. 1.

An understanding of the bistable branch failure model may be gained through the properties of $\theta(\gamma)$. Suppose the derivative of γ with respect to time is composed of $\theta(\gamma)$ minus the energy stored in the branch; i.e., suppose $(d\gamma/dt) = \theta(\gamma) - (1/2L)\gamma\lambda^2$. Assume that λ^2 were to begin from an initial condition near 0, and γ from an initial condition near 1, so that the initial γ lies along the steeply sloped region in Fig. 1 labeled as the “operational region.” Suppose further that dynamics of the circuit are such that $(1/2L)\lambda^2$ increases with time. In this scenario, γ

would maintain a stable, quasi-equilibrium value along the operational region of $\theta(\gamma)$ (thereby keeping $\gamma \approx 1$), so long as $(1/2L)\gamma\lambda^2$ had a magnitude less than 1. If $(1/2L)\gamma\lambda^2$ were to exceed the threshold value of 1, $(d\gamma/dt)$ would become negative, pushing γ to values less than 1. The nature of $\theta(\gamma)$, and the fact that $(1/2L)\lambda^2$ is nonnegative by construction, ensure that when γ drops below 1 by any appreciable margin, $d\gamma/dt$ becomes very strongly negative and γ further decreases until it is “captured” in a new quasi-equilibrium in the vicinity of $\gamma = 0$, along the portion of the plot denoted as the “failed region” in Fig. 1.

By appropriate scaling of $\theta(\gamma)$, the threshold value that initiates failure is parametrized in the model. Similarly, by suitable adjustment of the coefficients of the exponential terms, the slopes in the vicinity of the $\gamma = 0$ and $\gamma = 1$ critical points can be modified, forcing equilibrium values to lie as close to 0 or 1 as desired. The seemingly arbitrary choice in the construction of $\theta(\gamma)$ is the negative offset selected for its “flat” region; in the graph, one observes a negative offset of approximately 0.4. The impact of this choice of offset can be easily interpreted. As noted above, to obtain the structure of state equations that allows the potential function $\vartheta(\mathbf{x}, \gamma)$ to serve as a global Lyapunov function in our augmented model, a sign-reversed integral of $\theta(\gamma)$ must be added to the potential. Hence, our interest focuses on the integrated function below (note: a representative lower limit of integration of 1 is shown below; more precisely, this lower limit would be the exact equilibrium value of γ in the vicinity of 1):

$$\Theta(\gamma) := \int_1^\gamma -\theta(\eta) d\eta.$$

For the representative $\theta(\gamma)$ illustrated in Fig. 1, the resulting $\Theta(\gamma)$ is shown in Fig. 2. The impact of the -0.4 offset in constructing $\theta(\gamma)$ is clear; a transition of a branch from operational state to failed state contributes a net drop in energy of approximately -0.4 . Moreover, the graph of $\Theta(\gamma)$ illustrates the qualitative properties we have obtained: local minima about 1 and 0, with a potential barrier blocking the transition from 1 to 0. The key parameter is the height of the potential that must be overcome to escape the local minimum about the “operational” equilibrium point; this is set by the failure threshold and is *not* affected by the arbitrary choice of the -0.4 term in the definition of $\theta(\gamma)$.

Circuit Example with Branch Failure Model

Consider a failure-free linear R-L-C network in which each node is connected to ground through a parallel R-C combination, and nodes are interconnected through inductive branches to form a simply connected graph. Each node may optionally have a piece-wise constant current source attached. Standard textbook treatments [9] show that this circuit topology yields a direct analogy to a class of spring-mass-damper systems. We will consider a scenario in which the network is subject to an initial disturbance in the form of one or more simultaneous step changes at current inputs. To represent branch failure, an inductive branch is removed from service when a specified current threshold is exceeded (or equivalently, a specified $(1/2L)\lambda^2$ threshold). In this application, we wish to efficiently determine if a given disturbance causes one or more branch failures, and whether these failures are triggered in a relatively small number of branches, allowing an acceptable new steady state to emerge, or whether they continue to propagate, ultimately separating the network into disconnected components.

To illustrate these ideas in a circuit of modest complexity, an eight-node, 11-branch R-L-C network example is shown in Fig. 3. With the goal of focusing attention on the nonlinear behavior introduced by our branch failure model, this example is chosen to be purely linear in other respects; i.e., for a fixed set of inductors in service, it is governed by linear state equations. However, we wish to stress that the fact that the state equations display the special structure of (1) is *not* dependent upon linearity of energy storage elements. The model developed below easily generalizes to the case of nonlinear charge-controlled capacitors, and nonlinear flux-controlled inductors, while keeping the important special structure of (1).

