From Erlang formula to robust control

Matthieu Jonckheere

Queueing Systems (2022) 100:457–459
https://doi.org/10.1007/s11134-022-09757-8

Check for updates

From Erlang formula to robust control

Matthieu Jonckheere[1]

Received: 3 February 2022 / Accepted: 28 February 2022 / Published online: 22 March 2022

1 Introduction

The historical study and dimensioning of communication systems (communication networks but also supply chains) have been successful in great part due to the existence of robust performance evaluation tools. Even when the statistical traffic characteristics are unknown (or very poorly modeled), under certain conditions, the performance evaluation of the system might only depend on first moment estimations. The first theoretical evidence in that direction was given in 1917 with the celebrated Erlang formula [3] which gives the blocking probability of a set of servers (formerly telephone lines), relating in that way the quality of service to the number of available servers.

Many communication systems have nowadays reached very complex architectures and load management mechanisms needing the resolution of difficult control problems linked to resource allocations. Classical examples are bandwidth sharing problems in wired or wireless networks, while examples of more recent applications are data centers with redundancy and load balancing mechanisms. This complexity and the need of control schemes have been tackled since a few decades by research in Markov decision processes (MDP) when a Markovian model for the system is assumed to be known and reinforcement learning (RL)—the precise transitions are unknown—where learning is added to control in the model-free setting.

The robustness properties of performance evaluation formulas echo the huge literature in robust statistics and learning where researchers have looked not only at corrupted data but also at the notion of mismatch model (see the notion of minimax theory [9]). However, the theory of robustness in control of queueing systems and more generally of generalized MDP (for instance, partially observed MDP, MDP in random environment) is still to be fully developed.

Going beyond queueing theory, it is crucial to develop a minimax theory (see Sect. 3) for an extended class of MDP and RL problems where realistic statistical assumptions are much more general than the ones supposed (while incorporating these assumptions into a Markovian description would lead to an intractable model).

✉ Matthieu Jonckheere
matthieu.jonckheere@laas.fr

1 LAAS, CNRS, Toulouse, France

2 Insensitive policies

Many MDP representing queueing systems have been studied under the condition of (quasi-)reversibility of the Markov queue lengths process X_t (when service times are exponential). This property implies (and is equivalent in some cases to) the insensitivity of the stationary distribution of the queue lengths process which depends on the jobs size distribution through their mean only [1, 2]. We shall refer to policies leading to the quasi-reversibility of X_t as insensitive policies, and they were characterized in [1, 6, 7] in the context of dynamic load balancing problems.

It is interesting to note that depending on the specific problem considered (setting and cost criterion), the price of insensitivity (see Sect. 3 for a precise definition), though not characterized in general, was shown to have completely different values. For instance, it was shown for networks with a unique class of traffic and blocking that the best insensitive policies compare very accurately to optimal policies for the original MDP in terms of blocking, while they are less efficient in terms of delay [1, 2]. The performance penalization imposed by reversibility is significantly greater for multi-class networks [6, 7]. Hence, a small to moderate price has to be paid for robustness for finite state space. However, for models with infinite buffers, this price becomes very high. It was indeed proved for the problem of dynamic load balancing (according to queue lengths) that the optimal insensitive policy for any convex cost criterion is static (i.e., does not depend on the queue-lengths) if the state space is infinite [4]. It is hence much less efficient than a state-dependent but sensitive load balancing.

3 Minimax formulation for an extended queueing system control and the price of insensitivity

Consider a queueing system with Poisson arrivals where transitions depend on an action $a \in \mathscr{A}$. It could be, for instance, a routing decision in load balancing problems and/or a service policy. Consider first the case of exponential service time. We can then describe the control problem as a Markov decision process where $X_t, t \geq 0$ describes the queue lengths process, on a state space χ under the average cost criterion. To simplify, let us assume that the policies π are (possibly random) functions of X_t and that X_t is irreducible. Let $p^\pi(x)$ and $r(x)$, $x \in \chi$, denote the (unique) steady-state distribution and the mean reward in state x, respectively. The long-run expected cost is $C^E(\pi) := \sum_x r(x)p^\pi(x)$. The optimal policy satisfies the Bellman equation and is denoted π^* (see [8]), and the associated cost is denoted by $C^E = C^E(\pi^*)$.

