Lyapunov Functions and the Analysis of Queues

Peter W. Glynn

Stanford University

INI Satellite Workshop, ICMS, Edinburgh, UK
June 11, 2010

Outline of Talk

1. First Transition Analysis for Markov Chains
2. Operator-theoretic Bounds
3. Lyapunov Bounds
4. A G/G/1 Application
5. Bounds on Equilibrium Moments
6. Application to Importance Sampling
7. Connection to Supermartingales
8. Further Applications

First Transition Analysis

Example 1: (Expected Reward to a Hitting Time)

$$u^*(x) = \mathbb{E}_x \sum_{j=0}^{T-1} r(X_j) \qquad (r : S \to \mathbb{R}_+)$$

$$T = \inf\{n \geq 0 : X_n \in C^c\}$$

Then, u^* satisfies

$$u = r + Pu$$
$$\text{s/t} \quad u = 0 \quad \text{on } C^c$$

Equivalently, u^* on C satisfies

$$u = r + Bu$$

where $B = (B(x, \mathrm{d}y) : x, y \in C^c)$ and $B(x, \mathrm{d}y) = P(x, \mathrm{d}y)$

This linear system generally has multiple solutions.

But

$$u^*(x) = \sum_{n=0}^{\infty} \mathbb{E}_x r(X_j) I(T > j)$$
$$= \sum_{n=0}^{\infty} \int_C B^n(x, \mathrm{d}y) r(y)$$
$$= \sum_{n=0}^{\infty} (B^n r)(x)$$

i.e. the probabilistically meaningful solution is

$$u^* = \sum_{n=0}^{\infty} B^n r$$

Example 2: (Expected Infinite Horizon Reward)

$$u^*(x) = \mathbb{E}_x \sum_{j=0}^{\infty} e^{-\alpha j} r(X_j) \qquad (\alpha > 0, r : S \to \mathbb{R}_+)$$

Then, u^* satisfies

$$u = r + e^{-\alpha} P u$$

or, equivalently,

$$u = r + Bu$$

where $B = e^{-\alpha} P$

Also, the probabilistically meaningful solution is

$$u^* = \sum_{n=0}^{\infty} B^n r$$

Bounding Solutions to Linear Systems

$$u = r + Bu \qquad (r, B \geq 0)$$

$$u^* = \sum_{n=0}^{\infty} B^n r$$

How can we bound u^*?

Operator-theoretic Bounds

Given $w : S \to [1, \infty)$, let $\|f\|_w = \sup\{\frac{|f(x)|}{w(x)} : x \in C\}$ and $L_w = \{f : S \to \mathbb{R} : \|f\|_w < \infty\}$. Then, we can define the operator norm of B as

$$|||B|||_w = \sup_{f \in L_w, \|f\|_w > 0} \frac{\|Bf\|_w}{\|f\|_w}$$

If $|||B|||_w < 1$, then

$$\|u^*\|_w \leq \sum_{n=0}^{\infty} (|||B|||_w)^n \|r\|_w = \frac{\|r\|_w}{1 - |||B|||_w}$$

Note that $|||B|||_w < 1$ translates to

$$\mathbb{E}_x w(X_1) I(X_1 \in C) \leq \beta w(x) \quad \text{in } C$$

where $\beta < 1$.

Lyapunov Bounds

Assume there exists a finite-valued $g \geq 0$ such that $Bg \leq g - r$. Then:

$$u^* \leq g.$$

The proof...

Note that: Bg is finite-valued

$$B^n g \text{ is finite-valued}$$

So,

$$r \leq g - Bg$$

$$B^k r \leq B^k g - B^{k+1} g$$

Hence,

$$\sum_{k=0}^{n} B^k r \leq g - B^{n+1} g \leq g$$

So,

$$u^* \leq g$$

How to guess the Lyapunov function?