Note that the inductive branches in Fig. 3 are *purely* inductive, with no series resistive term. The reader familiar with circuit modeling will observe that this reduces the dimension of the state space, as each inductor no longer has associated a unique state variable (e.g., current or flux linkage). Rather, as we will show, it becomes natural to associate flux quantities with each node, and the differences of these nodal fluxes become elements of our state vector.

This “branch-resistance-free” property is not necessary for the construction; it is selected for the convenience of reduced state dimension and because of its resemblance to a class of models that appears in stability analysis of electric power systems.

The element values selected in Fig. 3 produce a lightly damped linear system with a large number of closely grouped oscillatory modes having natural frequencies slightly below 1 Hz. These characteristics present a challenging scenario in which to predict element failure, as disturbances damp out slowly and a current may first reach its maximum magnitude well after the initiating event. To illustrate a representative response, Fig. 4 displays the branch currents in L_9 and L_{10} over a 4-s interval following simultaneous step inputs of 10 mA at node 1 and -10 mA at node 8, with all inductors and capacitors having zero initial conditions. Given our choice of input, which balances a 10-mA current flow in at node 1 with equal flow out at node 8, key properties of the associated steady state can be predicted by inspection; all the node voltages will go to zero in steady state (i.e., the capacitors must fully discharge through their parallel resistances), and any cutset of branches separating node 1 from node 8 will carry a net flow of 10 mA, with currents dividing equally among branches in the cut. Hence, relevant to Fig. 4, currents in L_9 and L_{10} ultimately approach steady-state values of 5 mA.

To obtain further insights into the dynamic behavior of this circuit, and to confirm that it possesses the structure of (1) for fixed current inputs, we begin from basic circuit quantities:

$$\begin{aligned} \mathbf{q} &:= \text{nodal charges on capacitors; } \mathbf{q} \in \mathbf{R}^8 \\ \mathbf{v} &:= \text{node voltages relative to ground; } \mathbf{v} \in \mathbf{R}^8 \\ \mathbf{i}_L &:= \text{inductor branch currents; } \mathbf{i}_L \in \mathbf{R}^{11}. \end{aligned}$$

It will prove useful to reduce the vector of node voltages, as we will be interested only in differences of node voltage quantities. To this end, we define




Figure 4. Representative inductor currents following step input.




Figure 5. Branch current L_5 with γ_5 , versus time (s).




Figure 6. Branch current L_{10} with γ_{10} , versus time (s).

$$\mathbf{E}_r = \begin{bmatrix} -1 & -1 & -1 & -1 & -1 & -1 & -1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

and let

$$\mathbf{v}_r := \mathbf{E}_r^T \mathbf{v}.$$

Note that the elements of \mathbf{v}_r are simply the seven node voltage differences $[(v_2 - v_1); (v_3 - v_1); \dots; (v_8 - v_1)]$. We will asso-

ciate a vector of nodal flux differences with these \mathbf{v}_r , defining $\dot{\phi}$ via

$$\dot{\phi} = \mathbf{v}_r.$$

Important to the development of state equations will be a description of the interconnection of the inductive branches to nodes, through a reduced node-to-branch incidence matrix [10], denoted as \mathbf{H}_r . Node 1 is chosen as the reference node for this incidence matrix (note that the inductive branches make no connection to the ground node, so ground does not play a role in this reduced incidence matrix). The resulting \mathbf{H}_r for this circuit is given by

$$\mathbf{H}_r = \begin{bmatrix} -1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 & -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & -1 \end{bmatrix}.$$

This reduced incidence matrix is important to our state space formulation because it allows recovery of the vector of branch fluxes associated with inductors, λ , from nodal fluxes ϕ . In particular, one has:

$$\lambda = \mathbf{H}_r^T \phi.$$

Moreover, it is straightforward to show that through \mathbf{H}_r and \mathbf{E}_r one may obtain:

$$\mathbf{E}_r \mathbf{H}_r \mathbf{i}_L = \text{sum of inductor currents leaving each node.}$$

Relevant circuit parameter matrices are defined as:

