11ème Atelier en Evaluation de Performances
Toulouse, du 15 au 17 mars 2016

Efficient content delivery in the presence of impatient customers and multiple content types

M. Larrañaga[1], O. J. Boxma[2], R. Núñez-Queija[3], M. S. Squillante[4*]

[1] L2S UMR CNRS 8506, CentraleSupelec
[2] Eurandom, The Netherlands
[3] CWI, Amsterdam, The Netherlands
[4] Mathematical Sciences Department, IBM, USA
maialen.larranaga@supelec.fr

1. Introduction

In this paper we investigate a system that combines batch services with abandonment of customers. This model consists of a multi-class M/M/1 multi-server queue with an adapted service process. This service process consists in delaying the customers that demand the delivery of the same content for batching of service. Delays due to batching come at the cost of abandonment. In particular, customers may abandon the system while waiting to be served (expiration of their deadlines), for which we penalize the system at a fixed cost per abandoning customer. Such penalties can either represent the loss of the customer or the cost of serving the customer on an expensive back-up service.

The characterization of an optimal control for the case in which customers require a different content type is out of reach. We therefore develop approximations to tackle the problem, in particular, we study Whittle's index.

Index Rules have enjoyed a great popularity, since a complex control problem whose solution might, a priori, depend on the entire state space turns out to have a strikingly simple structure. In a seminal work, Whittle introduced the so-called Restless Multi Armed Bandit Problems (RMABP), see [2]. In a RMABP all bandits (in our study a bandit is a group of customers that demand a specific type of content) in the system incur a cost, the scheduler selects one bandit to be made active, but all bandits might evolve over time according to a stochastic kernel that depends on whether the bandit was active or *frozen*. The objective is to determine the control policy that, based on the entire state-space description, selects the bandit with the objective of optimizing the average performance criterion. Whittle introduced an approximate control policy of index-type, which is nowadays referred as Whittle's index.

The model under study can be cast as an RMABP. We will therefore aim at deriving the Whittle index to obtain a heuristic for the original problem.

2. Model Description

We consider a multi-class M/M/1 queue with batch service, M servers with infinite service capacity and customers abandonment. Customers that demand content type $k \in \{1, \ldots, K\}$ arrive according to a Poisson process with rate λ_k and have an exponentially distributed service requirement with mean $1/\mu_k$, which is independent of the batch size. Customers that are waiting in the queue abandon after an exponentially distributed amount of time with mean $1/\theta_k$. Furthermore, all interarrival times, service requirements and abandonment times are independent.

Each server can only deliver one content type at a time. In every decision epoch the policy ϕ chooses whether to process the demands for a content or not. Once a customer has been admitted for service we assume that it can not abandon the system. Let $N_k^\phi(t) \in \{0, 1, \ldots\}$ denote the number of customers with a demand for content k that are waiting in the queue at time t under the policy ϕ. We will denote them class-k customers. Let $S_k^\phi(\vec{N}^\phi(t)) \in \{0, 1\}$ denote the decision with respect to class-k customers at time t under policy ϕ when there are $\vec{N}^\phi(t)$ customers present in the system, with $\vec{N}^\phi(t) = (N_1^\phi(t), \ldots, N_K^\phi(t))$. Namely, $S_k^\phi(\vec{N}^\phi(t)) = 0$ if the server does not serve class k, and $S_k^\phi(\vec{N}^\phi(t)) = 1$ if the server decides to take a class-k batch into service. Due to the infinite capacity of the server we assume that, as soon as the server is activated, *i.e.*, $S_k^\phi(\vec{N}^\phi)(t) = 1$, all customers that are waiting in the queue k initialize their service. Hence, the batch size upon activation equals the number of customers waiting in the queue, $N_k^\phi(t)$.

We assume that the service requirements are exponentially distributed with rate $\mu_k < \infty$. Upon activation of queue k the server takes a batch of size $N_k^\phi(t)$ into service, and allocates an exponentially distributed amount of time to process it. While the server is busy, new customers might arrive to the queue. In this case, the server is not allowed to take a new batch into service until service completion of

[*] Ceci est une version courte des consignes pour AEP9, compilées par Jean-Marc Vincent.

