Controlling rare events

Joint work with Rami Atar, Paul Dupuis

Technion, May 2005

MOTIVATION and OUTLINE

General issue: combining control theory and large deviations.

Test case: controlled networks with Large Deviations scaling.

Outline:

- A re-entrant line network
- Controlled Markov processes and Large Deviations
- Control of rare events
- The exit problem

- Projected dynamics and the Skorokhod problem
- Dynamic programming and the differntial game
- Representation of the optimal cost
- Explicit solutions

“Re-entrant line” queueing network:

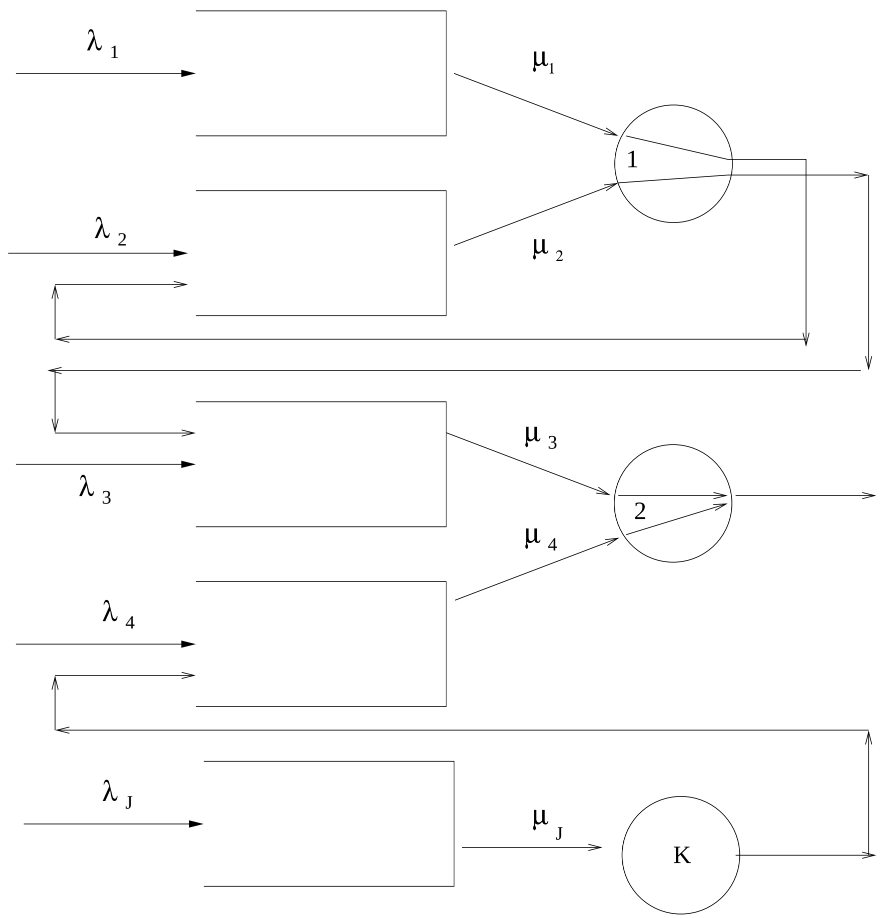

Controlled Markov processes: martingale formulation

We are given a generator $\mathcal{L}^u$ of a Markov process which depends on a control variable $u \in \mathbf{U}$.

Say X is a *controlled Markov process*, starting from x, with control process U if $\{X, U\}$ is adapted to $\mathcal{F}_t$, $P_x^U(X(0) = x) = 1$, and for every function f (in some class)

$$f(X(t)) - \int_0^t \mathcal{L}^{U(s)} f(X(s)) ds$$

is an $\mathcal{F}_t-$martingale.

Example: $p_{ij}(u)$ is, for each u, the transition matrix of a Markov chain, transition time from i to j is exponential $\lambda(i)$, and $U(t) = \tilde{U}(x(t))$ (feedback control).

The large Deviations scaling

Speed up time, shrink space: define

$$\tilde{\mathcal{L}}^{n,u} f(x) := n\tilde{\mathcal{L}}^{u} g(nx),$$

where $g(\cdot) = f(n^{-1}\cdot)$. For the scaled controlled process

$$f(X^n(t)) - \int_0^t \tilde{\mathcal{L}}^{n,u(s)} f(X^n(s))ds$$

is an $\mathcal{F}_t-$martingale.
Example: sum of J Poisson processes, with jump directions e_j

$$\mathcal{L}f(x) = \sum_{j=1}^{J} \lambda_j \left[f(x+e_j) - f(x) \right]$$

$$\mathcal{L}^n f(x) = \sum_{j=1}^{J} n\lambda_j \left[f\left(x+\frac{e_j}{n}\right) - f(x) \right].$$

Control of rare events: take 1.

