Impatient Customers and Optimal Control

Alain Jean-Marie[1]
in collaboration with Emmanuel Hyon[2]

[1]Inria

[2]Université Paris Ouest Nanterre la Défense
LIP6

WRQ 11, Amsterdam, 31 August 2016

Introduction: the optimal control of queues

In many situations, the operation of queueing systems involves decisions...

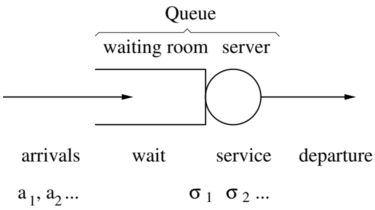

Arrivals:

- accept a customer?
- classify in a service class, priority?
- set service price

Introduction: the optimal control of queues

In many situations, the operation of queueing systems involves decisions...

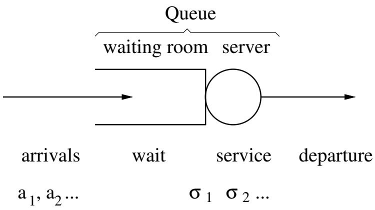

Service:

- start a service? go on a vacation?
- start/stop a server (machine, team, ...)?
- choose service speed

Introduction: the optimal control of queues

In many situations, the operation of queueing systems involves decisions...

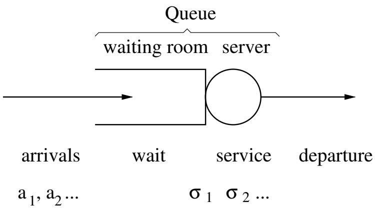

Customer:

- should I enter the queue?
- should I stay or should I go?
- how much should I pay for service?

More decision problems

See also Manufacturing Systems

- order parts? how much?
- accept order?

See also Call Centers

- add more servers?
- match customer to server?

See also (Wireless) Communications

- what packets to transmit?

See also Health Systems

- what “ressources” to match?

This presentation

- Review the Stochastic (Markovian) Optimal Control framework, which is suited for modeling some of these decision problems
- Discuss its application to some queues with impatience
- Present some advances in the methodology

Outline

Progress

Stochastic Optimal Control

The classical Stochastic Dynamic Optimal Control framework: a crash course.

The standard desciption of *Markov Decision Processes* has 6 elements:

- a state space $\mathcal{S}$;
- action spaces $\mathcal{A}(x)$ for all $x \in \mathcal{S}$;
- transition probabilities $p(x, a, y)$, $x, y \in \mathcal{S}$, $a \in \mathcal{A}(x)$;
- costs/rewards $c(x, a)$;
- an optimization criterion, *e.g.*

$$\mathbb{E}\left[\sum_{n=0}^{\infty} \theta^n c(X_n, A_n)\right], \qquad \liminf_{T} \frac{1}{T}\mathbb{E}\left[\sum_{n=0}^{T-1} c(X_n, A_n)\right].$$

- a class of *policies*

Questions for MDP

The theory typically addresses the following issues:

- assess the existence of *optimal policies*, or else of ε-optimal ones
- determine the amount of *information* these strategy need: knowledge of time, past actions, past states, ...?
- *characterize* mathematically optimal strategies
- find formulas and/or *algorithms* to compute them
- quantify errors made when using sub-optimal *approximations* ("heuristics").

Optimality Equations

Illustration of this research program: for the expected discounted cost:

$$V(x) \;=\; \inf_{\pi \text{ policy}} \mathbb{E}_x^\pi \left[\sum_{n=0}^{\infty} \theta^n \, c(x_n, a_n) \right] \; .$$

Bellman Equations

Under appropriate conditions, the (optimal) value function V is the unique solution to the equation: for all state x,

$$V(x) \;=\; \min_{a \in A(x)} \left\{ c(x,a) + \theta \sum_y p(x,a,y) V(y) \right\} .$$

Optimality Equations (ctd.)

Markov policies depend only on the current state.

Synthesis of control

Any Markov deterministic policy γ such that:

$$\gamma(x) \;\in\; \arg\min_{a \in A(x)} \left\{ c(x,a) + \theta \sum_y p(x,a,y) V(y) \right\}$$

is optimal.

Fixed points and iterations

The value function is the fixed point of a non-linear operator, the dynamic programming operator:

$$V = TV.$$

Value Iteration

Let V_0 be a function from $\mathcal{S}$ to $\mathbb{R}$. Consider the sequence of functions

$$V_{k+1} = TV_k.$$

This sequence converges to the value function.

