

Stochastic Analysis of Qualitative Dynamics¹

Jon Doyle

Laboratory for Computer Science
Massachusetts Institute of Technology
545 Technology Square
Cambridge, Massachusetts 02139
617-253-3512

Elisha P. Sacks

Department of Computer Science
Princeton University
Princeton, New Jersey 08544
609-452-4951

Also submitted to IJCAI-89

Abstract

We extend qualitative reasoning with estimations of the relative likelihoods of the possible qualitative behaviors. We estimate the likelihoods by viewing the dynamics of a system as a Markov chain over its transition graph. This corresponds to adding probabilities to each of the transitions. The transition probabilities follow directly from theoretical considerations in simple cases. In the remaining cases, one must derive them empirically from numeric simulations, experiments, or subjective estimates. Once the transition probabilities have been estimated, the standard theory of Markov chains provides extensive information about asymptotic behavior, including a partition into persistent and transient states, the probabilities for ending up in each state, and settling times. Even rough estimates of transition probabilities provide useful qualitative information about ultimate behaviors, as the analysis of many of these quantities is insensitive to perturbations in the probabilities. The algorithms are straightforward and require time cubic in the number of qualitative states. The analysis also goes through for symbolic probability estimates, although at the price of exponential-time worst-case performance.

¹Authors listed in alphabetical order. Jon Doyle is supported by the National Library of Medicine under National Institutes of Health Grant R01 LM04493. Address correspondence to Jon Doyle.

1 Introduction

Qualitative reasoning seeks to predict the global behavior of a complex dynamic system by partitioning its state space into a manageable number of qualitative states and characterizing its behavior by the sequences of qualitative states that it can go through. This methodology is too weak to describe the limiting behavior of dynamic systems. For example, a damped pendulum eventually must approach equilibrium either directly below or directly above its pivot (Figure 1). The first possibility is almost certain, whereas the second almost never occurs. Qualitative simulation discovers both equilibria, but can neither determine their relative likelihoods nor rule out the possibility that the pendulum will spin forever. Yet qualitative considerations suffice for both conclusions, independent of the numeric details of the system.

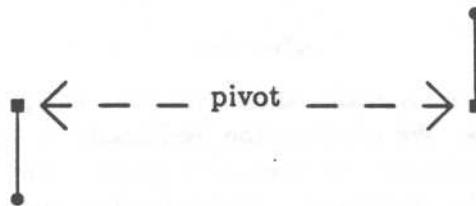


Figure 1: Equilibria of a damped pendulum.

Limiting behaviors are global characteristics of a system. To understand them, we must look beyond individual transitions to sequences of transitions. We must assign each sequence a likelihood, ranging from impossible to definite. In this paper, we describe one approach to this problem. We estimate the likelihoods of a system's asymptotic behaviors by viewing its dynamics as a Markov chain over its transition graph. This corresponds to adding probabilities to each of the transitions. The transition probabilities follow directly from theoretical considerations in simple cases such as the pendulum example. In the remaining cases, one must derive them empirically from numeric simulations, experiments, or subjective estimates.

The standard theory of Markov chains provides extensive information about asymptotic behavior, smoothly blending qualitative and quantitative information into a unifying framework that provides the best possible conclusions given the evidence. It derives some sorts of essentially qualitative conclusions that qualitative simulation does not, including a partition into persistent and transient states and a partition of the persistent states into the probable and the improbable. Many of these conclusions follow solely from qualitative considerations. The remainder, though numerically derived, are insensitive to perturbations in the probabilities. The theory also provides quantitative refinements of these qualitative conclusions, including the mean and variance of settling times. Unlike the qualitative conclusions, the quantitative results are in some cases sensitive to variations in the input probabilities. The

algorithms are straightforward, involving only a topological sort of the transition graph and a few matrix operations on the transition probabilities, and require time at most cubic in the number of qualitative states. The analysis goes through for symbolic probability estimates, although at the price of exponential-time worst-case performance.

