A closer look at the classical fixed-point analysis of WLANs

Rajesh Sundaresan

Indian Institute of Science

May 2017

Section 1

The classical fixed point analysis

DCF: The 802.11 countdown and its Markovian caricature

- N nodes accessing the common medium in a wireless LAN. Infinite backlog of packets. Attempts to transmit HOL packet.
- Each node's (backoff) state space: $\mathcal{Z} = \{0, 1, \cdots, m-1\}$. Backoff state determines attempt probability for a node in a slot.
- Transitions:

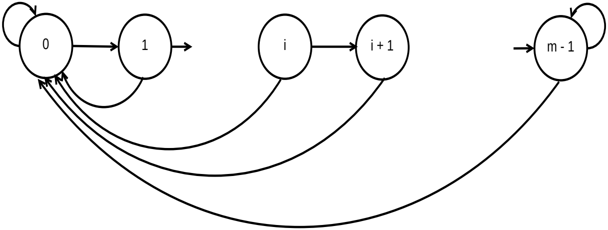

Example design – exponential backoff

- Assume three states, $\mathcal{Z} = \{0, 1, 2\}$ or $m = 3$.
- Attempt probability for node in state i is c_i/N.
- Aggressiveness of the transmission $c = (c_0, c_1, c_2)$.
- The scaling by $1/N$ ensures that the overall attempt probability of a single node is $O(1/N)$ so that the overall (system) attempt probability for the system is $O(1)$.
- Conventional wisdom: exponential backoff:

$$c_i = c_{i-1}/2.$$

Double the average waiting time after every failure.

Back-of-envelope analysis

- Observation: your collision probability depends only on the empirical measure of node states ... excepting you.
- ξ = current empirical measure of nodes across states.
- Number of nodes across states is $(N\xi_0, N\xi_1, \ldots, N\xi_{m-1})$.
- If you are in state 0, others states $(N\xi_0 - 1, N\xi_1, \ldots, N\xi_{m-1})$. Probability that no one else transmits is:

$$\left(1 - \frac{c_0}{N}\right)^{N\xi_0 - 1} \cdot \prod_{i=1}^{m-1} \left(1 - \frac{c_i}{N}\right)^{N\xi_i} = \left(1 - \frac{c_0}{N}\right)^{-1} \cdot \prod_{i=0}^{m-1} \left(1 - \frac{c_i}{N}\right)^{N\xi_i}$$
$$\to e^{-\langle c, \xi \rangle}.$$

- $\langle c, \xi \rangle$ is the attempt probability: $\sum_i (N\xi_i)(c_i/N)$.
- If N is small or if attempt probabilities don't scale, avoid the limit.

The classical fixed-point analysis

1. Conditional collision probability, when making an attempt, is the same for each node in each state.

   $$\gamma := 1 - e^{-\langle c,\xi \rangle} = 1 - e^{-(attempt)}$$

   This amounts to assuming that spatial distribution stabilises at ξ.

2. The system interactions decouple.

3. Focus on a node. Consider renewal instants of return to state 0. From the renewal-reward theorem:

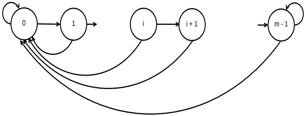

$$\frac{attempt}{N} = \frac{E[Reward]}{E[RenewalTime]} = \frac{1+\gamma+\gamma^2+\dots}{\frac{N}{c_0}+\gamma\frac{N}{c_1}+\gamma^2\frac{N}{c_2}+\dots} =: \frac{G(\gamma)}{N}.$$

4. Solve for the fixed point: $\gamma = 1 - e^{-G(\gamma)}$.

Goodness of the approximation (from Bianchi 1998)