Now let $\mathscr{P}$ be a set of service time distributions. We can extend the control problem to dynamics with distributions of service times in $\mathscr{P}$. For a fixed policy π and a fixed service time distribution (denoted θ), we obtain a new stochastic process $\tilde{X}_t = \tilde{X}_t(\theta, \pi)$ for the queue lengths which is not Markov but still has (under mild assumptions) a unique stationary measure. We denote by $C(\theta, \pi)$ the corresponding average cost. The minimum cost of the extended control on the set of distributions $\mathscr{P}$ is defined as:

$$C^{\mathscr{P}} = \inf_{\pi \in \Pi} \sup_{\theta \in \mathscr{P}} C(\theta, \pi).$$

When this cost is actually attained by a policy, we call the latter the optimal robust policy $\hat{\pi} = \arg\min_{\pi \in \Pi} \sup_{\theta \in \mathscr{P}} C(\theta, \pi)$. On the other hand, the set of insensitive policies Π^r considered in the previous section contains policies π^r such that $C(\theta, \pi^r) = C(E, \pi^r)$, for all θ (in the set of distributions of positive random variables with finite mean). We denote by C^r the optimal cost among insensitive policies $C^r = \inf_{\pi^r \in \Pi^r} C(E, \pi^r)$.

Open questions. Given the previous definitions, we can ask several crucial questions:

1. For a given control problem, can we compute or find generic bounds for $C^{\mathscr{P}}$ for a given set of service time distributions?
2. Can we bound for a given θ and $\mathscr{P}$ the price of robustness $|C^{\mathscr{P}} - \inf_\pi C(\theta, \pi)|$, and the price of insensitivity $|C^{\mathscr{P}} - C^r|$?
3. When they can be defined, do optimal reversible policies $\hat{\pi}^r$ coincide with optimal robust policies for some set of distributions $\mathscr{P}$?

To illustrate the problematic defined, let us come back to the case of the load balancing problem [1, 4] considered in the previous section where X_t is the number of customers in a set of K servers with state space χ and $r(x) = -1_{x \in B}$ where B is the set of blocking states.

- If $\chi = \mathbb{N}^K$, using [5], the optimal insensitive policies are the Bernoulli routing policies which have no state dependency. Hence, both C^E and C^r can be easily computed (at least numerically). The value of $c^{\mathscr{P}}$, however, is unknown. The optimal robust policy is also unknown.
- If χ is finite and there is a single class of traffic, using [1] the optimal insensitive policies are the so-called simple policies which can be computed explicitly. The price of robustness and insensitivity is unknown, but simulations suggest that (for the specific cost associated with blocking), the cost of insensitivity is small.

References

1. Bonald, T., Jonckheere, M., Proutière, A.: Insensitive load balancing. SIGMETRICS Perform. Eval. Rev. **32**(1), 367–377 (2004)
2. Bonald, T., Proutière, A.: Insensitive bandwidth sharing in data networks. Queue. Syst. Theory Appl. **44**(1), 69–100 (2003)
3. Erlang, A.K.: Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. Post Office Electr. Eng. J. **10**, 189–197 (1917)
4. Jonckheere, M.: Insensitive versus efficient dynamic load balancing in networks without blocking. Queue. Syst. **54**(3), 193–202 (2006)
5. Jonckheere, M., Borst, S. C.: Stability of multi-class queueing systems with state-dependent service rates. In *Proc. Valuetools'06*, (2006)
6. Jonckheere, M., Mairesse, J.: Towards an Erlang formula for multiclass networks. Queue. Syst. **66**(1), 53–78 (2010)
7. Leino, J., Virtamo, J.: Insensitive load balancing in data networks. Comput. Netw. **50**(8), 1059–1068 (2006)
8. Puterman, M.L.: Markov Decision Processes. John Wiley & Sons, Hoboken (1994)
9. Tsybakov, A.: Introduction to Nonparametric Estimation. Springer, New York (2006)