2 Qualitative dynamics in phase space

The first step in our analysis of a system is to derive its states and transitions from the phase space of the system. The phase space for a system of first-order differential equations

$$x'_i = f_i(x_1, \dots, x_n); \quad i = 1, \dots, n \quad (1)$$

is the Cartesian product of the x_i 's domains. One can convert higher-order equations to first-order ones by introducing new variables as synonyms for higher derivatives. Points in phase space represent states of the system. Curves on which equation (1) is satisfied, called *trajectories*, represent solutions. A *phase diagram* for a system depicts its phase space and trajectories graphically. The topological and geometric properties of trajectories characterize the qualitative behavior of solutions. For instance, a point trajectory, called a *fixed point*, indicates an equilibrium solution, whereas a closed curve indicates a periodic solution. A fixed point is *stable* if every nearby trajectory approaches it asymptotically and *unstable* otherwise. More generally, the *basin* of a fixed point is the set of trajectories that approach it asymptotically. (See Hirsch and Smale [3].)

For example, the standard model for a damped pendulum is

$$\theta'' + \frac{\mu}{m}\theta' + \frac{g}{l}\sin\theta = 0, \quad (2)$$

with θ the angle between the arm and the vertical, l the length of the (weightless rigid) arm, m the mass of the bob, g the gravitational constant, and μ the damping coefficient, as shown in Figure 2 with the pendulum's phase diagram. The symmetry and 2π periodicity of the pendulum equation make it natural to employ the cylindrical phase space obtained by identifying the lines $\theta = -\pi$ and $\theta = \pi$. Two trajectories spiral toward the unstable fixed point at $(\pi, 0)$; the rest spiral toward the stable fixed point at $(0, 0)$.

A complete qualitative description of a system consists of a partition of its phase space into sets of qualitatively equivalent trajectories. The equivalence criterion depends on the problem task. Mathematicians generally focus on topological equivalence, whereas coarser relations are more useful in engineering applications. We follow standard AI practice and equate all trajectories that go through a specific sequence of regions in phase space. Our qualitative dynamics consists of a partition of phase space into regions along with a graph of possible transitions between regions.

Sacks [7] shows how to translate traditional qualitative reasoning into our qualitative dynamics without loss of information or increase in complexity. Qualitative states correspond

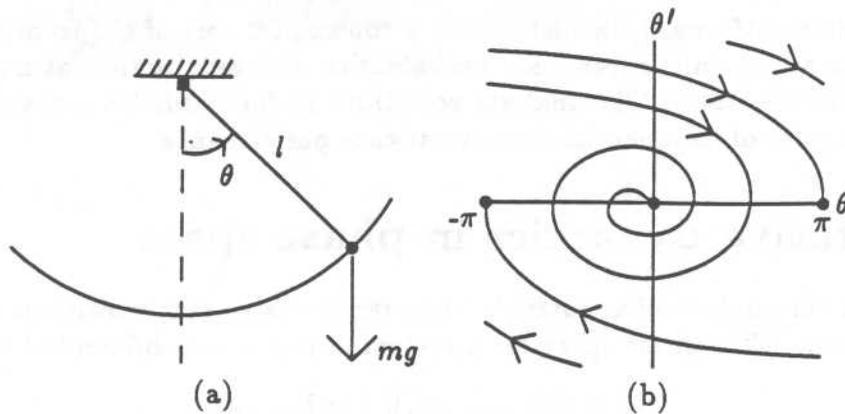


Figure 2: A damped pendulum (a) and its phase diagram (b).

to rectangular regions in phase space, and qualitative simulation amounts to finding the possible transitions between regions. For example, automatic analysis of the damped pendulum equation results in ten qualitative states corresponding to four rectangles ($\langle \pm, \pm \rangle$), four line segments ($\langle \pm, 0 \rangle$), and two fixed points ($\langle 0, 0 \rangle, \langle \pi, 0 \rangle$), as shown with the transition graph over these qualitative states in figure 3.