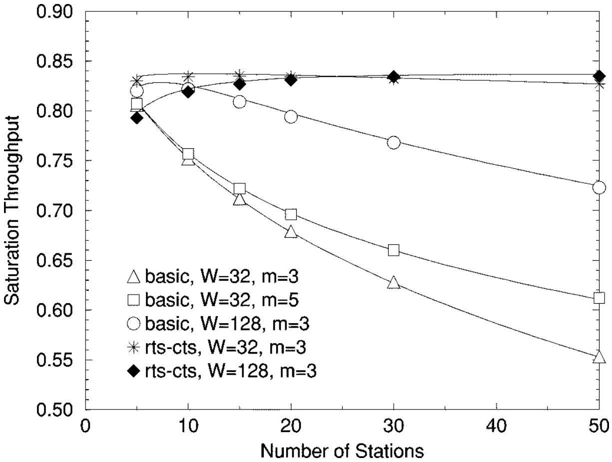

Plot for fixed-point analysis without taking $N \to \infty$.
W is the window size in the basic WLAN protocol.

In this talk:

We will see an overview of

- why decoupling is a good assumption;
- when node independent, state independent, conditional collision probability assumption holds;
- and going a little beyond

  what to do when the 'node/state independent collision probability' does not hold.

Section 2

The decoupling assumption

Mean-field interaction

- Coupled dynamics.
- Embed slot boundaries on $\mathbb{R}_+$. Assume slots of duration $1/N$.
- Transition rate = prob. of change in a slot / slot duration = $O(1)$.
- Transition rate for success or failure depends on the states of the other nodes, but only through the empirical measure $\mu_N(t)$ of nodes across states.
- At time t, node transition rates are as follows:
  - $i \rightsquigarrow i+1$ with rate $\lambda_{i,i+1}(\mu_N(t))$.
  - $i \rightsquigarrow 0$ with rate $\lambda_{i,0}(\mu_N(t))$.
  - In general, $i \rightsquigarrow j$ with rate $\lambda_{i,j}(\mu_N(t))$.

The transition rates

If $\mu_N(t) = \xi$, then

- Example:

$$\lambda_{0,1}(\xi) = \frac{(c_0/N)(1 - e^{-attempt})}{1/N} = c_0(1 - e^{-\langle c,\xi \rangle}).$$

- Write as a matrix of rates: $\Lambda(\cdot) = [\ \lambda_{i,j}(\xi)\ ]_{i,j \in \mathcal{Z}}$.

- For ξ, the empirical measure of a configuration, the rate matrix is

$$\Lambda(\xi) = \begin{bmatrix} - & c_0(1 - e^{-\langle c,\xi \rangle}) & 0 \\ c_1 e^{-\langle c,\xi \rangle} & - & c_1(1 - e^{-\langle c,\xi \rangle}) \\ c_2 e^{-\langle c,\xi \rangle} & 0 & - \end{bmatrix}.$$

For today's exposition, we will assume this continuous-time caricature with these instantaneous transition rates.

This is different, since at most one node can transit at any time.

The Markov processes, big and small

- $(X_n^{(N)}(\cdot),\ 1 \leq n \leq N)$, the trajectory of all the n nodes, is Markov
- Study $\mu_N(\cdot)$ instead, also a Markov process
  Its state space size is the set of empirical probability measures on N particles with state space $\mathcal{Z}$.

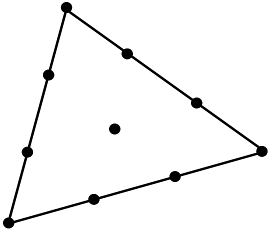

- Then try to draw conclusions on the original process.

The smaller Markov process $\mu_N(\cdot)$

- A Markov process with state space being the set of empirical measures of N nodes.

- This is a measure-valued flow across time.

- In the continuous-time version:
  the transition $\xi \rightsquigarrow \xi + \frac{1}{N}e_j - \frac{1}{N}e_i$ occurs at rate $N\xi(i)\lambda_{i,j}(\xi)$.

- For large N, changes are small, $O(1/N)$, at higher rates, $O(N)$. Individuals are collectively just about strong enough to influence the evolution of the measure-valued flow.

- Fluid limit : μ_N converges to a deterministic limit given by an ODE.

The conditional expected drift in μ_N

- Recall $\Lambda(\cdot) = [\ \lambda_{i,j}(\cdot)\ ]$ without diagonal entries. Then

$$\lim_{h \downarrow 0} \frac{1}{h} \mathbb{E}\left[\mu_N(t+h) - \mu_N(t) \mid \mu_N(t) = \xi\right] = \Lambda(\xi)^* \ \xi$$

with suitably defined diagonal entries.

