A Tutorial on Markov Chains

Lyapunov Functions, Spectral Theory
Value functions, and Performance Bounds

Sean Meyn

Department of Electrical and Computer Engineering
University of Illinois
and the Coordinated Science Laboratory

Joint work with R. Tweedie, I. Kontoyiannis, and P. Mehta

Supported in part by NSF (ECS 05 23620, and prior funding), and AFOSR

Markov Chains - Who Cares?

Goals for the week: Understanding the highlights on

- Stochastic Lyapunov Theory
- Dynamic Programming and Value Functions
- Spectral Theory and Model Reduction

Markov Chains - Who Cares?

Goals for the week: Understanding the highlights on

- Stochastic Lyapunov Theory
- Dynamic Programming and Value Functions
- Spectral Theory and Model Reduction

Why should you care?

Markov Chains - Who Cares?

Why I care:

- Optimal Control, Risk Sensitive Optimal Control
- Approximate Dynamic Programming
- Dynamic Economic Systems
- Finance
- Large Deviations
- Simulation
- Google

Markov Chains - Who Cares?

Why I care:

- Optimal Control, Risk Sensitive Optimal Control
- Approximate Dynamic Programming
- Dynamic Economic Systems
- Finance
- Large Deviations
- Simulation
- Google

Every one of these topics is concerned with computation or approximations of Markov models, particularly *value functions*

Objectives for Control

Nonlinear state space model $\equiv$ (controlled) Markov process,

state process $\boldsymbol{X}$

Typical form:

$$dX(t) = f(X(t), U(t))\,dt + \sigma(X(t), U(t))\,dW(t)$$

control

noise

Objectives for Control

Nonlinear state space model $\equiv$ (controlled) Markov process,

state process $\boldsymbol{X}$

Typical form:

$$dX(t) = f(X(t), \underbrace{U(t)}_{\text{control}})\, dt + \sigma(X(t), \underbrace{U(t)}_{\text{control}})\, \underbrace{dW(t)}_{\text{noise}}$$

Questions: For a given feedback law,

- Is the state process stable?
- Is the average cost finite? $\mathsf{E}[c(X(t), U(t))]$
- Can we solve the DP equations? $\min_u \Big\{ c(x,u) + \mathcal{D}_u h^*\,(x) \Big\} = \eta^*$
- Can we approximate the average cost η^* ? The value function h^* ?

Outline and Reading

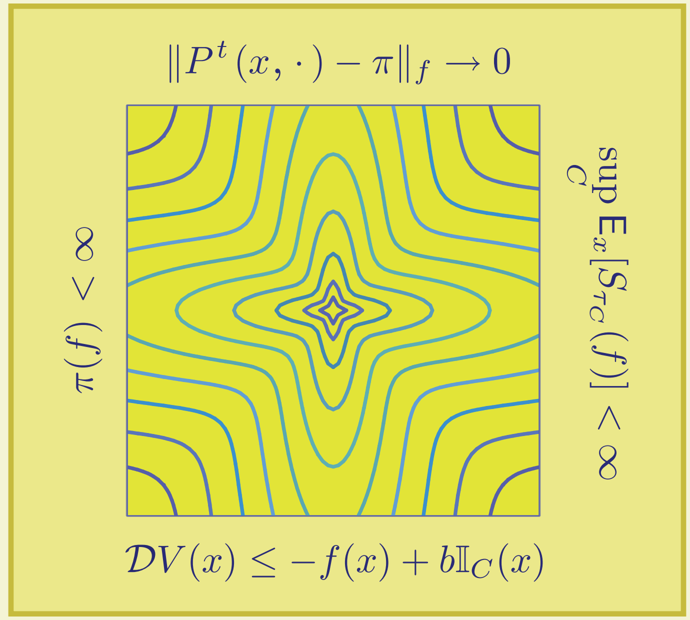

Monday: An Introduction

Motivation, and structural theory of
Markov models without control

Reading: Sections A1-A3 of CTCN

Tuesday: Value Functions

Lyapunov drift conditions and value functions
An introduction to dynamic programming

Reading: Sections A4-A6 of CTCN
See also Chs. 8 and 9, and Part III of MCSS

Thursday: Approximate Dynamic Programming

Approximations via deterministic ODE models
TD-learning and Q-learning algorithms

Reading: Section 11.5 of CTCN, Lecture Notes.
Recent publications `netfiles.uiuc.edu/meyn/www/spm_pubs.html`

Friday: Spectral Theory

Model reduction for Markov models based on spectral theory
Lectures based on joint with with Kontoyiannis, Huisinga, and Schuette,
`netfiles.uiuc.edu/meyn/www/spm_files/PhaseTransitions/PhaseTransitions.html`

References

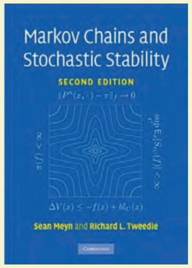

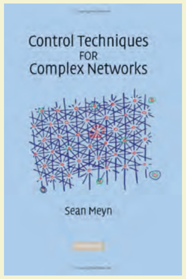

[1,4] ψ-Irreducible foundations

[2,11,12,13] Mean-field models, ODE models, and Lyapunov functions

[1,4,5,9,10] Operator-theoretic methods. *See also appendix of* [2]

[3,6,7,10] Generators and continuous time models

[1] S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, Cambridge, second edition, 2009. Published in the Cambridge Mathematical Library.

[2] S. P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, Cambridge, 2007. Pre-publication edition online: `http://black.csl.uiuc.edu/~meyn`.

[3] S. N. Ethier and T. G. Kurtz. *Markov Processes : Characterization and Convergence*. John Wiley & Sons, New York, 1986.

[4] E. Nummelin. *General Irreducible Markov Chains and Nonnegative Operators*. Cambridge University Press, Cambridge, 1984.

[5] S. P. Meyn and R. L. Tweedie. Generalized resolvents and Harris recurrence of Markov processes. *Contemporary Mathematics*, 149:227–250, 1993.

[6] S. P. Meyn and R. L. Tweedie. Stability of Markovian processes III: Foster-Lyapunov criteria for continuous time processes. *Adv. Appl. Probab.*, 25:518–548, 1993.

[7] D. Down, S. P. Meyn, and R. L. Tweedie. Exponential and uniform ergodicity of Markov processes. *Ann. Probab.*, 23(4):1671–1691, 1995.

[8] P. W. Glynn and S. P. Meyn. A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.*, 24(2):916–931, 1996.

[9] I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13:304–362, 2003. Presented at the INFORMS Applied Probability Conference, NYC, July, 2001.

[10] I. Kontoyiannis and S. P. Meyn. Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes. *Electron. J. Probab.*, 10(3):61–123 (electronic), 2005.

[11] W. Chen, D. Huang, A. Kulkarni, J. Unnikrishnan, Q. Zhu, P. Mehta, S. Meyn, and A. Wierman. Approximate dynamic programming using fluid and diffusion approximations with applications to power management. Accepted for inclusion in the 48th IEEE Conference on Decision and Control, December 16-18 2009.

[12] P. Mehta and S. Meyn. Q-learning and Pontryagin's Minimum Principle. Accepted for inclusion in the 48th IEEE Conference on Decision and Control, December 16-18 2009.

[13] G. Fort, S. Meyn, E. Moulines, and P. Priouret. ODE methods for skip-free Markov chain stability with applications to MCMC. *Ann. Appl. Probab.*, 18(2):664–707, 2008.

See also earlier seminal work by Hordijk, Tweedie, ... full references in [1].