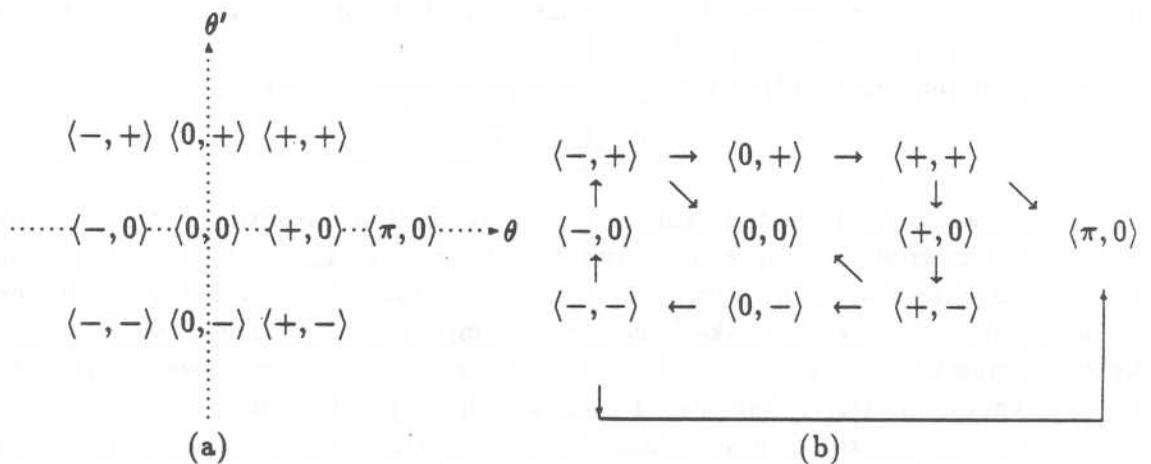


Figure 3: Phase space regions (a) and transition graph (b) of the qualitative states of the damped pendulum.

The probabilistic conclusions about the pendulum's asymptotic behavior follow from purely dimensional arguments, which in turn may be determined by inspection of the eigenvalues of the fixed points. By the stable manifold theorem [2, p. 13], the dimension of the

basin of a fixed point equals the number of its eigenvalues that have negative real parts.² The real parts of the eigenvalues are both negative at $(0, 0)$ and of opposite signs at $(\pi, 0)$. Hence, the basin of $(\pi, 0)$ has zero measure because it is a one-dimensional submanifold of the two-dimensional phase space, whereas the basin of $(0, 0)$ has positive measure. A similar argument shows the unstable equilibria of any system to be unlikely. But this approach is insufficient to yield all conclusions of interest in general, since systems may have multiple stable equilibria and other states with positive asymptotic probabilities.

3 Transforming flows into Markov chains

We estimate the likelihoods of a system's asymptotic behaviors by constructing a Markov process whose states are regions in phase space. The dynamic system itself is an uncountable Markov process whose states are the points of phase space,³ but direct analysis or use of this Markov process is impractical. Instead, we lump the point-states into a manageable number of regions in phase space, each of which represents a distinct qualitative state of the system. Sacks [5, 6, 8] presents a system that automatically identifies such regions and the possible transitions between them for second-order systems of ordinary differential equations. Most of the ideas extend directly to larger systems. One can also work with the rectangular regions that qualitative simulation implicitly defines.

The second step in constructing the Markov process is to associate probabilities with each of the possible transitions between regions. These probabilities reflect the likelihood of the system's state moving from one state-space region to another in unit time. We may obtain the probabilities directly from dimensional considerations in simple cases such as the pendulum example. Ordinarily, though, they will come from numeric simulations or physical experiments that sample representative points in each qualitative region and count how many go to each region. One can also obtain subjective estimates from domain experts. The qualitative analysis is insensitive to errors in these probabilities.

The transition probabilities represent the imprecision in the qualitative model of the dynamic system. If we were able to choose as regions the actual attractors of the system, there would be no imprecision and the transition rules would be perfectly deterministic. The great difficulty of determining the optimal set of regions for analysis helps motivate the stochastic approach to analyzing behaviors. Additional factors that transition probabilities can model include (1) uncertainty about initial conditions which induces a distribution of possible trajectories, (2) uncertainty about the parameters of the model equations, and (3) uncertainties (or noise sources) explicitly occurring in the system's equations, as in stochastic differential equations.

²This theorem applies to hyperbolic fixed points. The number of nonpositive real parts provides an upper bound for nonhyperbolic fixed points, which suffices for our analysis.

³In fact, this Markov process is deterministic: all state transition probabilities are zero or one.