\mathbf{C} := diagonal matrix of capacitance values,

indexed by node, $\mathbf{C} \in \mathbf{R}^{8 \times 8}$

\mathbf{G} := diagonal matrix of conductances to ground,

indexed by node, $\mathbf{G} \in \mathbf{R}^{8 \times 8}$

\mathbf{L} := diagonal matrix of inductance values,

indexed by branch, $\mathbf{L} \in \mathbf{R}^{11 \times 11}$

so that

$$\mathbf{v} = \mathbf{C}^{-1} \mathbf{q} \Rightarrow \mathbf{v}_r = \mathbf{E}_r^T \mathbf{C}^{-1} \mathbf{q}; \mathbf{i}_L = \mathbf{L}^{-1} \lambda = \mathbf{L}^{-1} \mathbf{H}_r^T \phi.$$

The model state variables are then ϕ and \mathbf{q} , with the expressions above defining the recovery of more standard quantities such as capacitor voltages and inductor currents from these states. The actual state equation for ϕ follows directly as

$$\dot{\phi} = \mathbf{E}_r^T \mathbf{C}^{-1} \mathbf{q}.$$

The state equation for \mathbf{q} is obtained by recognizing that $\dot{\mathbf{q}}$ represents current flow into each of the node capacitors, which must be balanced by $-\{\text{current into each node resistance}\}$, $-\{\text{sum of current flow out of each node into inductors}\}$, and $+\{\text{current source inputs at nodes}\}$, yielding

$$\dot{\mathbf{q}} = -\mathbf{G}\mathbf{C}^{-1}\mathbf{q} - \mathbf{E}_r\mathbf{H}_r\mathbf{L}^{-1}\mathbf{H}_r^T\phi + \mathbf{i}_{\text{in}}.$$

Rearranging terms, one obtains

$$\begin{bmatrix} \dot{\phi} \\ \dot{\mathbf{q}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{E}_r^T \\ -\mathbf{E}_r & -\mathbf{G} \end{bmatrix} \begin{bmatrix} \mathbf{H}_r\mathbf{L}^{-1}\mathbf{H}_r^T\phi \\ \mathbf{C}^{-1}\mathbf{q} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{i}_{\text{in}} \end{bmatrix}.$$

Letting

$$\mathbf{A} := \begin{bmatrix} \mathbf{0} & \mathbf{E}_r^T \\ -\mathbf{E}_r & -\mathbf{G} \end{bmatrix}$$

and

$$\vartheta(\phi, \mathbf{q}) := \frac{1}{2}\mathbf{q}^T\mathbf{C}^{-1}\mathbf{q} + \frac{1}{2}\phi^T\mathbf{H}_r\mathbf{L}^{-1}\mathbf{H}_r^T\phi + [\phi^T \mathbf{q}^T]\mathbf{A}^{-1}\begin{bmatrix} \mathbf{0} \\ \mathbf{i}_{\text{in}} \end{bmatrix},$$

direct calculation confirms that these state equations display the special structure of (1). Moreover, given the diagonal form of \mathbf{L} , it is clear that the quadratic form for inductive energy storage allows one to separate an individual term equivalent to $\lambda^2/2L$ for each inductor. We also reaffirm the earlier claim that this structure generalizes easily to the case of charge-controlled nonlinear capacitors and flux-controlled nonlinear inductors. To treat the case of nonlinear capacitors and inductors, the quadratic forms in \mathbf{q} and ϕ appearing in $\vartheta(\phi, \theta)$ are simply replaced by terms obtained by integrating capacitor voltage against charge and integrating inductor current against flux. Identical reasoning applies if one wishes to treat the analogous mechanical spring-mass-damper system, in the case of nonlinear spring (elastic) forces connecting masses.

Addition of Branch Failure to Circuit Model

As described previously, failure of a branch k will be represented by the inductance L_k becoming infinite, or equivalently, by L_k^{-1} going to zero. L_k infinite yields a circuit element whose impedance is unbounded, resulting in no flow of current along the branch k . This representation of branch failure is perhaps more intuitive if first described in the mechanical analogy. In the spring-mass-damper equivalent to the circuit of Fig. 3, L_k^{-1} is analogous to the spring constant relating spring force to spring displacement; L_k^{-1} going to zero describes a “broken” spring that no longer applies any forces to the masses it previously connected.