An interpretation

- The rate of change in the kth component is made up of increase

$$\sum_{i:i\neq k} (N\xi_i) \cdot \lambda_{i,k}(\xi) \cdot (+1/N)$$

- and decrease

$$(N\xi_k) \sum_{i:i\neq k} \lambda_{k,i}(\xi)(-1/N).$$

- Put these together:

$$\sum_{i:i\neq k} \xi_i \lambda_{i,k}(\xi) - \xi_k \sum_{i:i\neq k} \lambda_{k,i}(\xi) = \sum_i \xi_i \lambda_{i,k}(\xi) = (\Lambda(\xi)^* \xi)_k.$$

The conditional expected drift in μ_N

- Recall $\Lambda(\cdot) = [\ \lambda_{i,j}(\cdot)\ ]$ without diagonal entries. Then

$$\lim_{h \downarrow 0} \frac{1}{h} \mathbb{E}\left[\mu_N(t+h) - \mu_N(t) \mid \mu_N(t) = \xi\right] = \Lambda(\xi)^* \ \xi$$

  with suitably defined diagonal entries.

- Anticipate that $\mu_N(\cdot)$ will solve (in the large N limit)

$$\begin{aligned} \dot{\mu}(t) &= \Lambda(\mu(t))^* \ \mu(t), \quad t \geq 0 \qquad \text{[McKean-Vlasov equation]} \\ \mu(0) &= \nu \end{aligned}$$

- Nonlinear ODE.

A limit theorem

Theorem

Suppose that the initial empirical measure $\mu_N(0) \xrightarrow{p} \nu$*, where* ν *is deterministic.*

Let $\mu(\cdot)$ *be the solution to the McKean-Vlasov dynamics with initial condition* $\mu(0) = \nu$*.*

Then $\mu_N(\cdot) \xrightarrow{p} \mu(\cdot)$*.*

Technicalities:

- The McKean-Vlasov ODE must be well-posed.
- $\mu_N(0) \xrightarrow{p} \nu$: Probability of being outside a ball around ν vanishes.
- $\mu_N(\cdot) \xrightarrow{p} \mu(\cdot)$: For any finite duration, probability of being outside a tube around $\mu(\cdot)$ vanishes.

Back to the individual nodes

- Let $\mu(\cdot)$ be the solution to the McKean-Vlasov dynamics
- Choose a node uniformly at random, and tag it.
  - $\mu_N(\cdot)$ is the distribution for the state of the tagged node at time t.
  - As $N \to \infty$, the limiting distribution is then $\mu(t)$

Joint evolution of tagged nodes

- Tag k nodes.
  - If the interaction is only through $\mu_N(t)$, and this converges to a deterministic $\mu(t)$, the transition rates are just $\lambda_{i,j}(\mu(t))$.
  - Each of the k nodes is then executing a time-dependent Markov process with transition rate matrix $\Lambda(\mu(t))$.
  - Asymptotically, no interaction ... decoupling.
  - The node trajectories are (asymptotically) independent and identically distributed.

Section 3

The steady state assumption

The fixed-point analysis again

- Solve for the rest point of the dynamical system.

$$\Lambda(\xi)^*\xi = 0.$$

- If the solution is unique, predict that the system will settle down at $\xi \otimes \xi \otimes \ldots \otimes \xi$.

- Works very well for the exponential backoff.

- But not in general due to limit cycles.

