Uniformly Uniformly-ergodic Markov chains and BSDEs

Samuel N. Cohen

Mathematical Institute, University of Oxford

(Based on joint work with Ying Hu, Robert Elliott, Lukas Szpruch)

Centre Henri Lebesgue, Rennes, 22-24 May 2013

Research supported by the Oxford–Man Institute

Outline

1. Uniform Uniform-Ergodicity
2. Markov chain BSDEs
3. Markov chain Ergodic BSDEs
4. Markov chain BSDEs to stopping times

Outline

Markov chains

- Consider a time homogenous finite/countable-state continuous-time Markov chain X.
- X is a process in $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$, where $\{\mathcal{F}_t\}$ is the natural filtration of X.
- X takes values in $\mathcal{X}$, the standard basis of $\mathbb{R}^N$ where $N =$ number of states $\leq \infty$.
- X jumps from state X_{t-} to state $e_i \in \mathcal{X}$ at rate $e_i^* A X_{t-}$, so

$$X_t = X_0 + \int_{]0,t]} A X_{u-} du + M_t$$

  for some $\mathbb{R}^N$-valued martingale M.
- A is a $\mathbb{R}^{N\times N}$ matrix, with $[A]_{ij} \geq 0$ for $i \neq j$, and $\mathbf{1}^* A e_i = \sum_j [A]_{ij} \equiv 0$.

Example

For $N = 4$, we can consider the Markov chain with matrix

$$A = \begin{bmatrix} -3 & 1 & 0 & 0 \\ 1 & -3 & 2 & 2 \\ 2 & 2 & -3 & 1 \\ 0 & 0 & 1 & -3 \end{bmatrix}$$

Which can be drawn as a graph

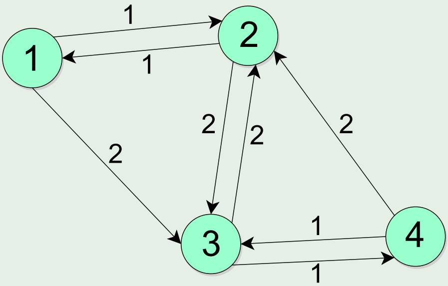

Uniform ergodicity

- If $X_0 \sim \mu$, for some probability measure μ on $\mathcal{X}$, we write $X_t \sim P_t\mu$.
- Writing μ as a vector, $P_t\mu \equiv e^{At}\mu$.
- We say X is uniformly ergodic if

$$\|P_t\mu - \pi\|_{\text{TV}} \leq Re^{-\rho t} \qquad \text{for all } \mu$$

  for some $R, \rho > 0$, some probability measure π on $\mathcal{X}$.
- $\|\cdot\|_{\text{TV}}$ is the total variation norm (and writing signed measures as vectors, $\|\cdot\|_{\text{TV}} = \|\cdot\|_{\ell_1}$).
- The measure π is the *ergodic* measure of the Markov chain and is unique.
- It also satisfies $P_t\pi = \pi$, so it is *stationary*.
- We call (R, ρ) the parameters of ergodicity.
- Irreducible finite state chains are always uniformly ergodic.

Stability of Uniform Ergodicity

- We can tie together the matrix A and the measure $\mathbb{P}$.
- Hence a change of measure corresponds to using a different (stochastic, time varying?) matrix B.
- We say A is uniformly ergodic if X is uniformly ergodic under $\mathbb{P}^A$

Question

Suppose the chain is uniformly ergodic under the measure $\mathbb{P}^A$, and has associated parameters (R_A, ρ_A).

If A and B are 'similar', can we say the chain is uniformly ergodic under $\mathbb{P}^B$, and can we say anything about (R_B, ρ_B)?

- We wish to find a relationship between A and B under which we can get a uniform estimate of ρ_B.

Definition

For $\gamma > 0$, we say 'A γ-controls B' (and write $A \preceq_\gamma B$) if $B - \gamma A$ is also a rate matrix.

If $A \preceq_\gamma B$ and $B \preceq_\gamma A$ we write $A \sim_\gamma B$.

Example

$$A = \begin{bmatrix} -3 & 1 & 0 & 0 \\ 1 & -3 & 2 & 2 \\ 2 & 2 & -3 & 1 \\ 0 & 0 & 1 & -3 \end{bmatrix} \qquad B = \begin{bmatrix} -4 & 1 & 0 & 0 \\ 2 & -3 & 2 & 2 \\ 2 & 2 & -3 & 1 \\ 0 & 0 & 1 & -3 \end{bmatrix}$$

Then for $\gamma \leq 1/2$, $A \sim_\gamma B$.