With this model of branch failure, and recalling that we will associate a (ideally) binary failure indicator variable γ_k with each branch, the first modification to our model, and to ϑ , is straightforward. The arguments of ϑ are augmented to

produce $\vartheta(\phi, \mathbf{q}, \gamma)$, with each L_k^{-1} term in $\vartheta(\phi, \mathbf{q})$ replaced by the product $\gamma_k L_k^{-1}$. It will prove convenient to define

$$\mathbf{N}(\gamma) := \text{diag}\{\gamma_1, \gamma_2, \dots, \gamma_{11}\} \times \mathbf{L}^{-1}.$$

Next, for each inductive branch in our circuit, the current magnitude failure threshold is assumed specified, and the equivalent threshold on inductive energy storage is computed; the resulting constant parameter describing the energy threshold is denoted v_k . The remaining modification to the energy function ϑ is simply an addition of a term of the form $v_k\Theta(\gamma_k)$ for each branch. In summary then, we define our augmented energy function as

$$\begin{aligned} \vartheta(\phi, \mathbf{q}, \gamma) := & \frac{1}{2}\mathbf{q}^T\mathbf{C}^{-1}\mathbf{q} + \frac{1}{2}\phi^T\mathbf{H}_r\mathbf{N}(\gamma)\mathbf{H}_r^T\phi \\ & + [\phi^T \mathbf{q}^T]\mathbf{A}^{-1}\begin{bmatrix} \mathbf{0} \\ \mathbf{i}_{\text{in}} \end{bmatrix} + \sum_{k=1}^{11} v_k\Theta(\gamma_k) \end{aligned} \quad (2)$$

and

$$\tilde{\mathbf{A}} := \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \frac{1}{\tau}\mathbf{I} \end{bmatrix}$$

where τ is a time constant parameter selected to govern the speed of transition in γ (e.g., $\tau = 0.4$ ms is selected for the simulations below). Direct calculation then verifies that the desired state equations for $(\phi, \mathbf{q}, \gamma)$ are recovered as

$$\begin{bmatrix} \dot{\phi} \\ \dot{\mathbf{q}} \\ \dot{\gamma} \end{bmatrix} = \tilde{\mathbf{A}}\nabla\vartheta(\phi, \mathbf{q}, \gamma).$$

Failure Simulations and Properties of the Potential Function

Before examining and interpreting our Lyapunov (energy) function ϑ , it is useful to verify the time domain behavior both of the branch failure model (specifically, the failure indicator variables, $\gamma(t)$), and the claimed nonincreasing property of ϑ along trajectories. As a simple illustrative case, we intentionally set two branch failure thresholds quite low: branches L_5 and L_{10} are set to fail at 0.396 mJ (equivalently, at 8.9 mA), while all the remaining branches are set to thresholds of 2.53 mJ (22.5 mA). Our test inputs remain current steps of +10 mA (in) at node 1, -10mA (out) at node 8. Figs. 5 and 6 provide plots of branch current L_5 with γ_5 versus time and branch current L_{10} with γ_{10} (all on a 12-s interval). As confirmation that no other current magnitudes exceeded their thresholds of 22.5 mA, Fig. 7 provides an ensemble plot of the remaining nine inductor currents.

The state equations as constructed guarantee that the time derivative of ϑ must be nonincreasing along trajectories; the light damping of our circuit suggests that the decay should be slow away from failure event intervals. To confirm these properties, Fig. 8 plots $\vartheta(\phi(t), \mathbf{q}(t), \gamma(t))$ versus time. Note that the transitions associated with failure of branches




Figure 7. Branch currents for all branches other than L_5 , L_{10} , versus time (s).




Figure 8. Decay of energy versus time.

5 and 10 do induce relatively rapid decreases in ϑ , in contrast to the slow decay away from failure zones. This is expected due to the structure of $\tilde{\mathbf{A}}$, whose symmetric part governs dissipation. The large diagonal entries formed by $(-1/\tau)$ terms, with our choice of $\tau = 0.4$ ms, suggest that any excitation of the γ state components away from their equilibria will induce relatively rapid dissipation.