A malware propagation example (from Benaim and Le Boudec 2008)

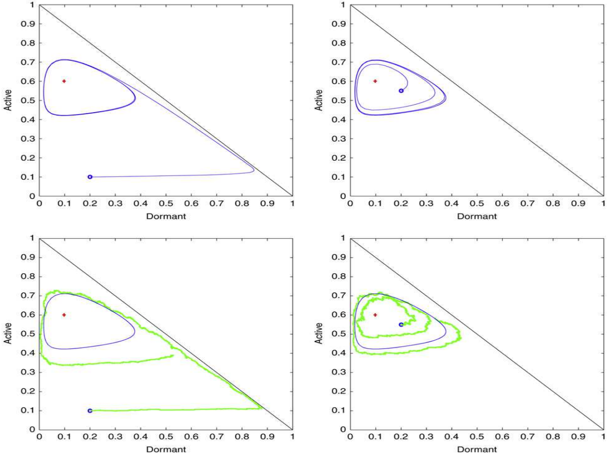

- The fixed point is unique, but unstable.
- All trajectories starting from outside the fixed point, and all trajectories in the finite N system, converge to the stable limit cycle.

What is the issue?

- Large time behaviour for a finite N system: $\lim_{t\to\infty} \mu_N(t)$.
  If N is large, we really want:

$$\lim_{N\to\infty} \left[\lim_{t\to\infty} \mu_N(t)\right].$$

- But we are trying to predict where the system will settle from the following:

$$\lim_{t\to\infty} \left[\lim_{N\to\infty} \mu_N(t)\right] = \lim_{t\to\infty} \mu(t).$$

- We need a little bit of robustness of the ODE for this work.

Does the method work?

Theorem

Let $\mu_N(0) \to \nu$ *in probability.*

Let the ODE have a (unique) globally asymptotically stable equilibrium ξ_f *with every path tending to* ξ_f.

Then $\mu_N(\infty) \xrightarrow{d} \xi_f$.

It is not enough to have a unique fixed point ξ_f.
But if that ξ_f is globally asymptotically stable, that suffices.

A sufficient condition

A lot of effort has gone into identifying when we can ensure a globally asymptotic stable equilibrium.

Theorem

If c is such that $\langle c, \xi \rangle < 1$ for all ξ, then the rest point ξ_f of the dynamics is unique, and all trajectories converge to it.

This is the case for the classical exponential backoff with $c_0 < 1$.

The case of multiple stable equilibria for the ODE

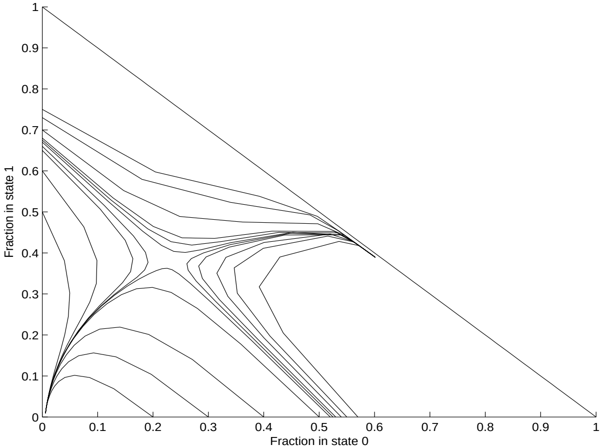

- Different parameters: $c = (0.5, 0.3, 8.0)$.
- There are two stable equilibria.
  One near $(0.6, 0.4, 0.0)$ and another near $(0, 0, 1)$.

The case of multiple stable equilibria: metastability

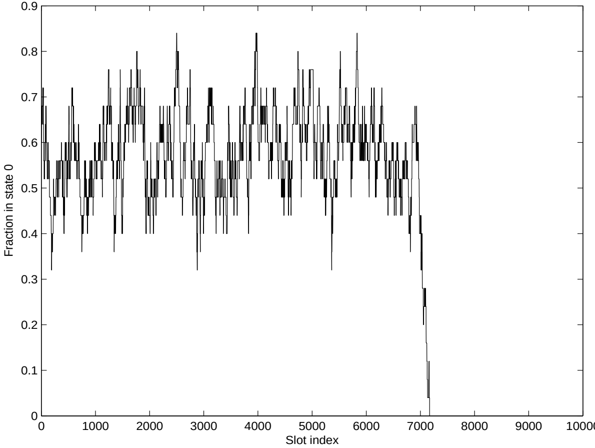

Fraction of nodes in state 0 is near 0.6 for a long time, but then moves to 0, and in a sequence of rapid steps.