Our analysis treats the Markov process constructed from a dynamic system as a *Markov chain*. That is, we assume that the transition probabilities remain constant over time and that they depend only on the qualitative state of origin, independent of past history. The first assumption holds for *autonomous* equations that are free of their independent variable. One can reduce any general system to an autonomous one by treating the independent variable, t , as a state variable governed by the equation $t' = 1$.

The second assumption holds to the extent that the future trajectory of the system is insensitive to its distant past. The most questionable case is that of conservative systems in which the volume of each region in phase space is preserved for all time by the flow, causing small differences between trajectories to retain their significance forever. Conservative systems pose problems for qualitative reasoning generally, not just for the stochastic analysis, as the regions of interest must be chosen carefully. Fortunately, most realistic systems are dissipative, hence volume shrinking, causing differences between trajectories to shrink exponentially.

Time-dependent transition probabilities imply that the current partition of phase space is too coarse: differences within a prior region express themselves in the current region because the distances between points in the prior region are too great to damp out in one time step. One approach, following Sacks [5, 6], involves iterative improvement of the model (though unlike that work, we have not automated this refinement process). If one observes time-dependent behavior in constructing the transition probabilities, one subdivides or otherwise refines the set of regions and starts over. In principle, the process ends when the chain assumption appears correct for all regions, but in practice the choice of when to accept a model as satisfactory involves a tradeoff of model complexity against model accuracy. The aptness of the chain assumption can also be tested against empirical observations or long term numeric simulations.

4 Analysis of Markov chains

In this section, we summarize the theory of Markov chains. The details appear in our longer paper [1] and in Roberts [4]. Let $S = \{s_1, \dots, s_n\}$ be the set of states of the qualitative dynamics, that is, the set of nodes of the dynamic digraph. Each of these will also be a *state* of the Markov chain. We describe the entire chain by specifying, for each nonexclusive choice of states s_i and s_j , the transition probability p_{ij} that if the system is in state s_i at one instant, it will be in state s_j after one time unit has passed. We write $P = \{p_{ij}\}$ to mean the $n \times n$ *transition matrix* of all transition probabilities. P is also called a *stochastic matrix*, which means that all entries are nonnegative and that each row sums to 1. Each row of P is called a *probability vector*. The *transition digraph* of a stochastic matrix is the graph over the states with a directed arc from s_i to s_j iff $p_{ij} \neq 0$.

The probability that the chain is in state s_j at time t given that it starts in state s_i

at time 0 is written $p_{ij}^{(t)}$ and called a *higher-order transition probability*. This probability is the i, j entry of P^t , the t 'th power of P . If we start the Markov chain at random, where the probabilities of starting in each state are given by an *initial probability vector* $p^{(0)} = (p_1^{(0)}, \dots, p_n^{(0)})$, then the probabilities of being in particular states at time t , $p^{(t)} = (p_1^{(t)}, \dots, p_n^{(t)})$, are given by the equation $p^{(t)} = p^{(0)} P^t$.

A set C of states is *closed* if $p_{ij} = 0$ for all $s_i \in C$ and $s_j \notin C$, that is, if once in C the chain can never leave C . A closed set C is *ergodic* if no proper subset is closed. A state is ergodic if it is in some ergodic set, and is *transient* otherwise. A state that forms an ergodic set by itself is called an *absorbing* state. Chains whose states form a single ergodic set are called ergodic chains, and chains in which each ergodic set is a singleton are called absorbing chains.

The mathematical analysis of the asymptotic behavior of a Markov chain is divided into two parts: the behavior before entering an ergodic set, and the behavior after entering one. One then combines these sub-analyses to get the overall asymptotic behavior.

1. In the first step, one creates an absorbing chain by lumping all states in each ergodic set into a single compound state. The transition probability from a transient state s to a compound state c is the sum of the transition probabilities from s to the members of c . The transition probability from c to other states is 0 by definition. The main result of the analysis is the long-term probability of entering each ergodic set when starting in each of the transient states. The probability of eventually entering some absorbing state is 1.
2. In the second step, one analyzes each ergodic set as a separate ergodic chain, unaffected by the other states. The result of the analysis is the long-term probability of being in each of the states of the set.

