Control of queueing systems in the **moderate** and **large** deviation regime

Rami Atar (Technion, Israel)
atar@ee.technion.ac.il

Based on work with

Anup Biswas (Technion), **Anindya Goswami** (IISER)
and **Adam Shwartz** (Technion)

The 2012 Stochastic Networks Conference

Talk based on the following papers

[1] *Risk-sensitive control for the parallel server model and an exponential version of the c mu rule*, with **A. Goswami and A. Shwartz**, Preprint

[2] *Asymptotic optimality of a fixed priority rule for a queueing problem in the large deviation regime*, with **A. Goswami and A. Shwartz**, Work in progress

[3] *Control of the multiclass G/G/1 queue in the moderate deviation regime*, with **A. Biswas**, Preprint

Ordinary, moderate, large deviations - i.i.d. r.v.s

Z_i iid, centered, finite exponential moments.

Fix $a > 0$. As $n \to \infty$,

$$P\Big\{\frac{Z_1 + \cdots + Z_n}{\sqrt{n}} > a\Big\} = O(1) \qquad \text{ordinary}$$

$$1 \ll b_n \ll \sqrt{n} \quad P\Big\{\frac{Z_1 + \cdots + Z_n}{b_n\sqrt{n}} > a\Big\} = e^{-\gamma b_n^2 + o(b_n^2)} \qquad \text{moderate}$$

$$P\Big\{\frac{Z_1 + \cdots + Z_n}{n} > a\Big\} = e^{-\gamma n + o(n)} \qquad \text{large}$$

Ordinary, moderate, large deviations - renewal process

A= a renewal process, finite exponential moments for interarrivals; mean interarrival=1

$$\hat{A}^n(t) = \frac{A(nt) - nt}{\sqrt{n}}, \quad \tilde{A}^n(t) = \frac{A(nt) - nt}{b_n \sqrt{n}}, \quad \bar{A}^n(t) = \frac{A(nt)}{n}$$

Fix R a subset of a function space $[0, \infty) \to \mathbb{R}$. As $n \to \infty$,

$P\{\hat{A}^n \in R\} = \gamma + o(1) \qquad \gamma = P(W \in R)$, W=Brownian motion

$P\{\tilde{A}^n \in R\} = e^{-\gamma b_n^2 + o(b_n^2)} \qquad \gamma = \inf\{\int_0^\infty L(\dot{\varphi}(t))dt : \varphi \in R, \varphi \in \mathcal{AC}\}$

$P\{\bar{A}^n \in R\} = e^{-\gamma n + o(n)} \qquad \gamma = \inf\{\int_0^\infty \tilde{L}(\dot{\varphi}(t))dt : \varphi \in R, \varphi \in \mathcal{AC}\}$

The parallel server model

A multiclass job allocation problem:

- Classes $i = 1, 2, \ldots, I$
- Multiple stations $j = 1, 2, \ldots, J$
- N_j identical servers in station j
- Arrival rates λ_i
- Service rate μ_{ij}

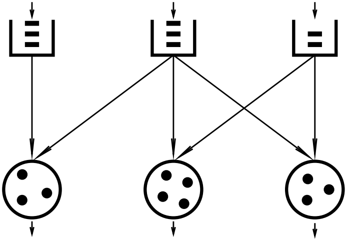

Allocation of jobs to servers is regarded a control

How different costs emphasize different asymptotic regimes

Let $X_i(t)$= number of class-i customers in the system at time t. Let $X = (X_1, \dots, X_I)$. Consider

$$\hat{X}^n(t) = \frac{X(nt)}{\sqrt{n}}, \quad \tilde{X}^n(t) = \frac{X(nt)}{b_n\sqrt{n}}, \quad \bar{X}^n(t) = \frac{X(nt)}{n}$$

For a path $\varphi : [0, T] \to \mathbb{R}^I$, consider $K(\varphi) = \int_0^T h(\varphi(t))dt + g(\varphi(T))$.