Note that u^* itself satisfies the Lyapunov inequality

We don't know u^* but we often can guess u^*'s asymptotic behavior,

i.e.

$$u^* \approx v$$

Guess:

$$g = v$$

To absorb approximation inaccuracies, try

$$g = av \qquad (a > 0)$$

Apply to Single-server Queue

Consider the delay sequence (aka the "waiting time sequence") for the FIFO G/G/1 queue:

$$W_{n+1} = [W_n + Z_{n+1}]^+$$

where $Z_1, Z_2, \ldots$ are iid and given by

$$Z_{n+1} = V_n - \chi_{n+1}.$$

When $W_n = x$ is large,

$$W_{n+1} \approx W_n + Z_{n+1}$$

i.e. the dynamics are approximately that of a random walk

For stability, we expect that we will need $\mathbb{E}Z_n < 0$ for stability

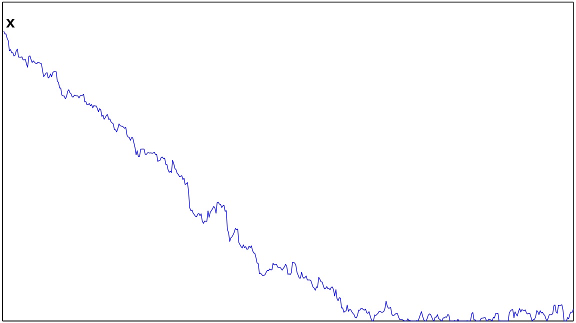

Given a Markov chain $X = (X_n : n \geq 0)$ living on S, the key to proving stability is:

Prove that the expected return time to compact sets is finite

Step 1: Bound the expected time to hitting K

$$u^* = e + Bu^* \qquad (u^*(x) = \mathbb{E}_x T_K; \quad e(x) = 1 \text{ on } K^c)$$

Lyapunov bound:

$$Bg \leq g - e$$

Simpler condition:

$$Pg \leq g - e \text{ on } K^c$$

To construct g, approximate u^* via "fluid analysis":

$$u^*(x) \approx v(x) = \frac{x}{|\mathbb{E}Z_1|} \text{ for } x \text{ large}$$

Try $g(x) = ax$ $(a > 0)$.

Note that

$$(Pg)(x) - g(x) = a\mathbb{E}([x + Z_1]^+ - x) \to a\mathbb{E}Z_1 \text{ as } x \to \infty$$

Put $a = 2/|\mathbb{E}Z_1|$, so that g satisfies the Lyapunov inequality for $x \geq x_0$.
Finally, put $K = [0, x_0]$.
Step 2: Bound the expected return time to K

$$\begin{aligned} \mathbb{E}_x\tau_K &= 1 + \mathbb{E}_x u^*(W_1) I(X_1 \in K^c), \qquad x \in K \\ &\leq 1 + \mathbb{E}_x g(W_1) \end{aligned}$$

So, if Z_1 is integrable and $\mathbb{E}Z_1 < 0$, then

$$\sup \mathbb{E}_x\tau_K \leq 1 + a \sup \mathbb{E}[x + Z_1]^+ \leq 1 + a(x_0 + \mathbb{E}|Z_1|)$$

Conclusion: $W = (W_n : n \geq 0)$ is a positive recurrent Markov chain if Z_1 is integrable and $\mathbb{E}Z_1 < 0$.

Relation to Operator-based Bounds

Recall that $|||B|||_w < 1$ is equivalent to requiring:

$$Bw \leq w - (1 - \beta)w$$

so this is a special type of Lyapunov bound with $r = cw$ $(c > 0)$

When $|||B|||_w < 1$, this implies (in great generality) that $X = (X_n : n \geq 0)$ is geometrically ergodic

For $W = (W_n : n \geq 0)$, geometric ergodicity demands that V have exponential moments.

Moral of the Story: To get the best possible conditions for stability, one can not depend on operator methods...one needs the additional flexibility offered via more general Lyapunov ideas.