Combining these results yields the long-term probability of being in each of the states of the chain. This is just the product of the probability of entering an ergodic set containing that state (this is zero if the state is transient) times the probability of appearing in that state once in the ergodic set. Stability analysis shows that these asymptotic probabilities are insensitive to variations in the input probabilities.

5 Examples

We present two examples to illustrate the techniques of stochastic analysis. The first makes precise the analysis of the damped gravitational pendulum. The second example, that of a charged pendulum in the presence of two other charges, is representative of a large class of everyday systems in which there are several asymptotic behaviors of nonzero probability. The analysis of the charged pendulum calculates these probabilities, and examines the dependence

of their sizes on the magnitude of the charges. We have applied stochastic analysis to other problems as well, including systems which exhibit chaotic behaviors in some regions, but space limitations do not permit us to present these examples here (see [1]).

With the Markov theory in hand, we can make precise our intuitive analysis of the damped pendulum equation

$$\theta'' + \frac{\mu}{m}\theta' + \frac{g}{l}\sin\theta = 0. \quad (3)$$

The transition graph forms an absorbing chain with absorbing states $\langle 0, 0 \rangle$ and $\langle \pi, 0 \rangle$. The pendulum must eventually approach one of these states; it cannot cycle between the remaining, transient states forever. The transition probabilities into $\langle \pi, 0 \rangle$ are 0 by dimensional analysis, as discussed in Section 2, so all trajectories end up at $\langle 0, 0 \rangle$ with probability 1. If the transition probabilities from $\langle -, + \rangle$ and $\langle +, - \rangle$ to $\langle 0, 0 \rangle$ are each $\alpha \ll 1$, then the settling time when starting in a transient state is approximately $4/\alpha$.

The damped pendulum model presupposes that the force on the bob is independent of the bob's location. The model for the variable attraction between a positively charged pendulum bob and two negative charges is more complicated (Figure 4). Each negative charge exerts a force

$$f(\alpha, k) = k(d+l)l^{-5}(\sin\alpha)(d^2 + dl + 2l^2 - 2l(d+l)\cos\alpha)^{-3/2} \quad (4)$$

along the line between it and the bob, with α the angle between that line and the pendulum, k the coefficient of electrostatic attraction between the bob and the charge, l the length of the arm, and d the vertical distance from the bob's orbit to the line connecting the two negative charges. The two-charge pendulum obeys the equation

$$\theta'' + \frac{\mu}{m}\theta' + f(\theta + a, k_1) + f(\theta - a, k_2) = 0 \quad (5)$$

with a the angle between each pole and the vertical, k_1 and k_2 corresponding to the left and right charges, μ the damping coefficient, and m the mass of the bob. Figure 5 contains the qualitative dynamics for the case of equal charges ($k_1 = k_2$). Saddles appear at $(\pi, 0)$ and $(0, 0)$ where the charges cancel each other. A sink appears where each charge is strongest. The pendulum can spin (A-B-C-D and E-F-G-H), oscillate around both negative charges (A-B-C-D-E-F-G-H), oscillate around the left charge (A-B-G-H), or oscillate around the right charge (C-D-E-F). It can also switch from spinning to oscillating and from oscillating around both charges to oscillating around either charge.

The sinks are located at the points on the θ axis where $f(\theta + a, k_1) + f(\theta - a, k_2) = 0$. We cannot locate them exactly because the equation has no closed-form solution and because the parameters are subject to measurement error. Instead, we bound each sink within a small region of phase space by numeric simulation or by experimentation. Unlike the sinks, which have zero exit probabilities a priori, the bounding regions can have positive exit probabilities, which must be estimated empirically. One can either incorporate the exit probabilities into the Markov model or shrink the bounding regions until their exit

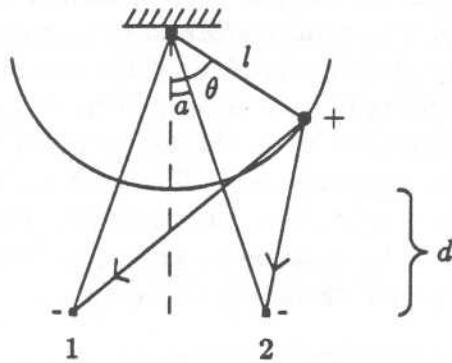


Figure 4: A positively charged pendulum attracted by two negative charges.