Problem 1. Minimize $\hat{J}_n(U) = E^U[K(\hat{X}^n)]$ under critical load

Problem 2. Minimize $\tilde{J}_n(U) = \frac{1}{b_n^2} \log E^U[e^{b_n^2 K(\tilde{X}^n)}]$ under critical load

Problem 3. Minimize $\bar{J}_n(U) = \frac{1}{n} \log E^U[e^{nK(\bar{X}^n)}]$

- Linear (in K) cost vs risk-sensitive cost
- Risk-sensitive cost motivation: quality of service assurance, buffer overflow considerations; robustness properties
- MD scale has further robustness properties
- A different formulation: minimize probability of buffer overflow

Role of the Laplace Principle

For a fixed control U, the limit is given by a **variational problem**

$$\frac{1}{b_n^2} \log E[e^{b_n^2 K(\tilde{X}^n)}] \approx \sup_{\varphi \in \mathcal{AC}} [K(\varphi) - \mathbb{I}^U(\varphi)]$$

Optimization over controls gives rise to a **game**

$$\inf_U \frac{1}{b_n^2} \log E[e^{b_n^2 K(\tilde{X}^n)}] \approx \inf_U \sup_{\varphi \in \mathcal{AC}} [K(\varphi) - \mathbb{I}^U(\varphi)]$$

Related work

- *On queueing models in MD scale:*
  - **Puhalskii** 1999 MD principle for single server and single class open queueing networks in heavy traffic
  - **Majewski** 2006 feedforward multiclass networks
  - **Wischik** 2001, **Cruise** 2012 links between queueing models at OD, MD, LD scales
  - **Remark:** Wischik 2001 conjectures about *similarities between control theoretic treatments at the diffusion and MD scales*
- *On control theory under LD regime, for small noise diffusions:* **Fleming-Souganidis** 1986, **Dupuis-Kushner** 1989, **Dupuis-McEneaney** 1997
- *On control theory for Markovian stochastic networks under LD regime:* **A-Dupuis-Shwartz** 2003

Overview of results

In LD scale

- A Markovian setting; the full parallel server model
- Convergence of the risk-sensitive control value to the value of a DG
- Characterization of the game's value as the unique solution to a PDE
- Solution in a special case

In MD scale

- General service time distributions; the single server case
- Convergence of the risk-sensitive control value to the value of a DG
- Solution of the game is great generality

Overview of the talk

- Markovian network in the LD regime - statement of convergence result and characterization of limit
- Proof idea
- Explicit solution of the DG in a special case (and lack of general explicit solution)
- The G/G/1 setting in MD regime
- Statement of convergence result
- Comparison of the two games and why the MD game is easier
- Explicit general solution of the MD game
- Derivation of the explicit solution

Markovian network in the LD regime

Model is parameterized by n.

Index set for customer classes: $\mathcal{I} = \{1, 2, \ldots, I\}$;
for service stations: $\mathcal{J} = \{1, 2, \ldots, J\}$.

$N_j(n)$ servers in station j.

Scaling assumptions:

- $\sum_j N_j(n) = o(n)$
- Arrival rates: $\lambda_i(n) = \lambda_i n$
- Service rates: $\mu_{ij}(n)$ satisfying $N_j(n)\mu_{ij}(n) = \mu_{ij} n$. (Thus total per-station service capacity scales like n.)

Denote $\Xi_t^n = (\Xi_t^{n,i})_{i \in \mathcal{I}}$; $\Xi_t^{n,i}$ the number of class-i customers at time t in the n-th system.

$X_t^n = n^{-1} \Xi_t^n$ takes values in $G^n := n^{-1}\mathbb{Z}_+^I$. Denote $G = \mathbb{R}_+^I$.

Admissible controls via the martingale formulation

$$U := \left\{ u \in \mathbb{R}^{I \times J} : u_{ij} \geq 0,\ i \in \mathcal{I}, j \in \mathcal{J}, \sum_i u_{ij} \leq 1,\ j \in \mathcal{J} \right\}$$

Consider the controlled generator acting on the space of functions $G^n \to \mathbb{R}$, for each $n \in \mathbb{N}$, $u \in U$,

$$\begin{aligned} \mathcal{L}^{n,u} f(x) = & \sum_i n\lambda_i (f(x + \frac{1}{n} e_i) - f(x)) \\ & + \sum_{i,j} n\mu_{ij} u_{ij} (f(x - \frac{1}{n} e_i) - f(x)) 1_{\{x - \frac{1}{n} e_i \in \mathbb{R}_+^I\}}, \qquad x \in G^n. \end{aligned}$$

(U^n, X^n) are regarded an admissible control and a corresponding controlled process starting from $(x, t) \in G^n \times [0, T]$ if