Let G be some domain and σ the exit time.
Asymptotically (Large Deviations), the probability of exit is

$$P_n^u(\text{exit}) \approx e^{-nI(u;G)}, \quad \text{where} \quad I(u;G) = \inf_{\phi \text{ exits } G} \int_0^\sigma l_u(\phi, \dot{\phi})\, dt$$

where u denotes the control policy.
If exit is undesirable, maximize $I(u;G)$ with respect to u.

Problems:

1. Large Deviations principle does not hold for most u.
2. Order of limits: why maximize the limit?

Control of rare events: take 2.

Fix $c > 0$ and a domain G. Let σ be the exit time from G. Define the value function for the stochastic control problem

$$V(x) := - \inf_{\text{controlled Markov Proc.}} \log E_x^u e^{-c\sigma} \quad x \in G.$$

The value function for the scaled process is

$$V^n(x) \doteq - \inf_{\text{controlled Markov Proc.}} n^{-1} \log E_x^{u,n} e^{-nc\sigma_n}.$$

Take limit of the extreme process!

Questions: existence, characterization of limit.

The queueing network:
J queues, each served by a single one of $K \leq J$ servers.
Arrivals λ_j, services μ_j, next station $r(j)$.
Control: $u_i = p$ if queue i is served with a fraction $0 \leq p \leq 1$ of the effort (w.p. p).
Generator: for $f : \mathbb{Z}^J \to \mathbb{R}$,

$$\mathcal{L}^u f(x) = \sum_{j=1}^{J} \lambda_j \left[f(x + e_j) - f(x) \right] \tag{1}$$

$$+ \sum_{j=1}^{J} u_j \mu_j 1_{x + e_{r(j)} - e_j \in \mathbb{Z}_+^J} \left[f(x + e_{r(j)} - e_j) - f(x) \right]. \tag{2}$$

The *unconstrained controlled process* Y is not restricted to $\mathbb{Z}_+^J$, and so its generator $\mathcal{L}_0^u$ is without $1_{x + e_{r(j)} - e_j \in \mathbb{Z}_+^J}$.

An exit problem:

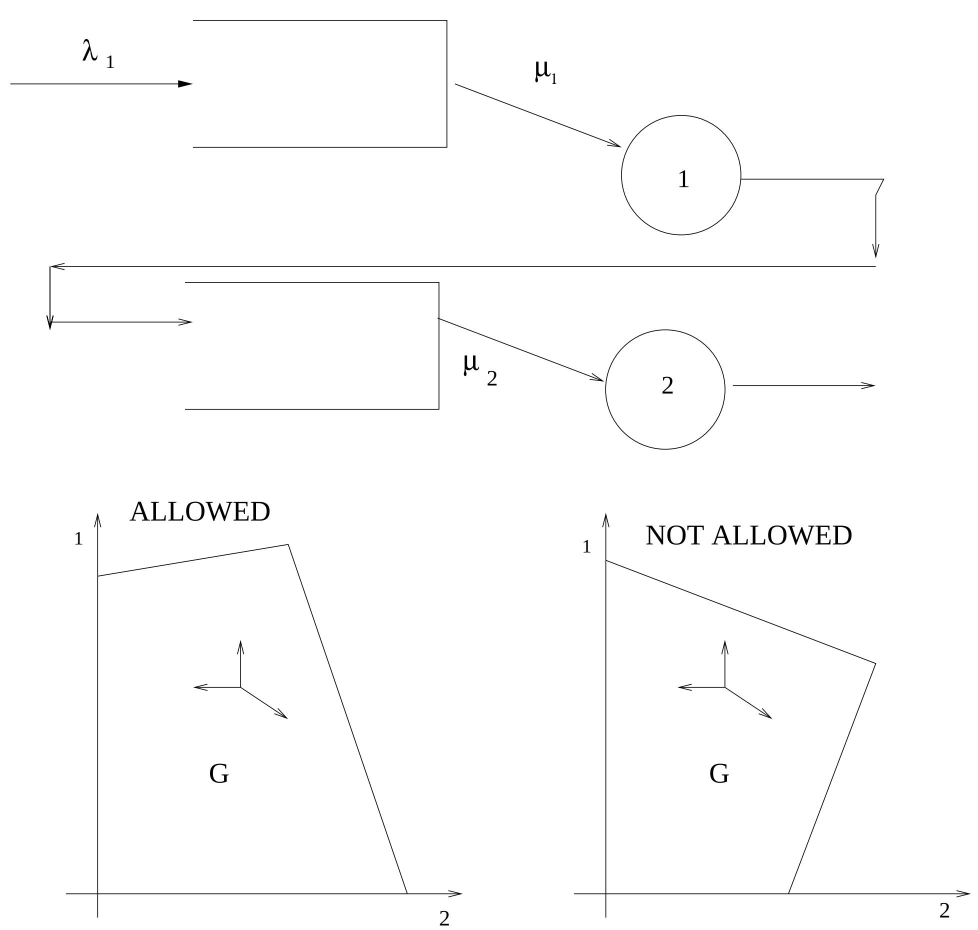

Let G be a piecewise smooth, bounded open set (with $0 \in G$). Assume: on every point in $\mathbb{R}^J$ of its boundary, it is possible to exit through an external arrival, *or* G is a hyper-rectangle. Define

$$\sigma^n := \inf \{t : X^n(t) \notin G\} .$$

Goal: keep process within the bounded region G for as long as possible: fix $c > 0$ and minimize over admissible controls

$$V^n(x) := -\inf_u n^{-1} \log E_x^{u,n} e^{-nc\sigma^n}.$$

Find the minimum and an "optimal" policy for large n.

Note: the Large Deviations Principle as well as representation of the rate function are available only for very few control policies, since the theory requires "smooth statistics" and rates which are bounded away from 0.

We need a way to describe the dynamics near the boundary of $\mathbb{R}_+^J$: not only for the queueing system but also for "limiting" processes. Suppose a queue is empty and the unconstrained process experiences a "service," changing its state from y to $y+v$. The state then changes by $e_{r(j)} - e_j$. Define for such y

$$\pi(y+v) = y = (y+v) - v.$$

There exists π so that for y in $\mathbb{Z}_+^J$ we have $\pi(y) = y$, and for y outside $\mathbb{Z}_+^J$ there is a unique $\pi(y)$ on the boundary—the projection along the directions $e_j - e_{r(j)}$. Then

$$1_{x+e_{r(j)}-e_j \in \mathbb{Z}_+^J} \left[f(x+e_{r(j)} - e_j) - f(x) \right] = f(\pi[x+e_{r(j)}-e_j]) - f(x).$$

This projection leads to the projected velocity

$$\pi(x,v) := \lim_{\Delta \downarrow 0} \frac{\pi(x+\Delta v) - \pi(x)}{\Delta}.$$

This projection defines a Skorohod map. For a continuous function ψ with $\psi(0) \in \mathbb{R}_+^J$ let $\Gamma(\psi)$ satisfy

- $\Gamma(\psi)(0) = \psi(0)$,

- $\Gamma(\psi)(t_2) - \Gamma(\psi)(t_1) = \psi(t_2) - \psi(t_1)$ whenever $\Gamma(\psi)(t)$ is in the interior of $\mathbb{R}_+^J$ for $t \in [t_1, t_2]$,

- while $\Gamma(\phi)(t)$ is on the boundary, if the dynamics of ψ point "out" of the boundary of $\mathbb{R}_+^J$, the path is "pushed back just enough" to stay inside, and along the specified directions.

Lemma 1 (Dupuis et. al.) Γ *is a Lipschitz continuous map from* $C^J([0,\infty))$ *(with the sup norm) to itself.*

The scaled processes X^n and Y^n:
time accelerates by a factor of n and space is shrunk by the same factor. The respective generators are

$$\mathcal{L}^{n,u} f(x) = n\mathcal{L}^u g(nx), \quad \mathcal{L}_0^{n,u} f(x) = n\mathcal{L}^u g(nx), \tag{3}$$

for $f : n^{-1}\mathbb{Z}^J \to \mathbb{R}$ and where $g = f(n^{-1}\cdot)$.
These define the scaled processes X^n and Y^n.

The optimal cost: dynamic programming.