We wish to illustrate the geometric insights available by viewing the state trajectory as a process of phase transition between potential wells about equilibria and how this viewpoint might be used to supplement information available from limited time horizon simulations. Consider the pertinent questions facing an analyst having in hand simulation data through the second branch failure above; e.g., on an interval of 0 to perhaps 4.25 s. A key question to be answered

is whether or not further branch failures will occur as the time interval of interest is extended. Based on the data plotted in Fig. 7, with simulation data out to 12 s, we have strong evidence to suggest that while our system trajectories pass moderately close to a failure threshold (with peak current magnitudes of approximately 19 mA versus the failure threshold of 22.5 mA), no further failures occur. Could we gain qualitative insight from data available at 4 s to support this hypothesis, without the costly simulation out to 12 s? The question would be one of determining whether the state at $t = 4$ s is yet “trapped” in the attractive set about the locally stable equilibrium produced when branches 5 and 10 are out of service. A rigorous examination of this question would involve computation of the lowest energy saddle exit \mathbf{x}^u in the vicinity of this stable equilibrium; conceptually, one expands constant contours of $\vartheta(\mathbf{x})$ from the stable equilibrium until the “closest” critical point \mathbf{x}^u at which $\nabla\vartheta(\mathbf{x}^u) = 0$ is encountered. The nonincreasing property of ϑ along trajectories ensures that the component of the set $\{\mathbf{x} | \vartheta(\mathbf{x}) < \vartheta(\mathbf{x}^u)\}$ containing the stable equilibrium will be invariant, and provided $\vartheta(\mathbf{x}(t))$ cannot be constant away from equilibria, convergence to the local stable equilibrium can be guaranteed. The attractive property of our formulation is that the lowest energy unstable \mathbf{x}^u s in the neighborhood of a stable equilibrium are each naturally associated with a single branch moving to the “cusp” of failure. These \mathbf{x}^u s are therefore easily computed.

To illustrate these ideas graphically in a single contour plot, one is unfortunately limited to examination of the behavior with respect to two degrees of freedom in a planar “cut” of the higher dimensional state space. Judicious selection of these two degrees of freedom is important, and any intuitive insights obtained must be checked against more rigorous conditions for the invariant set in the higher dimensional state space. Nonetheless, the geometric insights gained, even in a two-dimensional plot, are valuable, as our circuit example will illustrate. Given the information in hand following the two branch failures, a natural choice of a planar cut of interest presents itself. Select three points to determine the plane: i) the stable equilibrium with all branches in service; ii) the stable equilibrium with only branch 5 out of service; and iii) the stable equilibrium with branches 5 and 10 out of service. A rectangular planar cut embedded in the higher dimensional state space is constructed, with these equilibrium points at three of the corners. The figures to follow illustrate a variety of views of the energy $\vartheta(\mathbf{x})$ as contour plots or surface plots over this plane. Because the local minima associated with two of the three equilibria are quite “shallow,” two figures are used to display behavior for these two equilibria. The first, as found in Figs. 9(a) and 10(a), are “global” contour views, with axes scaled so that all three equilibria are contained in the range of the axes’ limits. The second, as found in Figs. 9(b) and 10(b), are expanded, local views, shown as surface plots to better illustrate the potential wells. Behavior about the