- $P(X_t^n = x) = 1$
- $\sum_j N_j(n) U_{ij}^n \leq \Xi^{n,i} = nX^{n,i}, \qquad i \in \mathcal{I}$
- For each bounded $f : G^n \to \mathbb{R}$, the following process is a martingale

$$f(X_s^n) - \int_t^s \mathcal{L}^{n, U^n(r)} f(X_r^n) dr, \quad s \geq t$$

LLN-scaled cost

$$J^n(t,x,U^n) = E_{t,x}\Big[\int_t^T h(X_s^n)ds + g(X_T^n)\Big], \qquad V^n(t,x) = \inf_{U^n} J^n(t,x,U^n)$$

We will always assume h and g are Lipschitz continuous and nondecreasing. Limiting dynamics

$$\begin{cases} \psi(s) = x + \int_t^s v(u(r))dr, & s \in [t,T], \\ \varphi = \Gamma[\psi], \end{cases}$$

where $v(u) = \sum_i \lambda_i e_i - \sum_{i,j} u_{ij}\mu_{ij} e_i$

$$\Gamma_1[\psi](s) = \psi(s) - \inf_{r\in[t,s]} \psi(r) \wedge 0, \qquad s \in [t,T],$$

$\Gamma = \Gamma^{t,T}$ mapping $C([t,T]:\mathbb{R}^I)$ to itself, $\Gamma[\psi]_i = \Gamma_1[\psi_i]$, for $i \in \mathcal{I}$.
Γ is called the *Skorohod map on G with normal constraint directions.*

Convergence

A well-known limit result: $V^n \to V = \inf_u c(t, x, u)$, where $c(t, x, u) = \int_t^T h(\varphi_s) ds + g(\varphi_T)$.

A well-known special case: A single service station, and either h or g given as $\sum_i c_i x_i$. The DG is solved explicitly. The optimal (feedback) control is of fixed priority type, namely

$$u_1 = 1_{\{\varphi_1 > 0\}}, \quad u_2 = 1_{\{\varphi_1 = 0, \varphi_2 > 2\}}, \ldots, u_I = 1_{\{\varphi_1 + \cdots + \varphi_{I-1} = 0, \varphi_I > 0\}},$$

where

$$c_1 \mu_1 \geq c_2 \mu_2 \geq \cdots \geq c_I \mu_I.$$

This is called the $c\mu$ rule.

There are various generalizations.

A risk-sensitive setting

Cost that penalizes heavily large perturbations

$$\bar{J}^n(t,x,U^n) = \frac{1}{n}\log E_{t,x}\exp\left\{ n\Big[\int_t^T h(X_s^n)ds + g(X_T^n)\Big]\right\}$$

$$\bar{V}^n(t,x) = \inf_{U^n} \bar{J}^n(t,x,U^n)$$

Notation

For $u \in U$ and $m = ((\bar{\lambda}_i)_{i \in \mathcal{I}}, (\bar{\mu}_{ij})_{i \in \mathcal{I}, j \in \mathcal{J}}) \in M = \mathbb{R}_+^I \times \mathbb{R}_+^{I \times J}$, let

$$v(u, m) = \sum_i \bar{\lambda}_i e_i - \sum_{ij} u_{ij} \bar{\mu}_{ij} e_i$$

$$\rho(u, m) = \sum_i \lambda_i \, \ell\Big(\frac{\bar{\lambda}_i}{\lambda_i}\Big) + \sum_{ij} u_{ij} \mu_{ij} \, \ell\Big(\frac{\bar{\mu}_{ij}}{\mu_{ij}}\Big),$$

where

$$\ell(x) = \begin{cases} x \log x - x + 1, & x \geq 0, \\ +\infty, & x < 0 \end{cases}$$

Definition of value in Elliott-Kalton sense

$$\bar{U} := \{u : [0, T] \to U; u \text{ is measurable}\},$$

$$\bar{M} = \{m : [0, T] \to M; m \text{ is measurable}, \ell \circ m \text{ is locally integrable}\}.$$

$$\begin{cases} \psi(s) = x + \int_t^s v(u(r), m(r))dr, \quad s \in [t, T], \\ \varphi = \Gamma[\psi] \end{cases}$$

$$c(t, x, u, m) = \int_t^T [h(\varphi(s)) - \rho(u(s), m(s))]ds + g(\varphi(T))$$

A mapping $\alpha : \bar{M} \to \bar{U}$ is called a *strategy for the minimizing player* if for every $m, \tilde{m} \in \bar{M}$ and $s \in [0, T]$,

$$m(r) = \tilde{m}(r) \text{ for a.e. } r \in [0, s] \text{ implies } \alpha[m](r) = \alpha[\tilde{m}](r) \text{ for a.e. } r \in [0, s].$$

Define the lower value $V^-(t, x) = \inf_{\alpha \in A} \sup_{m \in \bar{M}} c(t, x, \alpha[m], m)$.
Similarly, upper value $V^+(t, x) = \sup_{\beta \in B} \inf_{u \in \bar{U}} c(t, x, u, \beta[u])$.
The game is said to have value if $V^- = V^+$.