Note:

- '$\preceq_1$' is a partial order.
- If $A \sim_\gamma B$ for some $\gamma > 0$ then A and B have the same zero entries.

Our key result is:

Theorem

Suppose A is uniformly ergodic and $\gamma \in \,]0,1[\,$. Then there exist constants $\bar{R}, \bar{\rho}$ dependent on γ and A such that B is uniformly ergodic with parameters $(\bar{R}, \bar{\rho})$ whenever $A \preceq_\gamma B$.

The key is that these constants are uniform in B, so these measures are 'uniformly uniformly-ergodic'.

Outline

Markov chain BSDEs

- This result allows us to describe (uniformly) the long-run behaviour of Markov chains under a variety of measures.
- For infinite-horizon BSDEs, we can think of dynamically changing the measure, so such results are useful.
- Before we can make use of this connection, we need the basic theory of BSDEs when noise is driven by a Markov chain.

Martingale Representation

- Recall that

$$X_t = X_0 + \int_{]0,t]} AX_{u-} du + M_t$$

- M is a $\mathbb{R}^N$ martingale with predictable quadratic variation

$$d\langle M \rangle = \phi_t dt = (\mathrm{diag}(AX_{t-}) + AX_{t-}X_{t-}^* + X_{t-}X_{t-}^* A^*)dt$$

- We have $E\Big[\big(\int Z^* dM\big)^2\Big] = E\big[\int Z^* \phi_t Z dt\big] =: E\big[\int \|Z\|_{M_t}^2 dt\big]$
- We can give a martingale representation theorem with respect to M, namely for any scalar martingale L

$$L_t = L_0 + \int_{]0,t]} Z_t^* dM_t$$

and Z is unique up equality in $\|z\|_M^2 = z^* \phi_t z$. Note $\|\mathbf{1}\|_{M_t}^2 \equiv 0$.

BSDE Existence and Uniqueness

Theorem

Let f be a predictable function with

$$|f(\omega,t,y,z)-f(\omega,t,y',z')|^2 \leq c(|y-y'|^2+\|z-z'\|^2_{M_t})$$

for some $c>0$, and $E[\int_{]0,T]}|f(\omega,t,0,0)|^2 dt]<\infty$. Then for any $\xi\in L^2(\mathcal{F}_T)$ the BSDE

$$Y_t=\xi+\int_{]t,T]} f(\omega,t,Y_{u-},Z_u)du+\int_{]t,T]} Z_u^* dM_u$$

has a unique adapted solution (Y,Z) with appropriate integrability.

γ-balanced drivers

To get a comparison theorem, we need the following definition.

Definition

f is γ*-balanced* if there exists a random field λ, with $\lambda(\cdot, \cdot, z, z')$ predictable and $\lambda(\omega, t, \cdot, \cdot)$ Borel measurable, such that

- $f(\omega, t, y, z) - f(\omega, t, y, z') = (z - z')^*(\lambda(\omega, t, z, z') - AX_{t-})$,
- for each $e_i \in \mathcal{X}$,
$$\frac{e_i^* \lambda(\omega, t, z, z')}{e_i^* AX_{t-}} \in [\gamma, \gamma^{-1}]$$
for some $\gamma > 0$, where $0/0 := 1$,
- $\mathbf{1}^* \lambda(\omega, t, z, z') \equiv 0$ and
- $\lambda(\omega, t, z + \alpha\mathbf{1}, z') = \lambda(\omega, t, z, z')$ for all $\alpha \in \mathbb{R}$.

Comparison theorem

Lemma

If f is γ-balanced for $\gamma > 0$, then it is uniformly Lipschitz in z with respect to $\|\cdot\|_{M_t}$.

Lemma

If $f(u;\cdots)$ is γ-balanced for each u, then $g(\cdots) := \operatorname{ess\,sup}_u\{f(u;\cdots)\}$ is γ-balanced (given it is always finite).

Comparison theorem

Lemma

If $B \sim_\gamma A$ then $f(\omega, t, y, z) = z^*(B - A)X_{t-}$ is γ-balanced with $\lambda(\cdots) = BX_{t-}$, and $Y_t = E^B[\xi|\mathcal{F}_t]$.