Stability for Networks

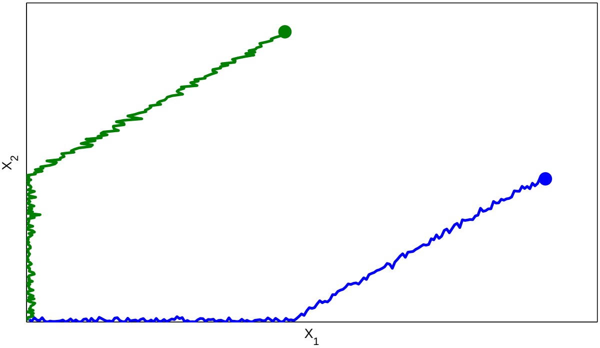

For large x,

$$\mathbb{E}_x T_K \approx v(x)$$

where $v(\cdot)$ is a piecewise linear function that describes the fluid dynamics of the network

Equilibrium Moment Estimates

Given a Markov chain $X = (X_n : n \geq 0)$, living on S, having an equilibrium distribution π and $f : S \to \mathbb{R}_+$, how do we bound πf?

$$\pi f = \mathbb{E}_{\pi_K} \sum_{n=0}^{\tau_K - 1} \frac{f(X_n)}{\mathbb{E}_{\pi_K} \tau_K} \leq \mathbb{E}_{\pi_K} \sum_{n=0}^{\tau_K - 1} f(X_n) \leq \sup_{x \in K} \mathbb{E}_x u^*(X_1)$$

where $u^*(x) = \mathbb{E}_x \sum_{n=0}^{\tau_K - 1} f(X_n)$

Lyapunov bound:

Find $g \geq 0$ s.t.

$$Pg \leq g - f \text{ for } x \in K^c$$

Then:

$$u^* \leq g \text{ on } K^c$$

Back to single-server queue with $\mathbb{E}Z_1 < 0$:

Suppose $f(x) = x^p$ $(p > 0)$

For x large:

$$u^*(x) \approx \sum_{n=0}^{\lfloor x/|\mathbb{E}Z_1| \rfloor} (x + n\mathbb{E}Z_1)^p \approx v(x)$$

where $v(x) = cx^{p+1}$.

So, choose $g(x) = ax^{p+1}$.

In order to use this choice of g, we need: $\mathbb{E}|Z_1|^{p+1} < \infty$.
(Best possible moment condition...)

A more efficient means of obtaining bounds on equilibrium moments:

Our Lyapunov bound requires:

$$Pg \le g - f \text{ on } K^c$$
$$Pg \le c \text{ on } K$$

Combine into one bound:

$$Pg \le g - f + \tilde{c}e \text{ on } S$$

Apply π to both sides:

$$\pi f \le \tilde{c}$$

This argument presumes that $\pi g < \infty$ (and hence suggests the need to use a second Lyapunov function to bound πg)

But a more careful argument shows such a second Lyapunov function is unnecessary.

Theorem

Let π be an equilibrium distribution of X. If $f : S \to \mathbb{R}_+$ and g is a non-negative function satisfying

$$Pg \leq g - f + ce$$

for some $c > 0$, then $\pi f \leq c$.

Lyapunov Functions and Importance Sampling

Setting: $X = (X_n : n \geq 0)$ Markov chain living on S

$$T = \inf\{n \geq 0 : X_n \in A\}$$

Compute $u^*(x) = \mathbb{P}_x(X_T \in B, T < \infty)$ $(B \subseteq A)$

Example: G/G/1 queue

$$\mathbb{P}_x(T < \infty) = ?$$

where $T = \inf\{n \geq 0 : S_n \geq 0\}$ and $x < 0$.

Simulate X under modified dynamics:

$$\tilde{P}(x, \mathrm{d}y) = P(x, \mathrm{d}y) r(x, y)^{-1}$$

Then:

$$u^*(x) = \tilde{\mathbb{E}}_x I(X_T \in B, T < \infty) \prod_{n=0}^{T-1} r(X_n, X_{n+1})$$

What is the variance of this estimator?

