Optimal control of service in a queue with impatience and setup costs

Emmanuel Hyon [2]

Alain Jean-Marie[1]

[1]Université Côte d'Azur/Inria,

[2]Université Paris Ouest Nanterre la Défense/ LIP6

PGMO Days, Paris, November 2017

The problem : a picture

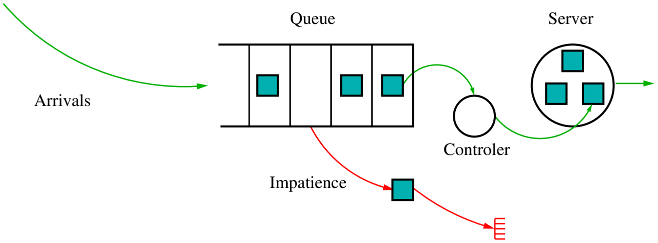

The goal

Find how to admit in service the customers

Application : Inventories with perishable Items, Scheduling with Deadlines, Call Centers, Location Data...

The problem : a queuing model with control

System

Consider the queueing model with infinite buffer:

- Poisson arrivals rate Λ
- single server, exponential service durations, rate μ
- exponential impatience durations, rate α per customer not in service

Costs

- 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

Decision: Action start a service or not, +FIFO discipline

The problem : a scheduling model

The problem can also be seen as an

on line stochastic scheduling problem

TASK=(A(w), D(w), PT(w) ; cL, cH, cB)

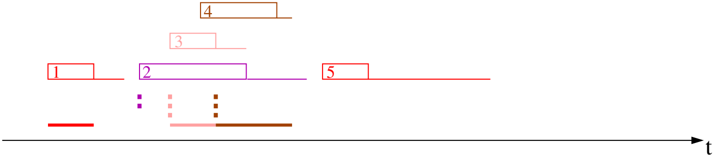

- inter arrivals have exponential durations (rate λ) ;
- exponential service and exponential impatience;
- batch cost c_B, loss cost c_L, holding cost c_H per unit time;
- Action : start the service, or not, of the first arrived present customer.
- Avaiable info to the controller (only the events)

The slotted multistage problem

The slotted time model has a constant time between decision.

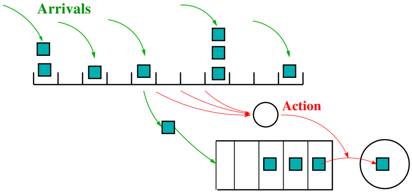

This gives a Stochastic Multistage Decision Problem

Theorem (H.-J.M. ([1] 2012))

In a slotted model the optimal policy is a threshold policy.

Proof.

Use the tool : Stochastic dynamic programming (Markov Decision Process framework). □

Continuous Time Multi Stage Stochastic Problem

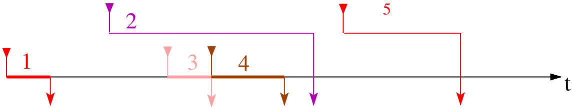

- The system evolves in a continuous time;
- Transitions occurs with random times between them;
- The "physical state" of the system follows changes randomly at transition epochs;
- At transition times one should take some *action* according a *decision*;
- The action modifies the "physical state" of the system.

A "canonical" notational system

Powell [1] [2] proposes an "*unified*" notation.

- State represents the "physical" state and some additional information required for the decision.
  Here: the state is $x = (n, \beta)$,
  - n number of customers in the queue
  - β activity of the server (0 idle, 1 busy).
- Action Space represents the set of action avaiable at the controler in a given state.
  Here: $\forall\, x \; A(x) = A = \{0, 1\}$: 1 serve, 0 no serve.
  - Decision rule a mapping π
    from the history to an action $\pi : \mathcal{H} \mapsto A$
    from the state to an action $\pi : S \mapsto A$
  - the Policy is a sequence of decision rules
    $\Pi = (\pi_1, \pi_2, \ldots, \pi_n, \ldots)$.

A “canonical” notational system Ctd

- Exogeneous Information W_t what is revealed at time t. This information can be random ($W_t(\omega_t)$) and depends on the realization of the hazard.
  Here: only occurences of departures, losses and arrival are known,
- Transition function describes the dynamic behavior. It takes state and action as inputs and gives the probability of the next states.
  Here: $\mathbb{P}(y|x, a) = \mathbb{P}(x_{t+1}|x_t, a_t)$.
- Objective function

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

A simple threshold solution

The optimal policy is threshold-based.

Theorem (Optimal policy (H. J.-M. 2017))

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

If

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

then the "never serve" (NS) policy is optimal.

If

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

then the "always serve" (AS) policy is optimal.

A structural characterization is given.

A complicated proof!

No simple proof seems to hold:

- sample path proof: fails
- stochastic optimal control $\rightarrow$ optimality equations
  - direct solution: fails
  - propagation in optimality equation: no way
  - propagation in approximate models: OK

Sample path comparisons

Intepretation of the optimal policy:

- c_B: cost of serving a customer immediately
- $\frac{c_Q}{\alpha+\theta}$: cost of never serving it, because is the value of:

$$\mathbb{E}\left[\int_0^D e^{-\theta t} c_H \mathrm{d}t + c_L e^{-\theta D}\right], \qquad D \sim \mathrm{Exp}(\alpha).$$

If the optimal policy is simple, and since it is based on individual quantities maybe the direct customer-per-customer comparison works.

$\implies$ Failure of direct sample path comparison :
It exists paths on which it is optimal to not serve at some points even though $c_B \leq c_Q/(\alpha+\theta)$.

Sample Path comparison works $\Leftrightarrow$ property is true for all scenarios.

Propagation of convexity fails !