The reverse move is a lot less frequent.

Metastability video

Separate file highlighting metastability.

Section 4

Multiple stable limit sets: A selection principle

A selection principle

- If unique globally asymptotically stable equilibrium ξ_f, then $\mu_N(\infty) \stackrel{d}{\to} \xi_f$. (Limit law).

- If we encounter multiple stable limit sets, look at probability of a large deviation.

- Characterise the exponent in

$$\Pr\{\mu_N(\infty) \in \text{ neighbourhood of } q\} \sim \exp\{-NV(q)\}.$$

- The locations $\{q : V(q) = 0\}$ should "select" the correct limit set.

- $V(q)$ is called a quasipotential (Freidlin-Wentzell).

Quasipotential $V(q)$

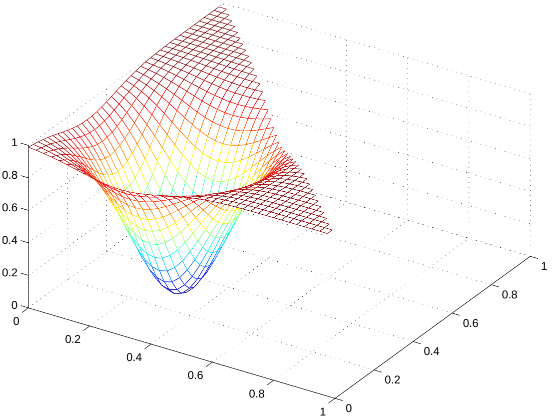

The case of a (unique) globally asymptotically stable equilibrium for the McKean-Vlasov dynamics: $V(\xi_f) = 0$.

Quasipotential $V(q)$

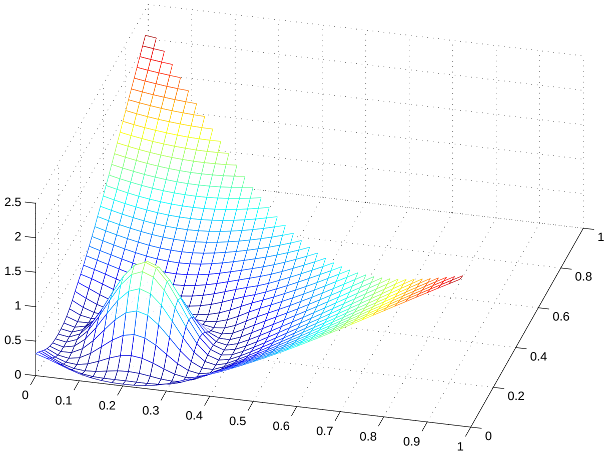

The case of a unique but unstable rest point. $V(\xi_f) > 0$.

All trajectories converge to the stable limit cycle.

Quasipotential $V(q)$

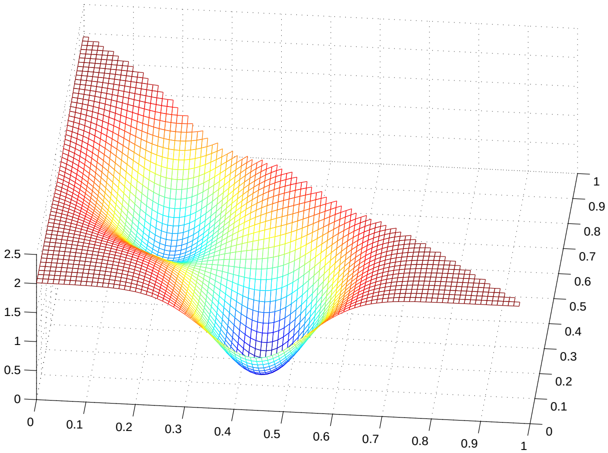

The case of two stable equilibria.

The selection is the one that has the deepest shade of blue ($V(\xi_{f1}) = 0$).