This property is extremely useful:

- theoretically
- algorithmically

Progress

The Model

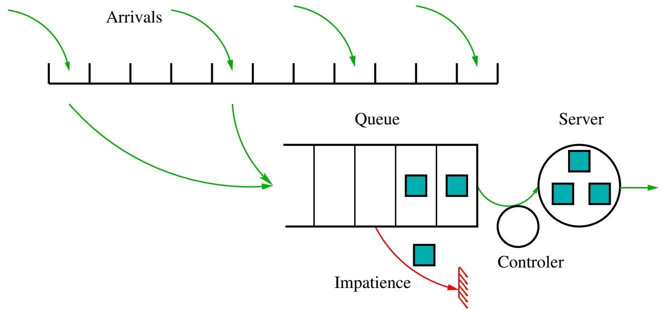

A discrete-time batch queue with geometric patience.

Model elements

Arrival

- Customers arrive to an infinite-buffer queue.
- Time is discrete.
- The distribution of arrivals in each slot A_t, arbitrary with mean λ (customers/slot), i.i.d.

Services

- Service occurs by *batches* of size B.
- Service time is one slot.

Model elements (ctd)

Deadline

Customers are *impatient*: they may leave before service.

- the individual probability of being impatient in each slot: α
- memoryless, geometrically distributed patience

Control

Service is *controlled*.

- The controller knows the number of customers but not their amount of patience: just the distribution.
- It decides whether to serve a batch or not.

The Question

What is the optimal *policy* π^* of the controller, so as to minimize the θ-*discounted* global cost:

$$v_\theta^\pi(x) \;=\; \mathbb{E}_x^\pi\left[\sum_{n=0}^{\infty} \theta^n \, c(x_n, q_n)\right] \;,$$

where:

- x_n: number of customers at step n;
- q_n: decision taken at step n;

and $c(x, q)$ is the cost incurred, involving:

- c_B: cost for serving a batch (*setup cost*)
- c_H: per capita *holding cost* of customers
- c_L: per capita *loss cost* of impatient customers.

Related Literature I

Control of queues and/or impatience (or reneging, abandonment) has a long history.

Optimal, deadline-based scheduling:

- Bhattacharya & Ephremides, 1989
- Towsley & Panwar, 1990

Optimal admission/service control (without impatience)

- Deb & Serfozo, 1973
- Altman & Koole, 1998 (admission)
- Papadaki & Powell, 2002 (service)

Optimal *routing* control with impatience

- Movaghar, 2005

Related Literature II

Optimal *service* control with impatience

- Koçaga & Ward, 2010
- Benjaffar & al., 2010 (inventory control)
- Larrañaga, Boxma, Núñez-Qeija and Squillante, 2015

Structure analysis of retrial queues

- Bhulai, Brooms and Spieksma, 2014

Adding to the state of the art...

Absent from the literature: optimal control of (finite) batch service in presence of stochastic impatience, with nonzero batch cost, discrete-time or continuous-time.

In the talk, we:

- give the solution to this problem, discrete-time, for $B = 1$
- explain what goes wrong when $B \geq 2$
- give the solution to this problem, continuous-time.

State dynamics

x_n: number of customers in the queue at time n.
$q_n = 1$ is service occurs, $q_n = 0$ if not, at time n.

Sequence of events (at each slot)

1. Begining of the slot
2. Admission in service
3. Impatience on remaining customers
4. Arrivals

State dynamics (ctd.)

The sequence of events leads to :

$$x_{n+1} \;=\; S\left([x_n - q_n B]^+\right) \;+\; A_{n+1} \;.$$

$S(x)$: the (random) number of "survivors" after impatience, out of x customers initially present.

$I(x)$: the number of impatient customers.
$\Longrightarrow$ binomially distributed random variables

Costs

The cost at step n is:

$$c_B q_n \; + \; c_L l([x_n - q_n B]^+) \; + \; c_H [x_n - q_n B]^+$$

Average Cost

$$c(x, q) \; = \; q\, c_B + (c_L\, \alpha + c_H)\,(x - qB)^+ \; = \; q\, c_B + c_Q\,(x - qB)^+ \; .$$

Optimization criterion:

$$v_\theta^\pi(x) \; = \; \mathbb{E}_x^\pi \left[\sum_{n=0}^{\infty} \theta^n \, c(x_n, q_n) \right] \; .$$

Dynamic programming equation

The optimal *value function* $V(x)$ is solution to:

The dynamic programming equation

$$V(x) = \min_{q \in \{0,1\}} \{c_B q + c_Q [x - Bq]^+ + \theta \mathbb{E} \left(V(S([x - Bq]^+) + A) \right) \}.$$

The optimal policy is *Markovian* and feedback:

$$\pi^* = (d, d, \ldots, d, \ldots)$$

and $d(x)$ is given by:

The optimal policy

$$d(x) \ = \arg \min_{q \in \{0,1\}} \{...\}.$$

Optimality Results

Theorem

The optimal policy is of threshold type: there exists a ν such that $d(x) = 1_{\{x \geq \nu\}}$.