The Hamilton-Jacobi-Isaacs equation

$$H(p) = \inf_{u \in U} \sup_{m \in M} [\langle p, v(u,m)\rangle - \rho(u,m)], \qquad p \in \mathbb{R}^I.$$

Let $\mathbf{I} : \partial G \to 2^{\mathcal{I}}$ be defined by

$$\mathbf{I}(x) = \{i \in \mathcal{I} : x_i = 0\}.$$

$$\begin{cases} V_t + H(D_x V) + h &= 0 \qquad \text{in } [0,T) \times G^o, \\ \langle D_x V(t,x), e_i \rangle &= 0 \qquad x \in \partial G, i \in \mathbf{I}(x), \\ V(T,x) &= g(x) \quad x \in G. \end{cases}$$

Limit results in LD scale

Theorem 1a:

- The game has value $V := V^- = V^+$.
- As $n \to \infty$, $\bar{V}^n(t, x_n) \to V(t, x)$ if $x_n \to x$.
- The value uniquely solves the HJI equation.

Proof

Step 1. $\liminf \bar{V}^n \geq V^+$ is shown as follows:

Consider the optimal feedback function u^n from state to control, for the prelimit problem. Consider an arbitrary strategy β. Construct a process X^n 'driven' by u^n and β. More precisely, find processes X^n, $\bar{u}^n$, $\bar{m}^n$, such that

- $\bar{m}^n = \beta[\bar{u}^n]$,
- $\bar{u}^n(s) = u^n(X^n(s))$,
- X^n solves the martingale problem for $\mathcal{L}^{n,\bar{u}^n,\bar{m}^n}$, where

$$\mathcal{L}^{n,u,m} f(x) = \sum_i n\bar{\lambda}_i (f(x + \frac{1}{n}e_i) - f(x)) + \sum_{i,j} n\bar{\mu}_{ij} u_{ij} (f(x - \frac{1}{n}e_i) - f(x)) 1_{\{x - \frac{1}{n}e_i \in \mathbb{R}^I_+\}}, \qquad x \in G^n.$$

Proof (continued)

Bellman's equation and the optimality of u^n imply the inequality

$$\mathcal{L}^{n,\bar{u}^n(t),\bar{m}^n(t)}\bar{V}^n(t,X^n(t))+\frac{d}{dt}\bar{V}^n(t,X^n(t))+h(X^n(t))-\rho(\bar{u}^n(t),\bar{m}^n(t))\leq 0.$$

Integrating up to T, taking the expectation gives

$$\begin{aligned}\bar{V}^n(t,x_n) &\geq E[\int_t^T (h(X^n(s)) - \rho(\bar{u}^n(s),\bar{m}^n(s)))ds + g(X^n(T))] \\ &\geq \inf_u c(t,x,u,\beta[u]) - \varepsilon_n,\end{aligned}$$

where martingale estimates can be used to obtain $\varepsilon_n \to 0$. Hence

$$\liminf \bar{V}^n(t,x_n) \geq \sup_\beta \inf_u c(t,x,u,\beta[u]) = V^+(t,x).$$

Proof (continued)

Step 2. Similarly, $V^+ \le \liminf \bar{V}^n \le \limsup \bar{V}^n \le V^-$

Step 3. Both V^- and V^+ are viscosity solutions to the HJI equation for which uniqueness holds.

A special case

- A single service station
- $h = 0$
- $g(x) = \sum c_i x_i$

Thus

$$\bar{J}^n(U^n) = \frac{1}{n} \log E^{U^n} \exp[\sum_i c_i \Xi_i(T)].$$

The DG cost:

$$c(t, x, u, m) = -\int_t^T \rho(u(s), m(s)) ds + \sum c_i \varphi_i(T).$$

This game can be solved.

Theorem 1b:

- A fixed priority policy, according to the index $(1 - e^{-c_i})\mu_i$, is optimal for the minimizing player.
- Preemptive priority according to the same index, is asymptotically optimal for the queueing model.

Remarks

- Reminiscent of the $c\mu$ rule.
- We can solve the game explicitly for multiple stations as well.
- We don't know of an explicit solution to the case of linear running cost or nonlinear terminal cost.

The game is subtle...