Define $W^n(x) = \exp[-nV^n(x)]$. Dynamic programming Eq.:

$$0 = \inf_{u \in U} \left[\tilde{\mathcal{L}}^{n,u} W^n(x) - ncW^n(x) \right], \quad x \in G^n, \tag{4}$$

and $W^n(x) = 1$, $x \notin G^n$. In terms of V^n (deviding by W^n)

$$\begin{aligned} 0 \;=\; & \inf_{u \in U} \left[\sum_{j=1}^{J} n\lambda_j \left(\exp\left[-nV^n\left(x + \frac{1}{n} v_j \right) + nV^n(x) \right] - 1 \right) - nc \right. \\ & \left. + \sum_{i=1}^{J} n\mu_i u_i \left(\exp\left[-nV^n\left(x + \frac{1}{n} \pi(x, \tilde{v}_i) \right) + nV^n(x) \right] - 1 \right) \right] \end{aligned}$$

for $x \in G^n$. Devide by n and use

$$[e^y - 1] = \sup_{x>0} [xy - (x \log x - x + 1)] \quad \text{yields}$$

$$0 = \sup_{u \in U} \inf_{m \in M} [\mathcal{L}^{n,u,m} V^n(x) + c + \rho(u,m)],$$

λ_i and μ_i denote arrival rate and service rate of queue i,
The control vector $u = (u_1, \ldots, u_J)$,
the "perturbed rates" vector $m = (\bar{\lambda}_1, \ldots, \bar{\lambda}_J, \bar{\mu}_1, \ldots, \bar{\mu}_J)$,
$l : \mathbb{R} \to \mathbb{R}_+ \cup \{+\infty\}$ and $\rho : U \times M \to \mathbb{R} \cup \{+\infty\}$ are

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

$$\rho(u,m) = \sum_{i=1}^{J} \lambda_i l\left(\frac{\bar{\lambda}_i}{\lambda_i}\right) + \sum_{i=1}^{J} u_i \mu_i l\left(\frac{\bar{\mu}_i}{\mu_i}\right).$$

$\mathcal{L}^{n,u,m}$ is the generator for the same process, but with rates m.

The Bellman equation for our control problem is therefore

$$\begin{cases} 0 = \sup_{u\in U} \inf_{m\in M} [\mathcal{L}^{n,u,m} V^n(x) + c + \rho(u,m)], & x \in G^n \\ V^n(x) = 0, & x \notin G^n. \end{cases} \quad (5)$$

This leads to a differential game.

u_i is the service rate to class i. Player 1 control space is

$$\bar{U} = \{u : [0,\infty) \to U \ ; \ u \text{ measurable}\}.$$

Denoting $M = [0,\infty)^{2J}$, the control space for player 2 is

$$\bar{M} = \{m = (\bar{\lambda}_1, \dots, \bar{\lambda}_J, \bar{\mu}_1, \dots, \bar{\mu}_J) : [0,\infty) \to M; \ m \text{ measurable}\}$$

For $u \in U$ and $m \in M$ define

$$v(u,m) = \sum_{j=1}^{J} \bar{\lambda}_j v_j + \sum_{i=1}^{J} u_i \bar{\mu}_i \tilde{v}_i,$$

where $\quad v_j = e_j, \quad \tilde{v}_i = e_{r(i)} - e_i.$

The dynamics are given by

$$\begin{cases} \dot{\phi}(t) = \pi(\phi(t), v(u(t), m(t))), \\ \\ \phi(0) = x. \end{cases}$$

Note that π is not continuous in ϕ. Nonetheless,

Lemma 2 (Dupuis et al) *Given u and m there exists a unique solution to the ODE.*

This is to be expected since Γ is Lipschitz, but is not trivial.

The exit time is $\sigma = \inf\{t : \phi(t) \notin G\}$.

The cost for the game is (recall $c \geq 0$)

$$C(x, u, m) = \int_0^\sigma [c + \rho(u(t), m(t))]dt.$$

$\alpha : \bar{M} \to \bar{U}$ is a *strategy for player P1* if

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

implies

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

In an analogous way, one defines a *strategy for player P2* $\beta : \bar{U} \to \bar{M}$. The set of all strategies for P1 [P2] is denoted A [B].

The lower value for the game is

$$V^{-}(x) = \inf_{\beta \in B} \sup_{u \in \bar{U}} C(x, u, \beta[u]),$$

and the upper value is

$$V^{+}(x) = \sup_{\alpha \in A} \inf_{m \in \bar{M}} C(x, \alpha[m], m).$$

Notes: In general the upper value may be smaller than the lower value: no symmetry!
The value may be discontinuous accross a boundary!