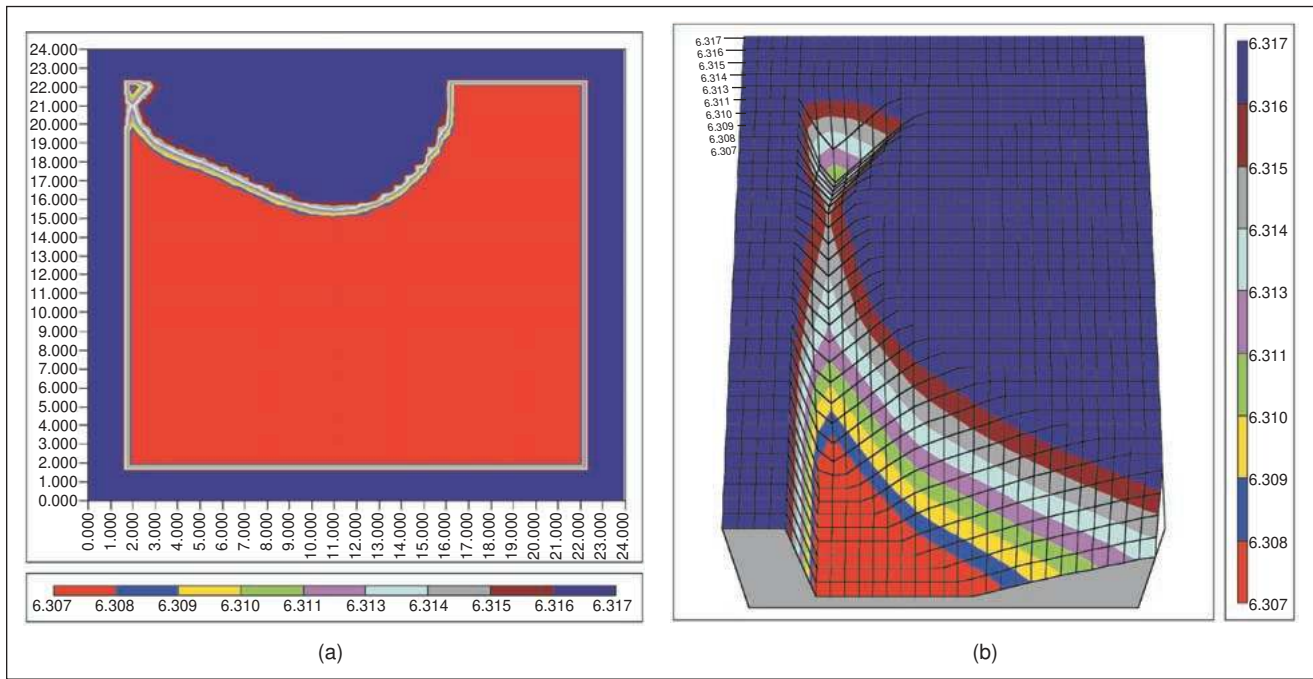


Figure 9. (a) Contour plot of $\vartheta(\mathbf{x})$: Focus on stable equilibrium at (2,22). (b) Surface plot of $\vartheta(\mathbf{x})$: Expanded detail about location (2,22).

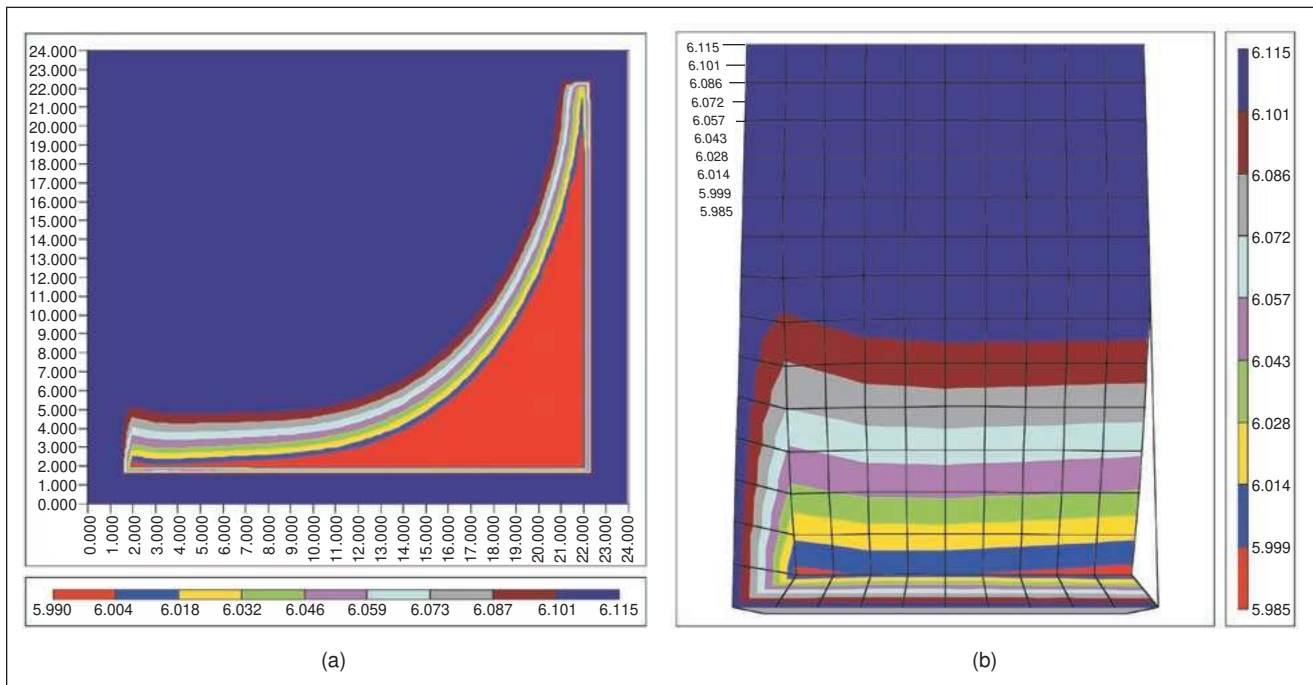


Figure 10. (a) Contour plot of $\vartheta(\mathbf{x})$: Focus on stable equilibrium at (2,2). (b) Surface plot of $\vartheta(\mathbf{x})$: Expanded detail about location (2,2).

third equilibrium of interest is shown in the single contour plot of Fig. 11, because the potential well shown is wide enough in expanse that an expanded view is unnecessary.