Bounding the Second Moment of IS Estimators

Put $$s^*(x) = \tilde{\mathbb{E}}_x I(X_T \in A, T < \infty) \prod_{n=0}^{T-1} r^2(X_n, X_{n+1})$$
$$= \mathbb{E}_x I(X_T \in A, T < \infty) \prod_{n=0}^{T-1} r(X_n, X_{n+1})$$

Then, s^* satisfies

$$s = r + Bs \text{ on } A^c$$

where $B(x, \mathrm{d}y) = P(x, \mathrm{d}y) r(x, y)$ and $r(x) = \mathbb{E}_x I(X_1 \in A) r(x, X_1)$

Lyapunov bound:

Find $g \geq 0$ s.t.

$$Bg \leq g - r \text{ on } A^c$$

Further simplification: The zero-variance change-of-measure is given by $P^*(x, \mathrm{d}y) = P(x, \mathrm{d}y)u^*(y)/u^*(x)$

Try $\tilde{P}$ of the form $\tilde{P}(x, \mathrm{d}y) \propto P(x, \mathrm{d}y)v(y)$, where $v \approx u^*$ (e.g. $u^*(x) \approx e^{\theta^* x}$ as $x \to \infty$ for light-tailed G/G/1)

$$\text{i.e.} \quad \tilde{P}(x, \mathrm{d}y) = P(x, \mathrm{d}y)\frac{v(y)}{w(x)}$$

For this class of $\tilde{P}$'s, we hope that $s^*(x)$ is close to $u^*(x)^2$. So, choose $g(x) = av(x)^2$ (or, more generally, write $g(x) = a(x)v(x)^2$ where we expect $a(\cdot)$ to be bounded above and below)

A Martingale Perspective on Lyapunov Functions

Consider problem of bounding

$$u^*(x) = \mathbb{E}_x \sum_{n=0}^{T-1} r(X_n)$$

Assuming that $u^*(\cdot)$ is finite-valued,

$$M_i = \mathbb{E}_x[\sum_{n=0}^{T-1} r(X_n)|X_0, \ldots, X_i]$$

must be a martingale adapted to $(\mathscr{F}_i : i \geq 0)$. But

$$M_i = \sum_{n=0}^{T\wedge i-1} r(X_n) + u^*(X_{T\wedge i}) \tag{1}$$

In order that this sequence be a MG, we need $\mathbb{E}[M_i - M_{i-1}|\mathscr{F}_{i-1}] = 0$:
On $\{T > i-1\}$,

$$u^*(X_{i-1}) = r(X_{i-1}) + (Pu^*)(X_{i-1}) \quad \text{a.s.}$$

If we replace u^* in (1) by a function $g \geq 0$ for which

$$g(X_{i-1}) \geq r(X_{i-1}) + (Pg)(X_{i-1}) \quad \text{a.s.}$$

then $(M_n : n \geq 0)$ becomes a non-negative supermartingale

$$\mathbb{E}_x M_i \leq \mathbb{E}_x M_0 = g(x)$$

i.e.

$$\mathbb{E}_x \sum_{n=0}^{T \wedge i-1} r(X_n) \leq g(x)$$

Send $i \to \infty$

$$u^*(x) \leq g(x)$$

Supermartingale Bounds in Continuous Time

$$u^*(x) = \mathbb{E}_x \int_0^T r(X(s))\,\mathrm{d}s$$

where X is Markov. Then, assuming finiteness of $u^*(\cdot)$,

$$M(t) = \int_0^{T\wedge t} r(X(s))\,\mathrm{d}s + u^*(X(T\wedge t))$$

must be a martingale adapted to $(\mathscr{F}_t : t \geq 0)$. In order to upper bound u^*, find a $g \geq 0$ such that

$$M(t) = \int_0^{T\wedge t} r(X(s))\,\mathrm{d}s + g(X(T\wedge t))$$

for which $M(\cdot)$ is a supermartingale.

e.g. one dimensional RBM

$$\mathrm{d}X(t) = \mu \mathrm{d}t + \sigma \mathrm{d}B(t) + \mathrm{d}L(t)$$

Then, if g is suitably smooth, on $\{T > t\}$,

$$\mathrm{d}M(t) = (r(X(t)) + \mathscr{L}g(X(t)))\mathrm{d}t + g'(X(t))\sigma \mathrm{d}B(t) + g'(0)\mathrm{d}L(t)$$

where

$$\mathscr{L} = \mu \frac{\mathrm{d}}{\mathrm{d}x} + \frac{1}{2}\sigma^2 \frac{\mathrm{d}^2}{\mathrm{d}x^2}$$