Theorem

Let ψ be the number defined by

$$\psi = c_B - \frac{c_Q}{1 - \overline{\alpha}\theta} .$$

Then,

1. *If $\psi > 0$, the optimal threshold is $\nu = +\infty$.*
2. *If $\psi < 0$, the optimal threshold is $\nu = 1$.*
3. *If $\psi = 0$, any threshold $\nu \geq 1$ gives the same value.*

Method of Proof

Framework: propagation of properties through the dynamic programming operator (Puterman, Glasserman & Yao).

Requirement 1 (Puterman)

$$\exists w(\cdot) \geq 0, \qquad \sup_{(x,q)} \frac{|c(x,q)|}{w(x)} \; < \; +\infty \; ,$$

$$\sup_{(x,q)} \frac{1}{w(x)} \sum_y \mathbb{P}(y|x,q) w(y) \; < \; +\infty \; ,$$

and $\forall \mu$, $0 \leq \mu < 1$, $\exists \eta$, $0 \leq \eta < 1$, $\exists J$, such that: $\forall$ J-uple of Markov Deterministic decision rules $\pi = (d_1, \ldots, d_J)$, and $\forall x$,

$$\mu^J \sum_y P_\pi(y|x) w(y) \; \leq \; \eta w(x) \; .$$

$\rightarrow$ works with $w(x) = C + c_Q x$

Method of Proof

Framework: propagation of properties through the dynamic programming operator (Puterman, Glasserman & Yao).

Requirement 2 (Puterman, Glasserman & Yao)

$\exists V^\sigma, \mathcal{D}^\sigma$

1. $v \in V^\sigma$ implies $Lv \in V^\sigma$,
2. $v \in V^\sigma$ implies there exists a decision d such that $d \in \mathcal{D}^\sigma \cap \arg\min_d L_d v$,
3. V^σ is a closed by simple convergence.

$\rightarrow$ works with:

$$V^\sigma = \{ \text{ increasing and convex } \} \text{ and}$$
$$\mathcal{D}^\sigma = \{ \text{ monotone controls } \}$$

Propagation of structure

Theorem

Let, for any function v, $\tilde{v}(x) = \min_q Tv(x, q)$. Then:

1. *If v increasing, then $\tilde{v}$ increasing*
2. *If v increasing and convex, then $\tilde{v}$ increasing convex*

Theorem

If v is increasing and convex, then $Tv(x, q)$ is submodular over $\mathbb{N} \times \mathcal{Q}$. As a consequence, $x \mapsto \arg\min_q Tv(x, q)$ is increasing.

Submodularity (Topkis, Glasserman & Yao, Puterman)

g submodular if, for any $\overline{x} \geq \underline{x} \in \mathcal{X}$ and any $\overline{q} \geq \underline{q} \in \mathcal{Q}$:

$$g(\overline{x}, \overline{q}) - g(\underline{x}, \overline{q}) \leq g(\overline{x}, \underline{q}) - g(\underline{x}, \underline{q}).$$

Optimal Threshold / 1

The system under threshold ν evolves as:

$$x_{n+1} = R_\nu(x_n) := S\left([x_n - 1_{\{x \geq \nu\}}]^+\right) + A_{n+1} .$$

A direct computation gives:

$$V_\nu(x) = \frac{c_Q}{1 - \theta\overline{\alpha}} \left(x + \frac{\theta\lambda}{1-\theta}\right) + \psi \, \Phi(\nu, x)$$

$$\Phi(\nu, x) = \sum_{n=0}^{\infty} \theta^n \mathbb{P}(R_\nu^{(n)}(x) \geq \nu)$$

$$\psi = c_B - \frac{c_Q}{1 - \overline{\alpha}\theta} .$$

Optimal Threshold / 2

Lemma

The function $\Phi(\nu, x)$ is decreasing in $\nu \geq 1$, for every x.

Proof by a coupling argument.

Finally,

- if $\psi \geq 0$, $\psi\Phi(\nu, x)$ is decreasing in ν: $\nu = +\infty$ is optimal;
- if $\psi \leq 0$, $\psi\Phi(\nu, x)$ is increasing in ν: $\nu = 1$ is optimal.

What goes wrong when $B \geq 2$

Numerical experiments and exact results in special cases reveal that:

- The value function $V(x)$ is not convex in general
- The function $TV(x, q)$ is not submodular in general

Examples with $B = 10$, $\alpha = 1/10$, $\theta = 8/10$: V not convex

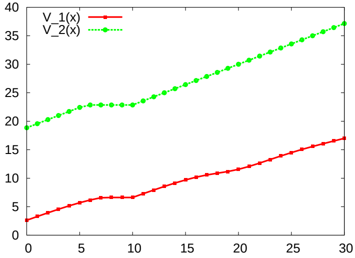

What goes wrong when $B \geq 2$, ctd.