The purpose of this definition is to obtain:

Theorem (Comparison theorem)

Let f be γ-balanced for some $\gamma > 0$. Then if

$$\xi \geq \xi' \text{ and } f(\omega, t, y, z) \geq f'(\omega, t, y, z),$$

the associated BSDE solutions satisfy $Y_t \geq Y'_t$ for all t, and $Y_t = Y'_t$ on $A \in \mathcal{F}_t$ iff $Y_s = Y'_s$ on A for all $s \geq t$.

Comparison theorem

The comparison theorem is easy to derive from

Lemma

Let f be γ-balanced. Then there exists a measure $\mathbb{Q} \sim \mathbb{P}$ such that

$$Y_t - Y'_t + \int_{]0,t]} \big(f(\omega, s, Y_{s-}, Z_s) - f'(\omega, s, Y'_{s-}, Z_s)\big) ds$$

is a $\mathbb{Q}$-martingale.

Proof.

$\mathbb{Q}$ is the measure where $\lambda(\omega, t, Z_t, Z'_t)$ is the compensator of X. $\square$

Markovian Solutions

We can also obtain a Feynman–Kac type result.

Theorem

Let $f : \mathcal{X} \times [0, T] \times \mathbb{R} \times \mathbb{R}^n$*, and consider the BSDE with*

$$Y_t = \phi(X_T) + \int_{]t,T]} f(X_{t-}, t, Y_{u-}, Z_u)du + \int_{]t,T]} Z_u^* dM_u$$

for some function $\phi : \mathcal{X} \to \mathbb{R}$*. Then the solution satisfies*

$$Y_t = u(t, X_t) = X_t^* \mathbf{u}_t, \qquad Z_t = \mathbf{u}_t$$

for $\mathbf{u}$ *solving*

$$d\mathbf{u}_t = -(\mathbf{f}(t, \mathbf{u}_t) + A^* \mathbf{u}_t)dt; \qquad e_i^* \mathbf{u}_T = \phi(e_i)$$

where $e_i^* \mathbf{f}(t, \mathbf{u}_t) := f(e_i, t, e_i^* \mathbf{u}_t, \mathbf{u}_t)$.

Application to Control

Consider the problem

$$\min_{\mathfrak{u}} E^{\mathfrak{u}}\left[e^{-rT}\phi(X_T) + \int_{]0,T]} e^{-rt} L(\mathfrak{u}_t; X_{t-}, t)dt\right]$$

where $\mathfrak{u}$ is a predictable process with values in U, and $E^{\mathfrak{u}}$ is the expectation under which X has compensator $\lambda_t(\mathfrak{u}_t) \in \mathbb{R}^N$.

- Suppose $\lambda_t(\cdot)$ satisfies $\frac{e_i^* \lambda_t(\cdot)}{e_i^* A X_{t-}} \in [\gamma, \gamma^{-1}]$ for some $\gamma > 0$.
- Define the Hamiltonian

$$f(\omega, t, y, z) = -rY_t + \inf_{u \in U}\{L(u; X_{t-}, t) + z^*(\lambda_t(u) - AX_{t-}(\omega))\}$$

- The dynamic value function Y_t satisfies the BSDE with driver f, terminal value $\phi(X_t)$.
- By the comparison theorem, we have existence and uniqueness of an optimal feedback control.

Outline

Ergodic BSDEs

- Hu, Fuhrman, Tessitore and collaborators have introduced ‘ergodic BSDEs’.
- These are known to relate to ergodic control problems.
- In a Brownian setting, existence/uniqueness of these equations depends on analysis of uniform ellipticity estimates of perturbations of the generator of the underlying X.
- These methods do not directly map to the Markov Chain case.
- Instead we will use the uniform ergodicity estimates we have derived.

Infinite-Horizon BSDEs

The first step is to establish existence of Infinite Horizon BSDEs

Theorem

Let $\alpha > 0$, and consider the equation

$$dY_t = -(f(\omega, t, Z_t) - \alpha Y_t)dt + Z_t^* dM_t \qquad t \in [0, \infty[.$$

If $|f(\omega, t, z|) < C$, this admits a solution (Y, Z) with $Y < C/\alpha$ for $t \in [0, \infty[$, and this is the unique bounded solution.

Proof.