Quasipotential $V(q)$

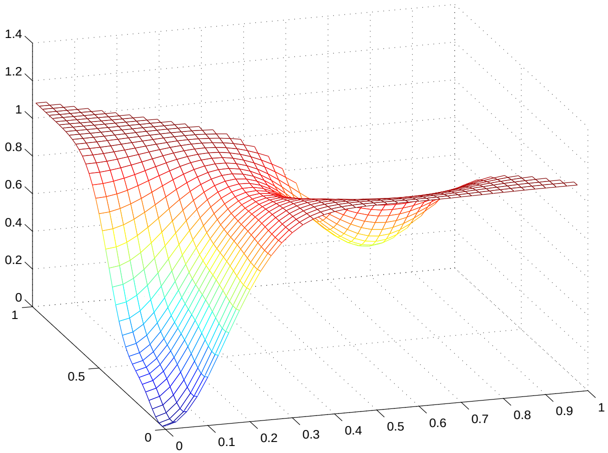

A qualitative picture for the case $c = (0.5, 0.3, 8.0)$.

The two stable points are $(0.6, 0.4, 0.0)$ and $(0.0, 0.0, 1.0)$.
The latter is a truer representative of the large time behaviour.

The quasipotential by inspection

The running cost of following a trajectory $\phi(\cdot)$

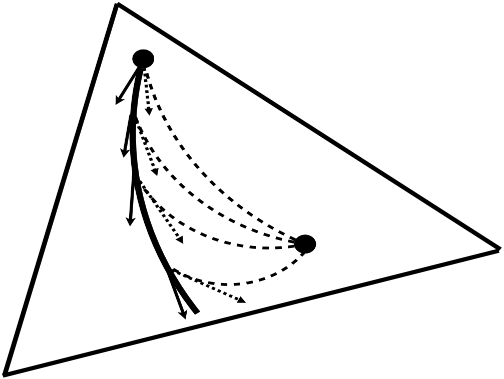

- At each time t, if the current state is $\phi(t)$, the natural tendency is to go along the tangent $\Lambda(\phi(t))^*\phi(t)$.
- To follow $\phi(t)$ however, the system needs to work against the McKean-Vlasov gradient and move along the tangent $\dot{\phi}(t)$.
- $L(\phi(t), \dot{\phi}(t))$.

Guessing the running cost

- Write $\dot{\phi}(t) = G(t)^* \phi(t)$.
- By decoupling, each node's state is iid $\phi(t)$.
- Natural tendency for the $N\phi(t)(i)$ nodes in state i is to have $i \rightsquigarrow j$ at current (instantaneous) rate $\lambda_{i,j}(\phi(t))$.
- But to move along $\phi(t)$ they must have an instantaneous rate of $G_{i,j}(t)$.
- The $N\phi(t)(i)$ Bernoulli($p = \lambda_{i,j}(t)\ dt$) random variables must have a large deviation and must have an empirical measure close to ($q = G_{i,j}(t)\ dt$). By Sanov's theorem, the negative exponent is:

$$N\phi(t)(i)D(q||p) \cong N\phi(t)(i)(q \log \frac{q}{p} - q + p)$$

- Sum over i and j and integrate over $[0, T]$ to get the action functional:

$$\int_0^T L(\phi(t), \dot{\phi}(t))\ dt.$$

The case of a globally asymptotically stable equilibrium ξ_f

Theorem

$V(q)$ is given by

$$V(q) = \inf \left\{ \int_0^T L(\phi(t), \dot{\phi}(t)) \; dt \mid \phi(0) = \xi_f, \phi(T) = q, T \in (0, \infty) \right\}.$$

- Any deviation that puts the system at q must have started its effort from ξ_f.
- $V(\xi_f) = 0$.

The path to q

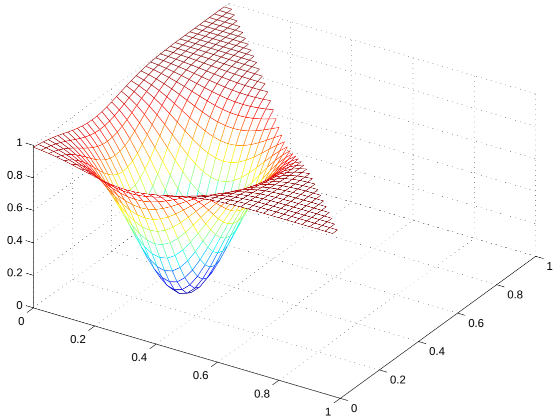

Can specify not only exponent $V(q)$ of the probability, but also the path.