Submodularity: if $Tv(x,1)$ is submodular, then $x \mapsto Tv(x,1) - Tv(x,0)$ is decreasing.
A counterexample with $B = 2$, $\lambda = 1/10$, $\alpha = 9/10$, $\theta = 9/10$.

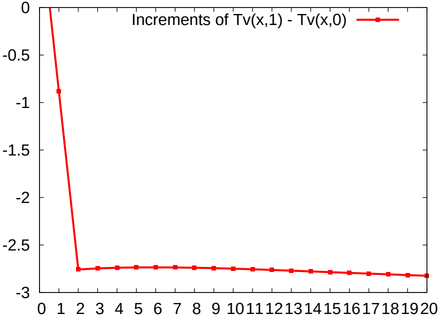

What goes wrong when $B \geq 2$, end.

Papadaki & Powell study the same problem without impatience.

Dynamics without impatience

$$x_{n+1} \;=\; [x_n - q_n B]^+ \;+\; A_{n+1} \;.$$

They show that the following "K-convexity" propagates:

***K*-convexity**

$$V(x+K) - V(x) \;\geq\; V(x-1+K) - V(x-1) \;.$$

Also used in Altman & Koole for batch arrivals.
$\Longrightarrow$ does not work here.

Extensions to the model

Average case / no discount: $\theta = 1$.
$\Longrightarrow$ should work as long as $\alpha \neq 0$ ($\overline{\alpha} \neq 1$)

Critical value:

$$\psi = c_B - c_Q \frac{1}{\alpha} = c_B - c_L - \frac{c_H}{\alpha} .$$

Branching processes: at each step, each customer is replaced by X customers. $\overline{\alpha} = \mathbb{E}X$, must be $\overline{\alpha} < \theta^{-1}$.
$\Longrightarrow$ same formula for the optimal policy

Critical value:

$$\psi = c_B - \frac{c_Q}{1 - \overline{\alpha}\theta} .$$

Progress

The model in continuous time

Consider now the queueing model with infinite buffer:

- Poisson arrivals rate Λ
- single server, exponential service durations, rate μ
- impatience rate α per customer not in service
- decision: start a service or not
- cost c_B for starting a service
- cost c_L for losing a customer by impatience
- holding cost c_H per customer in queue per unit time

Optimization criterion:

$$\mathbb{E}\left[\int_0^\infty e^{-\theta t} c_H(X_t)\mathrm{d}t + \sum_{n=0}^{\infty} e^{-\theta T_n}(c_L \mathbf{1}_{loss} + c_B \mathbf{1}_{service})\right] .$$

Optimality Equations

In order to obtain a “recursive-like” or “fixed point” equation, the trick is to go back to discrete time using an embeded process.

Value at $T_n \leftrightarrow$ value at T_{n+1}: forward reasoning with the strong Markov property.

Time-independence $\implies$ fixed point

Uniformizable models

When the set of all transition rates $|q(x, a, x)|$ is bounded, it is possible to transform the continuous-time problem into a discrete-time one. Technique attributed to Lippman (1975).
Let $\nu \geq \sup_{x,a}\{|q(x, a, x)|\}$. Define

$$\tilde{c}(x, a) = \frac{c(x, a)}{\nu + \theta}, \qquad p(x, a, y) = \frac{q(x, a, y)}{\nu}$$

and $p(x, a, x)$ to complete the transition distribution.

Uniformization equivalence

Then the optimal value and optimal policies for the discrete-time model are also optimal for the original continuous-time model.

Non-uniformizable models

What about non-uniformizable models?
Up to until quite recently:

- truncate model to “size N”
- solve for N as large as possible
- hope that the model is “reasonable”
  - ignore boundary effects
  - ignore multiplicity of solutions, discontinuities...

Numerical truncation effects occur almost always: Salch (2013), Bhulai, Brooms and Spieksma (2014), Larrañaga (2015), Blok and Spieksma (2015), ...

Non-uniformizable models

Thanks to theoretical contributions by Guo, Hernández-Lerma *et al.* and Blok, Spieksma *et al.*, the situation evolves

- validated optimality equations
- results for existence and uniqueness
- continuity results for approximated models
- smoothing technique to avoid boundary effects.

Bellman Equation for general models

Consider the controlled model with transition rates $q(x, a, y)$ and cost rates $c(x, a)$. Define $q(x, a) = \sum_{y \neq x} q(x, a, y)$.