Theorem 1 *The differential game has a value $V(x)$, so that*
$V^{+}(x) = V^{-}(x) = V(x)$, $x \in G$, $V^{+}(x) = V^{-}(x) = V(x) = 0$, $x \in \partial G$.

Proof outline:

- $V^{-}(x)$ is Lipschitz in G. The proof uses an explicit construction and pathwise arguments.

- By definition, $V^{-}(x)$ satisfies a variational equation.

- Therefore $V^{-}(x)$ satisfies the PDE: for $x \in G$,

  $H(DV^{-}(x)) = 0$, where $H(q) = \inf_{m} \sup_{u} [\langle q, v(u,m)\rangle + c + \rho(u,m)]$

  $V^{-}(x) = 0$ for $x \in \partial G$,
  $\langle DV^{-}(x), \gamma_i \rangle = 0$ for γ_i a direction of constraint if $x \in \partial \mathbb{R}^{J}_{+}$.

Actually, we show that V^- is both a viscosity super solution and a viscosity sub solution, hence a viscosity solution (analytic arguments). This is required since there is no C^2 solution (in fact the solution may not be even C^0!) Next we show

- $V^+(x)$ is also Lipschitz, and satisfies the same PDE (min-max theorem)

- there is uniqueness of solutions to the PDE. This involves showing that every sub solution is smaller than a super solution, and using the fact that a solution is by definition both a sub and a super solution.

This concludes the proof that the differential game has a value.

Theorem 2 *$V^n(x)$ is Lipschitz, uniformly in n, and*

$$\limsup_{n\to\infty} V^n(x) \le V^+(x), \qquad \liminf_{n\to\infty} V^n(x) \ge V^-(x).$$

Therefore $\lim_{n\to\infty} V^n(x) = V(x)$.

Proof outline:

- Using an explicit construction of a pair of processes with the same control (coupling) show Lipschitz property.

- Using the dynamic programming equation (of the control problem) and properties of the generators, derive the bounds.

Summary of main results:

Theorem 3 *(i) The limit* $\lim_{n\to\infty} V^n(x) = V(x)$ *exists: in fact,*

$$\lim_{\epsilon\downarrow 0}\limsup_{n\to\infty}\sup\{|V^n(x)-V(y)| : x\in G^n, y\in G, |x-y|\le\epsilon\} = 0.$$

(ii) The differential game has a value, and

$$V(x) = V^{+}(x) = V^{-}(x).$$

That is, The differential game with constrained dynamics has a Lipschitz value, which coincides with the limit of the optimal values of the exit control problem.

Example 1: competing queues

Single server, customers leave upon service completion.
Average cost problem: minimize $\sum c_i x_i$ where $c_i > 0$ and x_i is the size of queue i, service is exponnetial μ_i.
Result: Fixed priority, serve queue with maximal $\mu_i c_i$ (μc rule).

For exit from the rectangle with sides $[0, z_i]$, can solve PDE explicitely. For c large enough and some positive, computable α_i

$$V(x) = \min_i \alpha_i (z_i - x_i).$$

Serve the (weighted) closest to overflow.
Proof: verify via definition of viscosity-PDE.

Example 2: queues in series.

A sequence of queues, each with its own server. Upon service completion customers enter the next queue, and from the final queue they leave.
Note that overflow to any but the first queue can be prevented by stopping service to the preceding queue.

For exit from the rectangle with sides $[0, z_i]$, can solve PDE explicitely. For $c > 0$ and some positive, computable β_i

$$V(x) = \min_i b_i \cdot (z - x)$$

where $b_i = (\beta_i, \beta_i, \ldots, \beta_i, 0, \ldots, 0)$ has β_i in the first i coordinates.
The minimizer station is a "bottleneck" and must be served: service to other stations does not influence the cost.

Extensions (open):

- Relax conditions on G (partially done).

- Random routing (requires a new formulation for the reflection).

- Other costs:

$$V^n(x) := \inf_u n^{-1} \log E_x^{u,n} e^{c\tau^n}$$

where $\tau^n = \inf\{t \geq 0 : X^n(t) = 0\}$, or "Risk-seeking" cost

$$V^n(x) := \sup_u n^{-1} \log E_x^{u,n} e^{c\sigma^n}$$

- Solutions or structure of solutions to the differential game