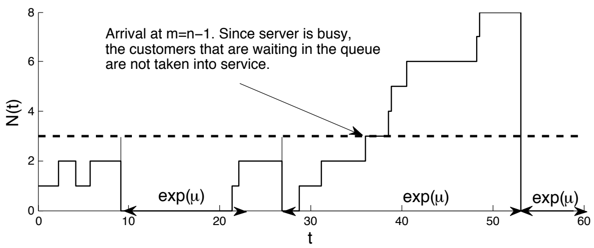

Figure 1: Simulation of process $N(t)$ under threshold $n = 3$. $exp(\mu)$ refers to the busy period of the server. $N(t)$ not only depends on n (the policy) but also on the length of each busy period.

the previous batch; see Figure 1 around $t = 37$.
Let us denote by c_k the cost per unit of time class-k customers are held in the queue, δ_k the penalty for class-k customers abandoning the queue and by c_k^s the cost per unit of time the server is busy. The objective of the present work is to find the policy ϕ so as to minimize

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T [\sum_{k=1}^K \tilde{c}_k N_k^\phi(t) + c_k^s S_k^\phi(\vec{N}^\phi(t))]\mathrm{d}t\right],$$

if $\mu < \infty$, where $\tilde{c}_k = c_k + \delta_k\theta_k$. Due to ergodicity of the system, the time-average optimal policy is equivalent to the optimal policy in steady-state, and hence we want to find ϕ such that

$$\min_\phi \sum_{k=1}^K \left(\tilde{c}_k\mathbb{E}[N_k^\phi] + c_k^s\mathbb{E}(\mathbf{1}_{\{S_k^\phi(\vec{N}^\phi)=1\}})\right), \quad (1)$$

subject to $\sum_{k=1}^K S_k^\phi(\vec{N}^\phi) \le M$. The problem described above is a Markov Decision Process and obtaining an optimal solution is out of reach. In the next section we propose a well-performing heuristic.

3. Whittle's index

The approach by Whittle is based on relaxing the original problem, allowing the constraint on the servers to be satisfied on average, that is,

$$\limsup_{T\to\infty} \frac{1}{T}\int_0^T \sum_{k=1}^K S_k^\phi(\vec{N}^\phi(t))\mathrm{d}t \le M. \quad (2)$$

This allows to decompose the original control problem into individual problems for each class of customers (we therefore drop the dependency on k from now on). Whittle's index can then be interpreted as the Lagrange multiplier of the constraint such that a given state joins the passive set.

The objective is now to determine the policy that solves (1) under Constraint (2). This can be solved by considering the uni-dimensional unconstrained control problems

$$\limsup_{T\to\infty} \left(\tilde{c}\mathbb{E}[N^\phi] + (c^s + W)\mathbb{E}(\mathbf{1}_{\{S^\phi(N^\phi)=1\}})\right), \quad (3)$$

where W is the Lagrange multiplier that can be interpreted as a subsidy for passivity.
We first prove that an optimal policy that solves (3) is of threshold type, that is, there exists n such that it is optimal not to allocate the server for all states $m \le n$ and it is optimal to allocate the server otherwise. The proof can be found in [1].

Proposition 1 *There exists $n = 0, 1, 2, \ldots$ such that $S^n(m) = 0$ for all $m \le n$ and $S^n(m) = 1$ otherwise.*

Having proven threshold type of policies to be optimal, one can define Whittle's index as described in the following algorithm.

Proposition 2 *Let $\mathcal{N}_i = \mathbb{N} \cup \{0\}\backslash\{0, \ldots, n_i\}$ for a given n_i, let $P^n = \mathbb{E}(\mathbf{1}_{\{S^n(N^n)=1\}})$ and be non-increasing, and define $\beta(\cdot)$ as follows: Define $n_0 = 0$ and*
Step i. Compute

$$\beta_i := \inf_{n\in\mathcal{N}_{i-1}} \tilde{c}\frac{\mathbb{E}(N^n) - \mathbb{E}(N^{n_{i-1}})}{P^{n_{i-1}} - P^n} - c^s, i \ge 1, \quad (4)$$

and denote by n_i the largest $n \in \mathcal{N}_{i-1}$ such that (4) is minimized and define $\beta(n) = \beta_i$ for all $n_i > n \ge n_{i-1}$. If $n_i = \infty$ stop and let $\beta(n) = \beta_i$ for all $n \ge n_i$, otherwise jump to step $i + 1$.
Then, β_i is strictly non-decreasing in i. and $\beta(n)$ defines Whittle's index.

Whittle's index policy prescribes to serve the M classes of customers with highest Whittle's index. We have numerically observed that Whittle's index policy behaves close to optimal for heavy-traffic and light-traffic regimes. The latter however has not been proven analytically.

Bibliographie

1. M. Larrañaga, O.J. Boxma, R. Nunez-Queija, M.S. Squillante – Efficient content delivery in the presence of impatient jobs. – Proceedings of ITC 2015.
2. P. Whittle – Restless bandits: Activity allocation in a changing world. – Journal of Applied Probability, 25:287-298, 1988.