Bellman Equation

Under appropriate conditions, the (optimal) value function V is the unique solution to the equation: for all state x,

$$V(x) = \min_{a \in A(x)} \left\{ \frac{c(x, a)}{q(x, a) + \theta} + \frac{1}{q(x, a) + \theta} \sum_{y \neq x} q(x, a, y) V(y) \right\} .$$

$$\theta V(x) = \min_{a \in A(x)} \left\{ c(x, a) + \sum_{y} q(x, a, y) V(y) \right\} .$$

Bellman Equation for general models

Consider the controlled model with transition rates $q(x, a, y)$ and cost rates $c(x, a)$. Define $q(x, a) = \sum_{y \neq x} q(x, a, y)$.

Bellman Equation, local uniformization

Let $\nu(x)$ be any function. Under the same appropriate conditions, the value function V is the unique solution to the equation: for all state x,

$$V(x) = \min_{a \in A(x)} \left\{ \frac{c(x,a)}{\nu(x)+\theta} + \frac{1}{\nu(x)+\theta} \sum_{y \neq x} q(x,a,y) V(y) + \frac{\nu(x) - q(x,a)}{\nu(x)+\theta} V(x) \right\}.$$

Bellman Equation, back to uniformizable models

Choose $\nu(x) = \nu$.

$$V(x) = \min_{a \in A(x)} \left\{ \frac{c(x,a)}{\nu + \theta} + \frac{\nu}{\nu + \theta} \sum_{y \neq x} \frac{q(x,a,y)}{\nu} V(y) + \frac{\nu}{\nu + \theta} \frac{\nu - q(x,a)}{\nu} V(x) \right\}.$$

Bellman Equation, back to uniformizable models

Choose $\nu(x) = \nu$.

$$V(x) = \min_{a \in A(x)} \Big\{ \tilde{c}(x,a) + \beta \sum_{y \neq x} p(x,a,y) V(y) + \beta p(x,a,x) V(x) \Big\}.$$

Application to the impatience queue

Bellman Equation

The value function of the problem is the unique solution to the Bellman equation:

$$V(n,0) = \min\Big\{c_B + \frac{1}{\Lambda + (n-1)\alpha + \mu + \theta}\big[k(n-1) + \Lambda V(n,1) + (n-1)\alpha V(n-2,1) + \mu V(n-1,0)\big], \frac{1}{\Lambda + n\alpha + \theta}\left[k(n) + \Lambda V(n+1,0) + n\alpha V(n-1,0)\right]\Big\}$$

for $n \geq 1$,

$$V(0,0) = \frac{1}{\Lambda + \theta}\,[k(0) + \Lambda V(1,0)],$$

$$V(n,1) = \frac{1}{\Lambda + n\alpha + \mu + \theta}\left[k(n) + \Lambda V(n+1,1) + n\alpha V(n-1,1) + \mu V(n,0)\right],$$

for $n \geq 0$.

Application to the impatience queue (ctd)

Define:

$$
\begin{aligned}
T_{AS}V(n,0) &= c_B + \frac{1}{\Lambda + (n-1)\alpha + \mu + \theta}\big[k(n-1) + \Lambda V(n,1) + (n-1)\alpha V(n-2,1) + \mu V(n-1, \\
T_{NS}V(n,0) &= \frac{1}{\Lambda + n\alpha + \theta}\left[k(n) + \Lambda V(n+1,0) + n\alpha V(n-1,0)\right]\big\}
\end{aligned}
$$

for $n \geq 1$,

$$
T_{AS}V(0,0) = T_{NS}V(0,0) = \frac{1}{\Lambda + \theta}\left[k(0) + \Lambda V(1,0)\right],
$$

$$
T_{AS}V(n,1) = T_{NS}V(n,1) = \frac{1}{\Lambda + n\alpha + \mu + \theta}\left[k(n) + \Lambda V(n+1,1) + n\alpha V(n-1,1) + \mu V(n,0)\right] ,
$$

for $n \geq 0$.

Bellman Equation, operator version

The value function of the problem is the unique solution to the Bellman equation:

$$
V = TV := \min\left\{T_{AS}V, T_{NS}V\right\} .
$$

Direct solution (mostly) fails

Idea: optimal policy is probably threshold-based.

$\Longrightarrow$ compute the value function of such policies and check whether they solve the Bellman equation... or not.

Even simpler: compute V_{AS} and V_{NS}:

- AS = Always Serve
- NS = Never Serve

Computing V_{NS}

Let $c_Q := c_H + \alpha c_L$.