Let T be the Bellman Operator
Example with: $\Lambda = 0.5$, $\mu = 5$, $\alpha = 1$ and $\theta = 0.1$. Costs: $c_B = 1.0$, $c_L = 2.0$ and $c_H = 2.0$.

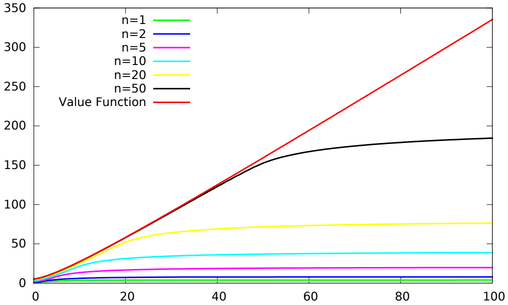

Plot of $n \mapsto V_k(n,0) := (T^{(k)}V_0)(n,0)$, for values of k, starting with $V_0 \equiv 0$ (a convex function...). Iterates not convex; limit is.

Derivation of optimality equations

The model is continuous-time, Markovian but with unbounded transition rates due to impatience and infinite state space.
$\Longrightarrow$ uniformization fails.
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 ([1] 2014), Larrañaga ([2] 2015), Blok and Spieksma ([1] 2015), ...

Non-uniformizable models

Thanks to theoretical contributions by

- Guo, Hernández-Lerma *et al.* (2009),
- Blok, Spieksma *et al.* (2015),
- Piunovskiy *et al.*,

the situation evolves

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

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 $n \geq 1$:

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

Approximate uniformizable model II

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

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

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:

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

Submodularity analysis

Another commonly found property is submodularity.
Consider: $n \mapsto T_{AS}V(n,0) - T_{NS}V(n,0)$.
Submodularity holds $\iff$ this function is decreasing $\iff$ optimal threshold policies.

Plot of this function for different values of α:

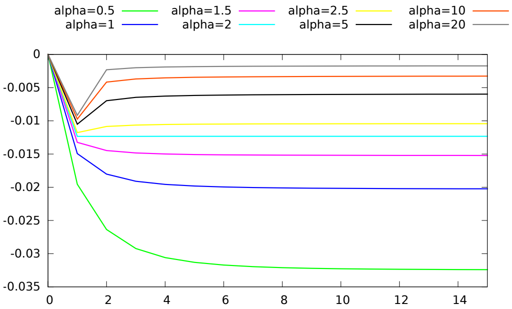

$\Lambda = 0.5$, $\mu = 2$ and $\theta = 1.5$, $c_B = 1.0$, $c_L = 2.0$ and $c_H = 2.0$. $N = 99$

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

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

Propagation framework

Theorem (Adapted from Theorem 6.11.1 in [1])

Assume that:

1. *for each $V \in \mathcal{V}$, there exists a deterministic Markov decision rule d such that $T^u = T_d v$,*
2. *$V \in \mathcal{V}_\sigma^u$ implies $T^u V \in \mathcal{V}_\sigma^u$,*
3. *$V \in \mathcal{V}_\sigma^u$ implies there exists a decision rule d such that $d \in \mathcal{D}_\sigma^u \cap \arg\min_d T_d V$,*
4. *$\mathcal{V}_\sigma^u$ is a closed subset of the set of value functions by simple convergence.*

Then, there exists an optimal stationary policy $(d^*)^\infty$ in the set of stationary Markov deterministic policies with $d^* \in \mathcal{D}_\sigma^u \cap \arg\min_d T_d V$.

Solution of approximate models

We found : 4 properties related with convexity, bounds on the increasing rate
We show : the 4 properties propagate through T^u

Using the Propagation theorem it follows:

Optimality for approximations

For the approximate model parametrized by N:

a) the policy "always serve" is optimal for $n \leq N$: there exists an optimal policy δ_N^* such that $\delta_N^*(n, 0) = AS$ for $n \leq N$.
b) $V_{AS}^{(u)}$ has the four properties above.

Limit when $N \to \infty$

We use the following result of Blok and Spieksma (2015). It applies to a collection of parametrized Markov processes $\{X(N,\delta);\ (N,\delta) \in \mathcal{N} \times \mathcal{D}\}$, on some denumerable S; $\mathcal{D} \equiv$ set of admissible stationary deterministic policies; $\mathcal{N} = \mathbb{N} \setminus \{0\} \cup \{\infty\}$.

Theorem ([1], Theorem 5.1)

Consider a collection of parametrized Markov processes $\{X(N,\delta);\ (N,\delta) \in \mathbb{N} \cup \{+\infty\} \times \mathcal{D}\}$ and cost $c : \mathbb{N} \cup \{+\infty\} \times \mathcal{D} \to \mathbb{N}$. Suppose that Assumptions below hold. Let v_N^α be the value function for the MDP $\{X(N,\delta)\}$ and δ_N^* an optimal policy. Then the following hold:

(i) $\lim_{N\to\infty} v_N^\alpha = v_\infty^\alpha$;

(ii) *any limit point of $(\delta_N^*)_{N \in \mathbb{N} \cup \{+\infty\}}$ is optimal for $X(\infty,\delta)$.*

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. Hernández-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.

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.

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

Bibliography (end)

More Control of queues with deadlines

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.

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

Stochastic optimzation modeling

Warren B. Powell.
Clearing the Jungle of Stochastic Optimization
Informs TutORials, 2014.

Warren B. Powell.
A Unified Framework for Optimization under Uncertainty
Informs TutORials, 2016.

Bibliography (end)

Continuity arguments

H. Blok and F. Spieksma.
Countable state Markov decision processes with unbounded jump rates and discounted cost: Optimality equation and approximations.
Adv. Appl. Prob., 47:1088–1107, 2015.

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.