As further explanation of the plots, the reader should note that the axes describing the two degrees of freedom are not aligned with any single physical state variable. Therefore, in the global views of Figs. 9(a), 10(a), and 11, the horizontal axis (first coordinate) and vertical axis (second coordinate) are arbitrarily labeled with scales of 0 to 24. The first

plot, Fig. 9(a), focuses on the stable equilibrium associated with all branches being in service; in the coordinate axes shown, this point occurs in the vicinity of the grid location (2,22), with an energy value of approximately 6.307 mJ. Note that Fig. 9(a) displays behavior over a very narrow range of energy variation, and the color assignments “saturate” for any energy value below or above that range. Therefore, all values below an energy of 6.308 are shaded red in Fig. 9(a); all points with energy above 6.316 are shaded in deep blue.




Figure 11. Contour plot of $\vartheta(\mathbf{x})$: Focus on stable equilibrium at (22,2).

The reader should note the small potential well that appears in red about the (2,22) point. The local behavior about that point is shown in more detail in the 3-D surface plot in Fig. 9(b). Note that the lowest saddle point that allows escape from the stable equilibrium appears at the energy level shown in light blue, a level of approximately 6.313, indicating that only a very modest addition of energy to the system can cause a transition to failure. Moreover, this saddle point, which appears near (2,21) in Fig. 9(a), is associated with branch 5 being on the cusp of transition to its failed state. The reader should recall Fig. 2; there $\Theta(\gamma)$ displays a local maximum in the vicinity of $\gamma=0.92$, which creates the saddle point shown in Fig. 9(a) and (b).

The next locally stable equilibrium is associated with the configuration that has branch 5 failed, but branch 10 still operational, and has an energy value of approximately 5.985. Fig. 10(a) displays this point, keeping the same coordinate axes as Fig. 9(a), but reassigns contour colors to focus on values at or slightly above 5.985. Here the local minimum associated with the stable equilibrium is very weak, as indicated by the *very* small region in red about (2,2). The saddle exit point that governs escape from this well is located in the vicinity of (4,2). Analogous to Fig. 9(b), Fig. 10(b) shows an expanded view of energy levels about this stable equilibrium and saddle point as a surface plot. Exit through the saddle point requires only a minuscule addition of energy, to overcome a boundary that has energy of approximately 6.005.

The final contour plot focuses on the stable equilibrium associated with the system configuration having branch 5 and branch 10 failed; this behavior is displayed in the single Fig. 11. Again, the scaling of the horizontal and vertical coordinate axes is identical to that of Figs. 9(a) and 10(a); the

only modification is in the range of contour levels and the assignment of color to level. The stable equilibrium here is associated with the red shaded region in the vicinity of the (22,2) point, with an energy value of approximately 5.40. In contrast to the equilibria seen in Figs. 9 and 10, this point is strongly stable, as indicated by a much deeper “energy well,” which does not offer an easy path of escape through a low energy saddle point in its neighborhood.

With these energy plots in hand, a qualitative interpretation of the time domain state trajectory data of Fig. 7 emerges. The starting energy associated with the system initial conditions is approximately 7.20, placing the initial condition in the energy range far above that of any of our stable equilibria. Therefore, it is quite plausible that the system would not be captured in the very shallow well associated with the “no failure” equilibrium at (2,22). Similarly, even after the energy dissipation of the first branch failure has reduced $\vartheta(\mathbf{x})$ to approximately 6.75, the energy is still not sufficiently small to suggest that the state would be captured by the shallow well about (2,2). Based on the energy contours, one would have the strong hypothesis that the state *would* be captured in the attractive well about equilibrium (22,2). However, an examination of contours with respect to two degrees of freedom within the state space is not a rigorous guarantee that this hypothesis holds; rather, one must search for the lowest energy saddle exit from this region, as described above. Such a search is straightforward, as low energy saddle points in this system are associated with a single branch on the cusp of failure, and are easily located by a Newton-Raphson iteration that initializes the associated branch at 0.92. Such a search here confirms that the saddle point at (4,2) is the lowest exit out of the potential well about (22,2), as the next lowest saddle point, \mathbf{x}^u , is associated with failure of branch L3, with $\vartheta(\mathbf{x}^u) > 7$. With this data in hand, obtained at a computational cost much lower than that of 8 s of additional time domain simulation, one could conclude that no additional branch failures occur.