Value of no service

The value of the "no service" policy is:

$$V_{NS}(n, \beta) \;=\; \frac{c_Q}{\alpha + \theta} \left(n + \frac{\Lambda}{\theta} \right) \;.$$

Optimality of no service

The "no service" policy is optimal if and only if:

$$c_B \;\geq\; \frac{c_Q}{\alpha + \theta} \;.$$

Computing V_{AS}

The function V_{AS} is defined by $V(n,1) = V(n+1,0) - c_B$ and

$$V(n,1) = \frac{1}{\Lambda + n\alpha + \mu + \theta}\Big[nc_Q + \Lambda V(n+1,1) + (n\alpha+\mu)V(n-1,1) + \mu c_B\Big].$$

$\Longrightarrow$ generating function analysis, but

$\Longrightarrow$ closed-form solution only for $\Lambda = 0$.

Solution via structure theorems

Second idea: use Value Iteration to show that

- V_{AS} has certain properties that implied it solves the Bellman Equation;
- γ_{AS} is a “limit point” of optimal policies for successive approximations.

Among these “certain properties”, one usually has monotony, convexity.

Let us see if it works.

Convexity analysis

Propagation of convexity fails!

Exemple with: $\Lambda = 0.5$, $\mu = 5$, $\alpha = 1$ and $\theta = 0.1$.

Costs: $c_B = 1.0$, $c_L = 2.0$ and $c_H = 2.0$.

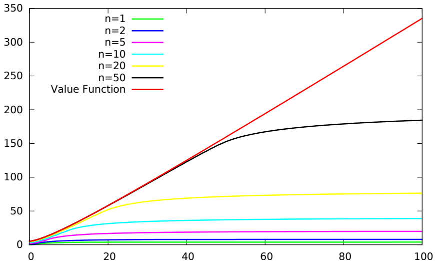

A plot of $n \mapsto V_k(n,0) := (T^{(k)} V_0)(n,0)$, for different values of k, starting with $V_0 \equiv 0$ (a convex function...).

Iterates are not convex, but the limit is.

Approximate uniformizable model I

Consider the model with:

- state-dependent arrival rate $\lambda(n)$
- state-dependent impatience rate $\alpha(n) \leq \Phi$.

Let $\nu := \Lambda + \Phi + \mu$.
Define, for $n \geq 1$:

$$
\begin{aligned}
T^{(u)}_{AS} V(n,0) &= c_B + \frac{1}{\nu+\theta}\big[(n-1)c_Q + \lambda(n-1)V(n,1) \\
&\quad + \alpha(n-1)V(n-2,1) + \mu V(n-1,0) \\
&\quad + (\nu - \lambda(n-1) - \alpha(n-1) - \mu)V(n,0)\big], \\
T^{(u)}_{NS} V(n,0) &= \frac{1}{\nu+\theta}\big[nc_Q + \lambda(n)V(n+1,0) + \alpha(n)V(n-1,0) \\
&\quad + (\nu - \lambda(n) - \alpha(n) - \mu)V(n,0)\big]
\end{aligned}
$$

Approximate uniformizable model II

$$T_{AS}^{(u)} V(0,0) = T_{NS}^{(u)} V(0,0)$$
$$= \frac{1}{\Lambda + \theta} \Lambda V(1,0)$$

$$T_{AS}^{(u)} V(n,1) = T_{NS}^{(u)} V(n,1)$$
$$= \frac{1}{\nu + \theta} \Big[n c_Q + \lambda(n) V(n+1,1) + \alpha(n) V(n-1,1) + \mu V(n,0)$$
$$+ (\nu - \lambda(n) - \alpha(n) - \mu) V(n,1) \Big] \ ,$$

for $n \geq 0$.

Approximate uniformizable model III

Bellman equation for the approximate model

The value function of the problem is the unique solution to the Bellman equation:

$$V \;=\; T^{(u)} V \;:=\; \min \left\{ T_{AS}^{(u)} V, T_{NS}^{(u)} V \right\} .$$

Let us propagate

Following Bhulai, Brooms and Spieksma (2014), we are particularly interested in:

Specific arrival/impatience functions

There exists some integer N such that:

a) The function $\alpha(\cdot)$ is given by

$$\alpha(n) \;=\; \min(n, N)\;\alpha;$$

b) The function $\lambda(\cdot)$ is given by

$$\lambda(n) \;=\; \frac{\Lambda}{N}\max(N-n, 0).$$

Let's start propagating properties!

Submodularity analysis

Even in truncated models, submodularity (partly) fails!