- For the cost $K(\varphi) = c \cdot \varphi(T)$, the optimal strategy u prioritizes according to $(1 - e^{-c})\mu$, **does not depend on** T
- In addition, the strategy for the maximizing player **does not depend on** T: $\bar{\lambda}_i(t) = \lambda_i e^{c_i}$, $\bar{\mu}_i(t) = \mu_i e^{-c_i}$
- One expects that the same strategies should be optimal for costs such as

$$\int_0^T c \cdot \varphi(t) dt, \qquad c \cdot \varphi(T_1) + c \cdot \varphi(T_2)$$

...but they're not! (we don't know what the optimal strategies are)

The multiclass G/G/1 queue in MD regime

The model is now driven by renewal processes. Denote

- $1/\lambda_i$ = mean interarrival for class i
- $\theta_i = 1/\mu_i$ = mean service time for class i

Critical load condition: $\sum_i \lambda_i/\mu_i = 1$

Limit results in MD scale

Theorem 2: Assume interarrival and service time distribution satisfying $E[e^{a\zeta^r}] < \infty$, where $a > 0$ and $r \in (0, 1]$ is such that $b_n^{r-2} n^{r/2} \to \infty$. Assume

$$w \mapsto \underset{y \in \mathbb{R}_+^I : \theta \cdot y = w}{\arg\min} \; h(y), \quad \text{is continuous.}$$

Then

- As $n \to \infty$, $b_n^{-2} \log \inf_U \tilde{J}_n(U) \to V$
- V can be characterized as the lower value of a differential game
- In the case where K is linear, the $c\mu$ rule is optimal for the game and asymptotically optimal for the queueing system

About the technique

We no longer assume Markovity, therefore the martingale formulation and the form of the generator are not available

We do not use the PDE

- A crucial element: explicit solution to the differential game
- the rate function is quadratic, and so the same sort of game characterizes the behavior under all distributions

Comparison of the DGs

LD game: maximizing [minimizing] player controls $(\bar{\lambda}, \bar{\mu})$ [resp. u]

Dynamics: $\varphi = \Gamma[\psi], \qquad \psi(t) = x + \int_0^t (\bar{\lambda}(s) - u(s) \bullet \bar{\mu}(s))ds$

Cost: $K(\varphi) - \int_0^T [1 \cdot \ell(\bar{\lambda}(s)) + u(s) \cdot \hat{\ell}(\bar{\mu}(s))]ds$

MD game: maximizing [minimizing] player controls $(\tilde{\lambda}, \tilde{\mu})$ [resp. η]

Dynamics: $\varphi(t) = x + \int_0^t (\tilde{\lambda}(s) - \tilde{\mu}(s))ds + \eta(t) \in \mathbb{R}_+^I$

Cost: $K(\varphi) - \int_0^T \sum [a_i \tilde{\lambda}_i(s)^2 + b_i \tilde{\mu}_i(s)^2]ds$

where a_i and b_i are positive constants, and $\theta \cdot \eta$ is nondecreasing

- MD game is simpler; in fact **completely solvable**

Game solution

Recall

$$\varphi(t) = x + \int_0^t (\tilde{\lambda}(s) - \tilde{\mu}(s))ds + \eta(t) \in \mathbb{R}_+^I$$

Multiply by θ,

$$\theta \cdot \varphi(t) = \theta \cdot x + \int_0^t \theta \cdot ((\tilde{\lambda}(s) - \tilde{\mu}(s))ds + \theta \cdot \eta(t) \geq 0$$

Given $(\tilde{\lambda}, \tilde{\mu})$, $\theta \cdot \eta$ given by the 1d Skorohod map minimizes $\theta \cdot \varphi$. Denoting the minimal response $\theta \cdot \varphi$ to $(\tilde{\lambda}, \tilde{\mu})$ by w, for the case $K(\varphi) = \int_0^T h(\varphi(t))dt$, for all φ,

$$K(\varphi) \geq K(\varphi^*), \qquad \varphi^*(t) = \operatorname*{arg\,min}_{y \in \mathbb{R}_+^I : \theta \cdot y = w(t)} h(y).$$

Concluding remarks

- In LD scale we presented an explicit optimal strategy in a special case; one would like to understand if it can be extended.
- Control theory in MD scale has not been considered before; we find that
  - it leads to games that can be solved in great generality
  - it is closely related to control theory at the diffusion scale
- We are interested in control theory at the MD scale for **more general stochastic networks**. It is potentially as rich as the existing theory for heavy traffic at diffusion scale.