Conclusions

In this work, we have considered a special structure of dynamic system model that admits a very tractable inclusion of element failure phenomena, for which a global system Lyapunov function can be constructed. This class includes Hamiltonian systems as a special case, with a wide class of R-L-C circuits and mechanical spring-mass-damper systems in which branch failures are induced by exceeding thresholds of inductor current or spring force magnitude. Using a detailed R-L-C circuit as our illustrative example, this article described how geometric features of the global Lyapunov function constructed, along with partial trajectory information from time domain simulations, can be used to more efficiently predict which branches are subject to failure in a specific disturbance scenario. The underlying concepts are closely related to techniques of merging families of Lyapunov functions in hybrid system analysis. However, the

nature of our application, and the construction that yields a clear algorithmic procedure for producing the appropriate global Lyapunov function, present unique features. It is hoped that these techniques will add to the set of tools available for predicting and preventing cascading failure in large-scale networks.

Acknowledgments

This work was supported under the Complex Interactive Networks/Systems Initiative of the Electric Power Research Institute and the U.S. Army Research Office, through project WO 83333-04, subcontract 35352-6087. This support and the untiring efforts of Dr. Massoud Amin of EPRI to foster this initiative are hereby gratefully acknowledged.

References

- [1] P.J. Smith, M. Shafi, and H. Gao, "Quick simulation: A review of importance sampling techniques in communication systems," *IEEE J. Select. Areas Commun.*, vol. 15, pp. 597-613, May 1997.
- [2] J.A. Bucklew, *Large Deviation Techniques in Decision, Simulation, and Estimation*. New York: Wiley, 1990.
- [3] K. Bae and J.S. Thorp, "An importance sampling application: 179 bus WSCC system under voltage based hidden failures and relay misoperations," in *Proc. 31st Hawaii Int. Conf. Syst. Sci.*, Kona, HI, vol. 3, 1998, pp. 39-46.
- [4] C. Asavathiratham, "The influence model: A tractable representation for the dynamics of networked Markov chains," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., MIT, Cambridge, MA, Oct. 2000.
- [5] B. Brogliato, *Nonsmooth Impact Mechanics: Models, Dynamics and Control*. Berlin, Germany: Springer-Verlag, 1996.
- [6] R.A. DeCarlo, M.S. Branicky, S. Pettersson, and B. Lennartson, "Perspectives and results on the stability and stabilizability of hybrid systems," *Proc. IEEE*, vol. 88, pp. 1069-1082, July 2000.
- [7] R. Ortega, A. van der Schaft, I. Mareels, and B. Maschke, "Putting energy back in control," *IEEE Control Syst. Mag.*, vol. 21, pp. 18-33, Apr. 2001.
- [8] A. van der Schaft, *L_2 -Gain and Passivity Techniques in Nonlinear Control*. London, U.K.: Springer-Verlag, 1996.
- [9] R.C. Dorf and R.H. Bishop, *Modern Control Systems*, 9th ed. Upper Saddle River, NJ: Prentice-Hall, 2001.
- [10] C.A. Desoer and E.S. Kuh, *Basic Circuit Theory*. New York: McGraw-Hill, 1969.

Christopher L. DeMarco is Professor of Electrical and Computer Engineering at the University of Wisconsin–Madison, where he has been a member of the faculty since 1985. He received his B.S. degree from the Massachusetts Institute of Technology (MIT), graduating Phi Beta Kappa in 1980. His M.S. and Ph.D. degrees were awarded from the University of California, Berkeley, in 1982 and 1985, respectively, all in electrical engineering. He was a Visiting Associate Professor with the Laboratory for Electromagnetic and Electronic Systems at MIT. He received the University of Wisconsin Chancellor's Award for Outstanding Teaching in May 2000 and was a recipient of the National Science Foundation Presidential Young Investigator Award in 1988.