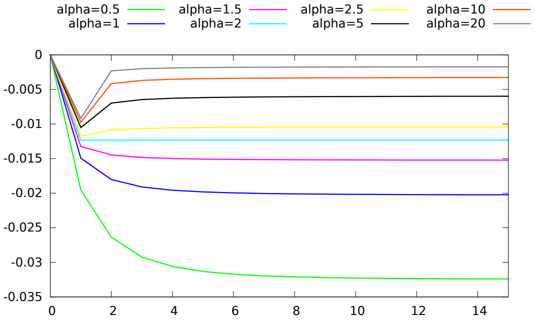

$\Lambda = 0.5$, $\mu = 2$ and $\theta = 1.5$, $c_B = 1.0$, $c_L = 2.0$ and $c_H = 2.0$. $N = 99$

A plot of $n \mapsto T_{AS} V(n,0) - T_{NS} V(n,0)$, for different values of α.
Submodularity $\iff$ this function is decreasing.

What would make AS optimal?

Submodularity is too strong. What else?

Lemma:

If the value function V_{AS} of the "always serve" (AS) policy satisfies:

$$c_B \;\leq\; \Delta_n V_{AS}(n,0) \;\leq\; \frac{c_Q}{\alpha+\theta}$$

for all $n \geq 0$, and if

$$c_B(\mu+\theta) \;\leq\; c_Q$$

then the AS policy is optimal.

What would make AS optimal? (cdt)

The function V_{AS} is defined by equations

$$V(n,0) = c_B + \frac{(n-1)c_Q + \Lambda V(n,1) + (n-1)\alpha V(n-2,1) + \mu V(n-1,0)}{\Lambda + (n-1)\alpha + \mu + \theta}$$

and $V_{AS}(n+1,0) = c_B + V_{AS}(n,1)$.
Now, V_{AS} solves the Bellman equations:

$$c_B(\Lambda+(n-1)\alpha+\mu+\theta)+(n-1)c_Q+\Lambda V(n,1)+(n-1)\alpha V(n-2,1)+\mu V(n-1,0)+\alpha V(n,0)$$

$$\leq \; nc_Q + \Lambda V(n+1,0) + n\alpha V(n-1,0) + \mu V(n,0)\ .$$

Eliminating terms $V(m,1) = V(m+1,0) - c_B$ and rearranging, this is equivalent to:

$$\begin{aligned} c_B(\mu+\theta) - c_Q + (\alpha-\mu)\Delta_n V(n-1,0) &\leq 0, \\ \underbrace{c_B(\alpha+\theta) - c_Q}_{\leq 0} + (\alpha-\mu)(\Delta_n V(n-1,0) - c_B) &\leq 0\ . \end{aligned}$$

What would make AS optimal? (end)

Observe the term $\alpha - \mu$.

Two cases:

- $\mu \leq \alpha$: it is sufficient that $\Delta_n V(n-1, 0) \geq c_B$
- $\mu \geq \alpha$: it is sufficient that $\Delta_n V(n-1, 0) \leq c_Q/(\alpha + \theta)$.

Propagable set of properties

Properties that propagate

If N large enough, the following set of properties are propagated by the Dynamic Programming operator $T^{(u)}$:

a) $n \mapsto \Delta_n V(n, 0)$ is positive and increasing for $0 \le n \le N$
b) $\Delta_n V(0, 0) \ge c_B$
c) $\Delta_n V(n, 0) \le c_Q/(\alpha + \theta)$ for all $0 \le n \le N$
d) $V(n+1, 0) = V(n, 1) + c_B$, for all $0 \le n \le N$
e) $(T^{(u)}_{NS} V)(n, 0) \ge (T^{(u)}_{AS} V)(n, 0)$ for all $0 \le n \le N$.

Necessity of smoothing

Why "for N large enough? Because:

$$
\begin{aligned}
&(\nu+\theta)[(T_{AS}^{(u)}V)(n,0)-(T_{NS}^{(u)}V)(n,0)] \\
&\quad = \quad c_B\,(\mu+\theta)-c_Q+(\alpha-\mu)\,(\Delta_n V)(n-1,0) \\
&\qquad\quad +[\lambda(n-1)-\lambda(n)]\,(\Delta_n V)(n,0)\;.
\end{aligned}
$$

Necessity of smoothing

Why "for N large enough? Because:

$$
\begin{aligned}
& (\nu+\theta)[(T_{AS}^{(u)}V)(n,0) - (T_{NS}^{(u)}V)(n,0)] \\
= \quad & c_B\,(\mu+\theta) - c_Q + (\alpha-\mu)\,(\Delta_n V)(n-1,0) \\
& + \frac{\Lambda}{N}\,(\Delta_n V)(n,0)\ .
\end{aligned}
$$

Necessity of smoothing

Why "for N large enough? Because:

$$
\begin{aligned}
&(\nu+\theta)[(T_{AS}^{(u)}V)(n,0) - (T_{NS}^{(u)}V)(n,0)] \\
&= \quad c_B\,(\mu+\theta) - c_Q + (\alpha-\mu)\,(\Delta_n V)(n-1,0) \\
&\qquad + \frac{\Lambda}{N}\,(\Delta_n V)(n,0)\ .
\end{aligned}
$$

Why not $\lambda(n) = \Lambda \mathbf{1}_{\{n \leq N\}}$? Because not convex.