Let Y_t^T solve the BSDE

$$dY_t = -(f(\omega, t, Z_t) - \alpha Y_t)dt + Z_t^* dM_t; \qquad Y_T^T = 0.$$

Then show $Y_t^T \leq C/\alpha$ for all T and $Y^T \to Y$ (uniformly on compacts). □

EBSDE Existence/Uniqueness

Theorem (SC & Hu 2012)

The ergodic BSDE

$$dY_t = -(f(X_t, Z_t) - \lambda)dt + Z_t^* dM_t$$

admits a bounded solution (Y, Z, λ) *with* $Y_t = u(X_t)$ *and* $u(\hat{x}) = 0$, *whenever* f *is* γ*-balanced with* $|f(x, 0)| < C$.

Any other bounded solution has the same λ, *any other bounded Markovian solution has* $Y'_t = Y_t + c$.

Our Feynman–Kac type result gives, with $Y = u(X_t) = X_t^*\mathbf{u}$, $Z = \mathbf{u}$,

$$\mathbf{f}(\mathbf{u}) - \lambda = -A^*\mathbf{u}$$

which is readily calculable for small dimensions.

Proof.

Denote the infinite horizon BSDE solution $Y_t^\alpha = u^\alpha(X_t)$. Then for some rate matrices $B^T \sim_\gamma A$,

$$u^\alpha(x) = \lim_{T\to\infty} E^{B^T}\Big[\int_{]0,T]} e^{-\alpha u} f(X_{u-},0)du | X_0 = x\Big]$$

So, with $(\bar{R},\bar{\rho})$ the uniform convergence bounds on $P_u^B\mu$,

$$\begin{aligned}|u^\alpha(x) - u^\alpha(x')| &\leq C \lim_{T\to\infty}\int_{]0,T]} e^{-\alpha u}\|P_u^{B^T}\delta_x - P_u^{B^T}\delta_{x'}\|_{TV}du \\ &\leq C\bar{R}(\alpha+\bar{\rho})^{-1} \leq C\bar{R}/\alpha < \infty\end{aligned}$$

A diagonal argument implies $u^\alpha(x) - u^\alpha(\hat{x}) \to u(x)$ and $\alpha u^\alpha(\hat{x}) \to \lambda$, as $\alpha \downarrow 0$. □

Properties

The comparison theorem does not hold for Y for EBSDEs. However, it does hold for λ.

Theorem

Let f, f' be γ-balanced and $f \geq f'$. Then we have $\lambda \geq \lambda'$ in the EBSDE solutions.

Theorem

Let π^A denote the ergodic measure when X has matrix A. Then for some $B^{\mathbf{u}} \sim_\gamma A$,

$$\lambda = \int_{\mathcal{X}} f(x, \mathbf{u}) d\pi^A(x) = \int_{\mathcal{X}} f(x, 0) d\pi^{B^{\mathbf{u}}}(x)$$

$$Y_t + c = \int_{\mathcal{X}} f(x, \mathbf{u}) \mu^A_{X_0}(x) = \int_{\mathcal{X}} f(x, 0) \mu^{B^{\mathbf{u}}}_{X_0}(x)$$

where $\mu^A_x = \int_{\mathbb{R}+} (P^A_u \delta_x - \pi^A) du$.

Applications

- Application to Ergodic control similar to the finite-horizon case, with value

$$\lambda = \min_u \left\{ \limsup_{T\to\infty} E^u \left[\frac{1}{T} \int_{]0,T]} L(u_t; X_{t-}) dt \right] \right\}$$

- Write f as the Hamiltonian, comparison theorem gives optimal feedback control, etc...

- Instead, we will look at creating spatially stable nonlinear probabilities on graphs.
- These attempt to generalize the ergodic distribution of a Markov chain to a nonlinear setting.

Applications

- Consider the driver $f(x,z) = I_{\{x\in\Xi\}} + g(x,z)$. Then if $g \equiv 0$, the EBSDE solutions will be $\lambda = \pi(\Xi)$.
- If g is convex, we will have a 'convex ergodic probability',
- If $g(x,z) = \sup_{B\in\mathcal{B}}\{z^*(B-A)x\}$, when $\mathcal{B}$ has 'column-exchangeability', then we can show that
$$\lambda^{\Xi} = \sup_B \pi^B(\Xi)$$
- The global balance equation $A\pi \equiv 0$ becomes the Isaac's condition
$$\sup_B \left\{ \int_{\mathcal{X}} \inf_{B'}\{v^* B' x\} d\pi^B(x) \right\} \equiv 0 \qquad \text{for all } v \in \mathbb{R}^N.$$

Nonlinear Pagerank

- The Pagerank algorithm is modelled on a Markov chain randomly following links in a graph (e.g. the WWW) at a constant rate.
- Nodes are ranked according to the ergodic probability the chain is at each node.
- What happens if we instead dynamically take the minimum over a range of rate matrices?