In order for $M(\cdot)$ to be a supermartingale, select $g \geq 0$ so that:

$$r + \mathscr{L}g \leq 0 \text{ on } C$$
$$g'(0) \leq 0$$

e.g. d-dimensional RBM on $\mathbb{R}^d_+$:

$$\mathrm{d}X(t) = \mu \mathrm{d}t + \Sigma^{1/2}\mathrm{d}B(t) + R\mathrm{d}L(t)$$

In order for $M(\cdot)$ to be a supermartingale, select $g \geq 0$ so that:

$$\begin{aligned} & r + \mathscr{L}g \leq 0 \text{ on } \mathbb{R}^d_+ \\ \text{s/t} \quad & \nabla g \cdot R_i \leq 0 \text{ on } F_i, 1 \leq i \leq d \end{aligned}$$

Paired Lyapunov Functions

For the CLT

$$t^{\frac{1}{2}}\left(\frac{1}{t}\int_0^t f(X(s))\,\mathrm{d}s - \pi f\right) \Rightarrow \eta\mathcal{N}(0,1) \quad \text{as } t \to \infty$$

Need a bound on η^2:

$$\eta^2 = \frac{\mathbb{E}_\nu(\int_0^\tau f_c(X(s))\,\mathrm{d}s)^2}{\mathbb{E}_\nu \tau}$$

where $f_c(\cdot) = f(\cdot) - \pi f$. Note that

$$\begin{aligned}\mathbb{E}_\nu\left(\int_0^\tau f_c(X(s))\,\mathrm{d}s\right)^2 &= 2\mathbb{E}_\nu \int_0^\tau f_c(X(t)) \int_t^\tau f_c(X(s))\,\mathrm{d}s\mathrm{d}t \\ &= 2\mathbb{E}_\nu \int_0^\tau f_c(X(t))u^*(X(t))\,\mathrm{d}t\end{aligned}$$

Choose g_1, g_2 non-negative so that

$$\begin{aligned} Pg_1 &\le g_1 - f \text{ on } C^c \\ Pg_2 &\le g_2 - fg_1 \text{ on } C^c \\ \sup_{x \in C} P(g_1 + g_2) &< \infty \end{aligned}$$

Return to G/G/1 waiting time sequence $(W_n : n \ge 0)$ with $f(x) = x$, $g_1(x) = a_1 x^2$, and $g_2(x) = x^4$

$$\eta^2 < \infty \text{ if } \mathbb{E}|Z_1|^4 < \infty$$

Condition is best possible...

Smoothness of Equilibrium Distributions

$P(\theta, x, \mathrm{d}y)$, $\pi(\theta, \mathrm{d}y)$, $f : S \to \mathbb{R}_+$

Goal: Establish smoothness of $\pi(\cdot)f$ (and obtain a formula)

$$\begin{aligned} \pi(\theta+h)f - \pi(\theta)f &= \pi(\theta+h)f_c \\ &= \pi(\theta+h)(I - P(\theta)u^*) \quad (u^* \text{ satisfies } u = f_c + Pu) \\ &= \pi(\theta+h)(P(\theta+h) - P(\theta))u^* \end{aligned}$$

Need to make sure $\frac{1}{h}(P(\theta+h) - P(\theta))u^*$ is well-behaved...requires pair of Lyapunov functions

For G/G/1 queue with $f(x) = x$, this approach leads to close to optimal conditions ($\mathbb{E}|Z|^3 < \infty$); previous operator-theoretic methods required exponential moments on Z

Concluding Remarks

Lyapunov methods/supermartingale ideas are key to establishing many key estimates and bounds for Markov chains/processes

While there is a systematic mechanism for guessing such functions, their actual verification and construction can require a LOT of creativity!! (even in the presence of the random walk structure typical of queues)