Any deviation that puts the system near q must have started from ξ_f, and must have taken the least cost path.

When there are multiple stable limit sets

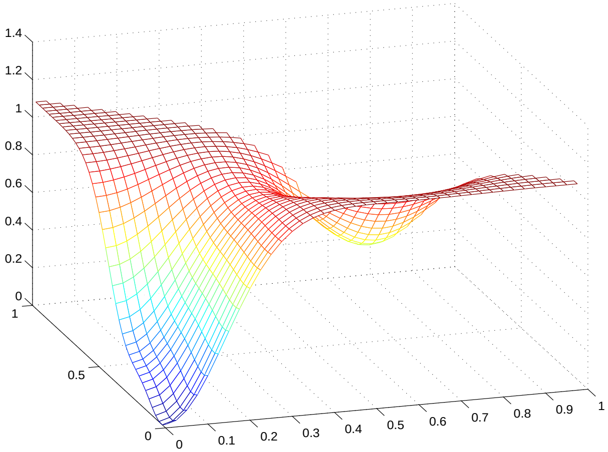

The case of two stable equilibria is easy to describe.

- V_{12} = cost of moving from ξ_{f1} to ξ_{f2}.
- V_{21} = cost of the reverse move.
- If $V_{12} > V_{21}$, then $v_1 = 0$ and $v_2 = V_{12} - V_{21}$.

When there are multiple stable limit sets

Theorem

$V(q)$ is given by

$$V(q) = \inf \left\{ v_i + \int_0^T L(\phi(t), \dot{\phi}(t))\ dt \mid i,\ \phi(0) = \xi_{fi}, \phi(T) = q, T \in (0, \infty) \right\}.$$

- Start from the global minimum ξ_{f1} and move to the attractor in the basin in which q lies along the least cost path.
- Then move to q along the least cost path.

If you are interested in the methodology

- Finite time horizons again, but this time to study large deviation from the McKean-Vlasov limit.
- Large deviation of the stationary measure when there is a globally asymptotic stable equilibrium.
- Analysis of the Markov chain of equilibrium neighbourhoods at hitting times of these neighbourhoods, and associated large deviation principles, when there are multiple stable limit sets.

The take-away picture

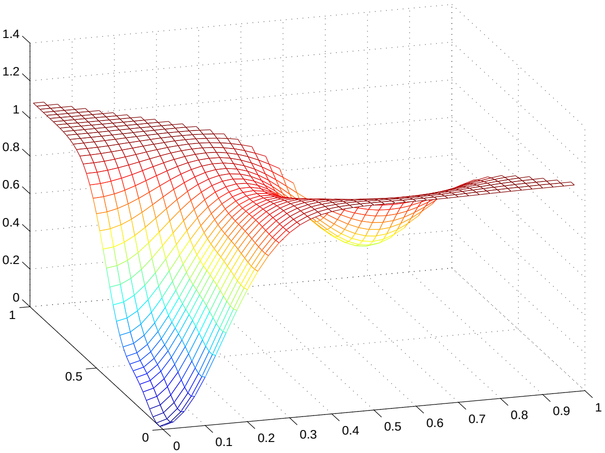

$V_{1\to 2} > V_{2\to 1}$

Background material and acknowledgements

- Today's talk is a synthesis of ideas from:

  Freidlin and Wentzell (1984), Sznitman (1989), Anantharam and Benchekroun (1993), Leonard (1995), IEEE 802.11 standard (1997), Graham and Meleard (1997), Bianchi (1998), Graham (2000), Altman et al. (2006), G. Sharma et al. 2006), Vvedenskaya and Suhov (2007), Benaim and Le Boudec (2008), Ramaiyan et al. (2008), Bordenave et al. (2010), Duffy (2010), Borkar and Sundaresan (2012).

- Thanks to many people for discussions and collaborations. In particular:

  Anurag Kumar, Vivek S. Borkar.