Pagerank example

Consider this undirected graph.

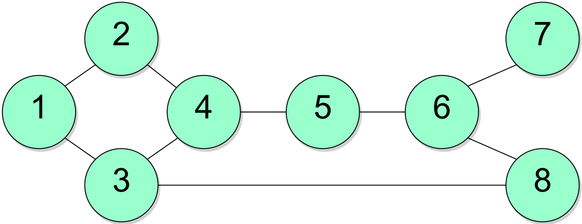

Our Markov chain moves between states randomly, at a constant (exit) rate.

Pagerank example

The ergodic probabilities/Pagerank under this basic model can be simply calculated.

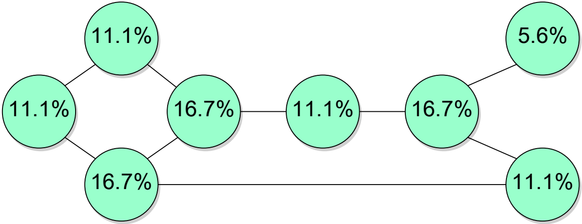

We have three states which are equally highest ranked, even though they have different positions in the graph.

Now suppose a bias can be introduced dynamically, so that the relative rate of jumping to the right/left can be doubled. Then the EBSDE can be solved with drivers

$$f(x,z) = I_{\{x\in\Xi\}} + \min_B\{z^*(B-A)x\}$$

for each $\Xi \subseteq \mathcal{X}$. The corresponding values λ^Ξ (for Ξ singletons) are

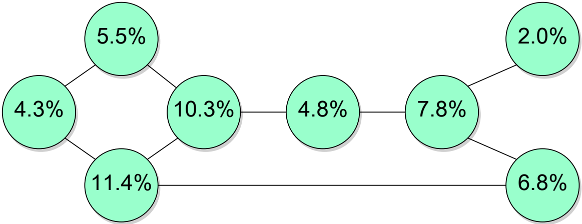

Pagerank example

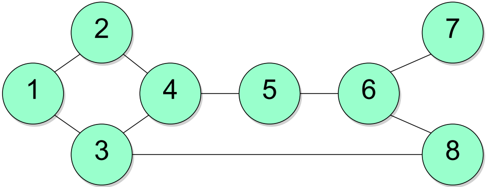

- We can also assess collections of states, so if
  - $\Xi = \{1, 3, 4\}$, then $\lambda^\Xi = 0.3067$
  - $\Xi = \{2, 3, 4\}$, then $\lambda^\Xi = 0.3164$
  - $\Xi = \{3, 4, 5\}$, then $\lambda^\Xi = 0.3531$
  - $\Xi = \{4, 5, 6\}$, then $\lambda^\Xi = 0.2899$.
  - Note all these collections have ergodic measure 0.444 under the basic model.
  - Peculiarly, $\lambda^{\{2\}} = 0.055 > 0.048 = \lambda^{\{5\}}$, but $\lambda^{\{2,3,4\}} < \lambda^{\{3,4,5\}}$.
- In this way we can assess the 'centrality' of a collection of states.

Outline

BSDEs to stopping times

- Another interesting class of BSDEs is where the terminal time is a stopping time.
- Here we want

$$Y_t = \xi + \int_{]t,\tau]} f(\omega, u, Y_{u-}, Z_u)du + \int_{]t,\tau]} Z_u^* dM_u$$

when τ is a stopping time, ξ is $\mathcal{F}_\tau$-measurable.

- If

$$\frac{f(\omega,t,y,z) - f(\omega,t,y',z)}{y - y'} \leq -\alpha \quad \text{for some } \alpha > 0,$$

then we can use the discounted methods considered earlier.

- This does not cover the interesting case $f(x,z) = z^*(B-A)x$.

BSDEs to stopping times

- We will assume sufficient bounds on ξ and τ that our BSDE admits a unique solution with polynomial growth.
- This is similar to the backwards heat equation admitting only one polynomial growth solution.
- Of use will be the family $\mathcal{Q}_\gamma$ of measures where X has compensator λ with
$$\frac{e_i^* \lambda(\omega, t)}{e_i^* A X_{t-}} \in [\gamma, \gamma^{-1}] \quad \text{for all } i.$$
- These measures are singular on $\mathcal{F}_\infty$ but equivalent on $\mathcal{F}_\tau$.