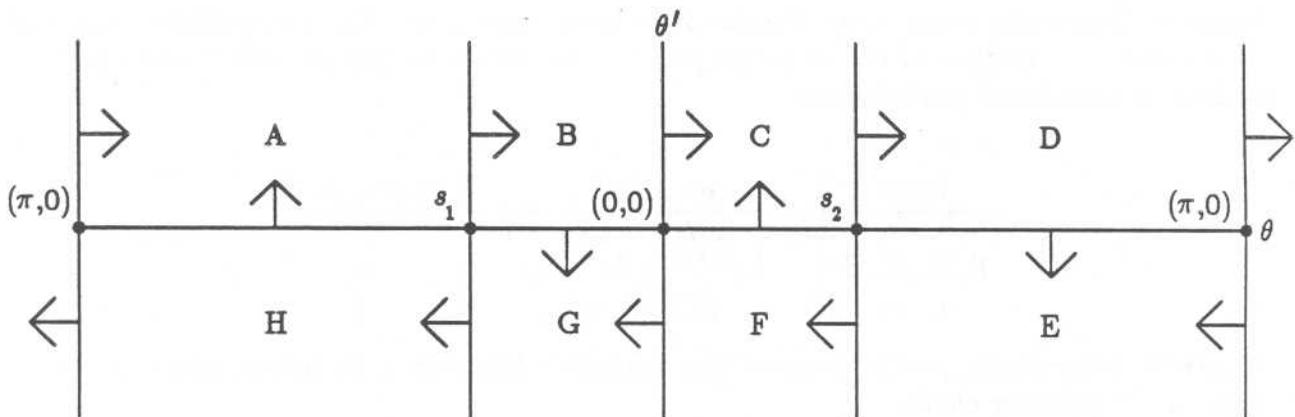


Figure 5: Qualitative dynamics for the two-charge pendulum.

probabilities become negligible. The choice involves a tradeoff between modeling time and analysis time.

Figure 6 illustrates the first strategy. We assign each bounding region entrance probability p and exit probability q , assign the saddles zero probabilities, and assume the remaining transitions to be equally likely. The resulting Markov chain consists of three ergodic chains: the two saddles and a primary chain containing all the remaining regions. The asymptotic probabilities of the pendulum being in each region, given that it starts in the primary chain, appear in Figure 7. As q approaches zero, the probabilities of the sinks approach .5 and those of the remaining regions approach zero. The sink probabilities depend only on the ratio p/q ; for example given $p/q = 10$, they are each .37. The second strategy derives the asymptotic probabilities directly by treating s_1 and s_2 as absorbing states, but incurs the difficulty of constructing regions on which q is negligible.

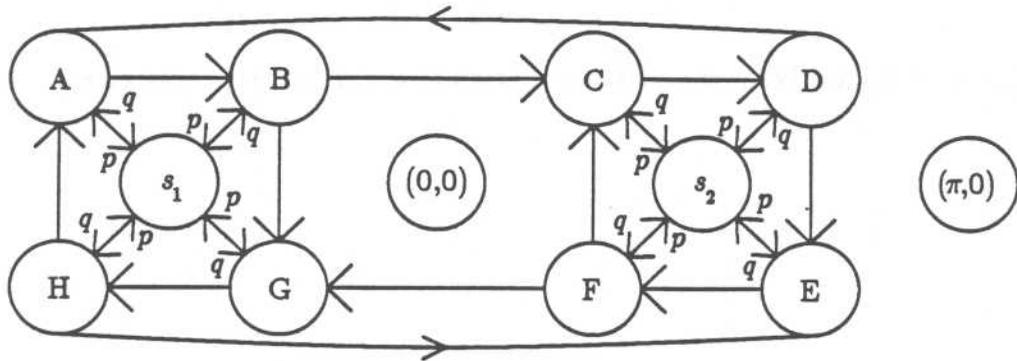


Figure 6: Transition probabilities for the two-charge pendulum. The probabilities from and to the bounding regions of the sinks are p and q respectively. The unmarked transitions of each node have equal probabilities.

state	probability	limit $q \rightarrow 0$
A, C, E, G	$q(4-p)/(7-4)(p+4q)$	0
B, D, F, H	$3q/(7-p)(p+4q)$	0
s_1, s_2	$p/2(p+4q)$.5

Figure 7: Asymptotic probabilities of the pendulum being in each region, given an initial state in the primary chain.