Optimality of always serve

Then by the structure theorem:

Optimality for approximations

For the approximate model parametrized by N:

a) the policy "always serve" is optimal
b) $V_{AS}^{(u)}$ has the five properties above.

Next, by the continuity results of Blok and Spieksma (2015):

Optimality of always serve

The "always serve" policy is optimal if and only if:

$$c_B \leq \frac{c_Q}{\alpha + \theta} .$$

Progress

Conclusions

- Impatience (*a fortiori* retrials) challenge the established techiques for Markov Decision Processes
- Need more structural results for dynamic programming operators
  Koole (2006) and Koçağa & Ward (2010) mention the incompatibility of impatience with structure theorems.
  Blok and Spieksma (2015) argue that structure theorems are possible for smoothed/truncated approximations.
- Exploit better the multiplicity of Bellman equations satisfied by the value function
- Structural MDP analysis generally needs help for identifying properties that propagate: theory and computer tools

Open problems

Some open problems we have left along the way (for both the discrete and continuous models):

- batch service $B \geq 2$
- general (non-linear) costs
- phase-type impatience and optimal control of population models

Bibliography

References on optimal Markovian control theory

M. Puterman.
Markov Decision Processes Discrete Stochastic Dynamic Programming.
Wiley, 2005.

P. Glasserman and D. Yao.
Monotone Structure in Discrete-Event Systems.
Wiley, 1994.

X. Guo and O. Hernández-Lerma.
Continuous-Time Markov Decision Processes – Theory and Applications.
Springer, 2009.

Bibliography (ctd)

Essential surveys

G. Koole.
Monotonicity in Markov reward and decision chains: Theory and applications.
Foundation and Trends in Stochastic Systems, 1(1), 2006.

X.P. Guo, O. Hernndez-Lerma and T. Prieto-Rumeau
A Survey of Recent Results on Continuous-Time Markov Decision Processes
Top, Volume 14, Number 2, 177–257, December 2006

H. Blok and F.M. Spieksma.
Structures of optimal policies in Markov Decision Processes with unbounded jumps: the State of our Art.
Draft, December 2015.

Bibliography (ctd)

References on the control of queues

R. K. Deb and R. F. Serfozo.
Optimal control of batch service queues.
Advances in Applied Probability, 5(2):340–361, 1973.

E. Altman and G. Koole.
On submodular value functions and complex dynamic programming.
Stochastic Models, 14:1051–1072, 1998.

K. P. Papadaki and W. B. Powell.
Exploiting structure in adaptative dynamic programming algorithms for a stochastic batch service problem.
European Journal of Operational Research, 142:108–127, 2002.

Bibliography (ctd)

Control of queues with deadlines

P. P. Bhattacharya and A. Ephremides.
Optimal scheduling with strict deadlines.
IEEE Trans. Automatic Control, 34(7):721–728, July 1989.

D. Towsley and S. S. Panwar.
On the optimality of minimum laxity and earliest deadline scheduling for real-time multiprocessors.
In *Proc. IEEE EUROMICRO-90 Real Time Workshop*, pages 17–24, June 1990.

A. Movaghar.
Optimal control of parallel queues with impatient customers.
Performance Evaluation, 60:327–343, 2005.

Y. L. Koçağa and A. R. Ward.
Admission control for a multi-server queue with abandonment.
Queuing Systems, 65: 275–323, 2010.

Bibliography (end)

More Control of queues with deadlines

S. Benjaafar, J.-P. Gayon and S. Tepe.
Optimal control of a production-inventory system with customer impatience.
Operations Research Letters 38 (2010) 267–272

E. Hyon and A. Jean-Marie.
Scheduling in a queuing system with impatience and setup costs.
The Computer Journal, Volume 55, Issue 5, pp. 553–563, may 2012.
Technical Report RR-6881, version 2, INRIA, Feb. 2010.

M. Larrañaga, O. J. Boxma, R. Núñez-Queija and M.S. Squillante.
Efficient Content Delivery in the Presence of Impatient Jobs.
ITC 2015, Ghent, Belgium

M. Larrañaga.
Dynamic control of stochastic and fluid resource-sharing systems,
PhD thesis, University of Toulouse INP, 2015.

Bibliography (end)

The truncation+smoothing technique

S. Bhulai, A.C. Brooms, and Spieksma F.M.
On structural properties of the value function for an unbounded jump Markov process with an application to a processor sharing retrial queue.
Queueing Systems, 76(4):425–446, 2014.