BSDEs to stopping times

Imagine that we can prove the following bounds.

Assumption

For some nondecreasing $K, \tilde{K} : \mathbb{R}^+ \to [1, \infty[$*, constants* $\beta, \tilde{\beta} > 0$*,*

$$E^{\mathbb{Q}}[|\xi||\mathcal{F}_t] \leq K(t),$$
$$E^{\mathbb{Q}}[(1+\tau)^{1+\beta}|\mathcal{F}_t] \leq K(t),$$
$$E^{\mathbb{Q}}[K(\tau)^{1+\tilde{\beta}}|\mathcal{F}_t] \leq \tilde{K}(t),$$

all $\mathbb{P}$*-a.s. for all* $\mathbb{Q} \in \mathcal{Q}_\gamma$ *and all* t*.*

BSDEs to stopping times

Theorem

Let ξ, τ satisfy the previous bounds, and suppose f is γ-balanced, and such that for any y, y', z,

$$|f(\omega, t, 0, 0)| \leq c(1 + t^{\hat{\beta}})$$

$$\frac{f(\omega, t, y, z) - f(\omega, t, y', z)}{y - y'} \in [-c, 0]$$

for some $c \in \mathbb{R}$, some $\hat{\beta} \in [0, \beta[$. Then the BSDE

$$Y_t = \xi + \int_{]t,\tau]} f(\omega, u, Y_{u-}, Z_u)du + \int_{]t,\tau]} Z_u^* dM_u$$

admits a unique adapted solution satisfying the bound

$$|Y_t| \leq (1 + c)K(t).$$

BSDEs to stopping times

Proof.

Approximate with finite horizon BSDE

$$Y_t = I_{T \leq \tau}\xi + \int_{]t,\tau\wedge T]} f(\omega, u, Y_{u-}, Z_u)du + \int_{]t,\tau\wedge T]} Z_u^* dM_u$$

and let $T \to \infty$, using Hölder and Markov inequalities to prove convergence on compacts. □

We can also prove the comparison theorem, which implies...

Theorem

If the conditions on f hold only locally, and $\xi, f(\omega, t, 0, 0)$ are bounded, by the comparison theorem everything works given conditions on τ.

This isn't too interesting if we can't verify the assumptions on τ.

Hitting times work!

Theorem

If τ is the first hitting time of a set $\Xi \subseteq \mathcal{X}$, and $\xi = g(\tau, X_\tau)$ for some $g(t,x) \leq c(1+t^{\hat{\beta}})$, then the assumptions are satisfied.

Proof.

As we have uniform ergodicity estimates, we can show that τ admits exponential moments uniformly bounded under all Markovian measures in $\mathcal{Q}_\gamma$. By applying the comparison theorem to finite-time approximations, we can extend this to all measures in $\mathcal{Q}_\gamma$. □

By extension, this also holds for nth hitting times, etc...

Markovian Structure

We can also give a Feynman–Kac type result

Theorem

Let τ be the first hitting time of a set $\Xi \subseteq \mathcal{X}$ and let f satisfy the conditions. Consider the BSDE

$$Y_t = \phi(X_\tau) - \int_{]t,\tau]} f(X_{u-}, Y_{u-}, Z_u)du + \int_{]t,\tau]} Z_u dM_u$$

where ϕ is a bounded function $\mathcal{X} \to \mathbb{R}$. Then there exists a bounded vector $\mathbf{u} \in \mathbb{R}^N$ such that $Y_t = X_t^* \mathbf{u}$, $Z_t \equiv \mathbf{u}$ and

$$f(x, x^*\mathbf{u}, \mathbf{u}) = -\mathbf{u}^* A x \qquad \text{for } x \in \mathcal{X} \setminus \Xi,$$

with boundary values $x^*\mathbf{u} = \phi(x)$ for $x \in \Xi$.

Applications: Network Reliability

- Consider a model for transmission of messages over our graph.

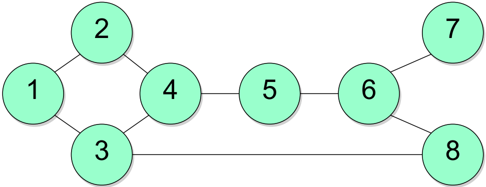

- Messages are passed randomly between nodes, we wish to assess whether a message created at node 1 will reach node 8 before being killed at the faulty node 5.
- This probability corresponds to a terminal value $\phi(x) = I_{\{x=e_8\}}$ at a time $\tau = \inf\{t : X_t = e_5 \text{ or } e_8\}$.