When the two nearby charges are of unequal magnitude, the unstable equilibria move away from $(0,0)$ and $(\pi,0)$ and the stable equilibria are positioned asymmetrically around the θ' axis, but otherwise the qualitative dynamics appears just as in Figure 5. The transition probabilities, however, change to reflect the greater and lesser attractions of the unequal charges, with entry probabilities p_1 and p_2 and exit probabilities q_1 and q_2 . Let p_1^∞ and p_2^∞

respectively represent the asymptotic probabilities of appearing in states s_1 and s_2 , averaged over all possible transient starting states. For simplicity, we assume that trajectories cannot escape the trapping regions, so that $q_1 = q_2 = 0$. Calculating the ratio $r^\infty = p_1^\infty / p_2^\infty$ yields

$$r^\infty = \left(\frac{p_1}{p_2} \right) \left(\frac{(p_1 - 1)p_2^2 - (3p_1 + 21)p_2 + 2p_1 - 14}{(p_1^2 - 3p_1 + 2)p_2 - p_1^2 - 21p_1 - 14} \right) \quad (6)$$

The dependence of r^∞ on p_1 and p_2 agrees with our intuitions. The ratio increases monotonically from 0 as p_1 increases from 0 to 1 and decreases monotonically from ∞ as p_2 increases from 0 to 1. It equals 0 when $p_1 = 0$, ∞ when $p_2 = 0$, and 1 when p_1 equals p_2 .

6 Conclusions and future work

We apply the theory of Markov chains to estimate the likelihoods of possible behaviors of a system, thereby filling a serious gap in the predictions of qualitative simulation. This theory enables us to draw the best possible conclusions from the available information. We can determine the possible long term behaviors of a system directly from its qualitative dynamics. More detailed information, such as the likelihoods of the possible behaviors and the expected settling times for each initial state, require estimates of the transition probabilities between qualitative states. The estimates can be numeric or symbolic; the analysis is formally identical in both cases, but has $O(n^3)$ time-complexity in the former and exponential time-complexity in the latter. We exhibit the utility of our method in several examples, and analyze the robustness of its conclusions to perturbations in the transition probabilities. The likelihoods of the long term behaviors are never sensitive to perturbations in the transition probabilities, whereas the expected settling times can be sensitive.

Our current analysis is only a first step towards full exploitation of the stochastic approach to qualitative reasoning. We have not fully explored the potential of Markov theory. Further investigation may yield simple ways of determining other qualitative properties of systems through application of known techniques. One might also relax the chain assumption underlying our treatment and instead view the qualitative dynamics as describing a more general Markov process in which transitions depend on past states. There is a rich theory of these processes which may support many of the same conclusions as above in the more general setting. Incorporating global properties of flows, such as energy conservation, into stochastic analysis is another topic for future research. Finally, automating the model refinement is a major machine learning challenge.

References

- [1] Jon Doyle and Elisha P. Sacks. Stochastic analysis of qualitative dynamics. Submitted for publication, 1988.
- [2] John Guckenheimer and Philip Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1986.
- [3] Morris W. Hirsch and Stephen Smale. *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press College Division, Orlando, Florida, 1974.
- [4] Fred S. Roberts. *Discrete Mathematical Models*. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1976.
- [5] Elisha P. Sacks. Automatic qualitative analysis of ordinary differential equations using piecewise linear approximations. TR 416, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, March 1988. Also appears as AI-TR-1031.
- [6] Elisha P. Sacks. Automatic qualitative analysis of ordinary differential equations using piecewise linear approximations. To appear in *Artificial Intelligence*, 1988.
- [7] Elisha P. Sacks. A dynamic systems perspective on qualitative simulation. To appear in *Artificial Intelligence*, 1988.
- [8] Elisha P. Sacks. Piecewise linear abstraction of intractable dynamic systems. *International Journal of Artificial Intelligence in Engineering*, July 1988.