Applications: Network Reliability

- Without control, the probability is given by

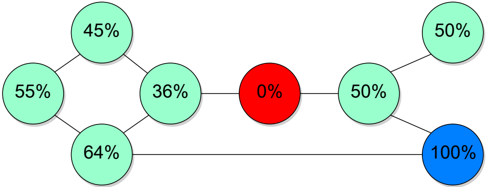

- With our control, we have

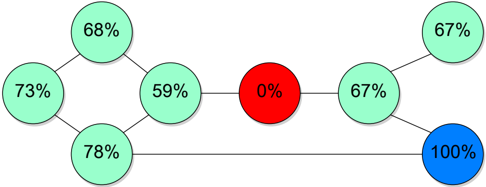

Applications: Non-Ohmic Circuits

- Markov chains can be used to model electronic circuits of resistors.
- If $w_{ij} = 1/R_{ij}$ is the resistance of a connection, then the voltage potential v satisfies
$$\begin{cases} v(i) = \phi(i), & i \in \Xi \\ v(i) \sum_j w_{ij} = \sum_j w_{ij} v(j), & i \notin \Xi \end{cases}$$
for Ξ a source set.
- This depends on Ohm's Law and Kirtchoff's current laws.
- For finite circuits, the voltage potential is the expected value at the first hitting time of Ξ, for a Markov chain with $A_{ij} = w_{ji}$ for $i \neq j$.

Applications: Non-Ohmic Circuits

- Diodes, Transistors, etc. do not satisfy Ohm's law.
- For a diode, the effective resistance is

$$R_{ij}^{v} = \frac{v(j) - v(i)}{c_i(\exp([v(j) - v(i)]/c_v) - 1)} > 0$$

- Write B^v for the rate matrix with entries $B_{ij}^v = 1/R_{ji}^v$.
- We can write the voltage potential as solving the BSDE to the first hitting time of Ξ with driver

$$f(x, z) = z^*(B^z - A)x$$

  where A is generated from a reference resistor circuit.
- One can show the required conditions are (locally) satisfied, so for any finite source potential we have unique solutions.

Applications: Non-Ohmic Circuits: Example

Consider an AND logic gate.

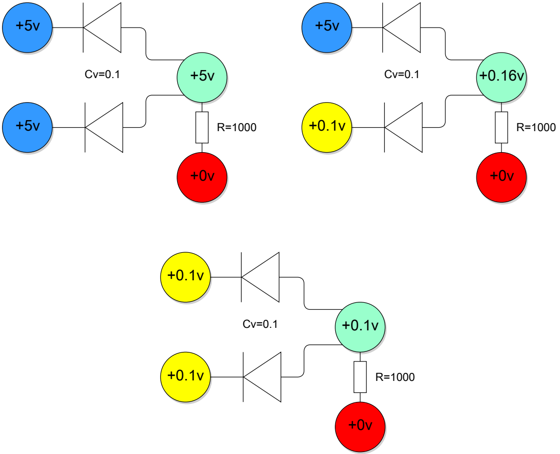

Connection to EBSDE

The following theorem gives a relation between Ergodic BSDEs and BSDEs to stopping times.

Theorem

Let $\hat{f}(x,z) = f(x,z) - \int_{\mathcal{X}} f(x,z) d\pi^A(x)$. Then the BSDE to $\tau = \inf\{t : X_t = \hat{x}\}$ with driver $\hat{f}$ and terminal value $\xi = 0$ agrees with the EBSDE with driver f and $u(\hat{x}) = 0$.

Hence the corresponding BSDEs to $\tau \wedge T$ converge uniformly to the EBSDE solution.

This suggests an alternative Monte-Carlo approach to numerical computation of EBSDE solutions, which is of importance for large/infinite-state chains.

Conclusions

- We have seen that one can obtain uniform uniform-ergodicity estimates for Markov chains under perturbations of the rate matrix.
- With these, one can study solutions to Ergodic BSDEs and BSDEs to hitting times
- Solutions are readily calculable for small examples, and can be used to model various problems.
- Open problems include removing the γ-balanced requirement, including y in the driver of EBSDEs, better numerics...