Report 99-021

Optimal control of batch service queues with finite service capacity and linear holding costs

Samuli Aalto

ISSN: 1389-2355

OPTIMAL CONTROL OF BATCH SERVICE QUEUES WITH FINITE SERVICE CAPACITY AND LINEAR HOLDING COSTS

SAMULI AALTO

ABSTRACT. We consider the optimal control problem of certain batch service queueing systems with compound Poisson arrivals and linear holding costs. The control problem involves the determination of the epochs at which the service is initiated as well as the sizes of the batches served. The service times are assumed to be independent and identically distributed, however, with a general distribution. A quite natural operating policy is to start the service as soon as the number of customers reaches some threshold and serve always as many customers as possible. Assuming infinite service capacity Deb [3] proved that under some mild conditions the optimal operating policy is of this type. In this paper we show that a similar result is valid even if the service capacity is finite. In this case the threshold is never greater than Q, the service capacity (the maximum number of customers that can be served at the same time).

1. INTRODUCTION

In this paper we consider the optimal control problem of $M^X/G(Q)/1$ batch service queueing systems with a single server, compound Poisson arrivals and general i.i.d. service times. The service capacity (i.e. the maximum number of customers that can be served together in a batch) is denoted by Q. The control problem involves the determination of the epochs T_n at which the service is initiated as well as the sizes B_n of the batches served. Costs are usually charged both for serving the customers (service costs) and for holding them in the system (holding costs). An optimal operating policy minimizes, for example, the discounted costs among all the admissible operating policies.

In a seminal paper by Deb and Serfozo [2] sufficient conditions were found for the following two types of operating policies to be optimal:

(i) Operating policy π_∞: No customers are served.
(ii) Operating policy π_x: After a service completion, as many customers as possible are served as soon as the queue length reaches a certain fixed level x.

The latter one is called a *queue length threshold policy.* As regards the holding costs, Deb and Serfozo assumed that they depend just on the number of customers in the system (but not, for example, on the times the customers have been waiting). According to [5], such holding costs are called *linear.* Deb and Serfozo further assumed that customers arrive according to a Poisson process.

In a later paper [3] Deb proved that similar optimality results are valid even when the customers arrive according to a compound Poisson process (still assuming linear

Date: June 22, 1999.

1991 *Mathematics Subject Classification.* 60K25.

Key words and phrases. Optimal control, batch service queue, finite service capacity.

holding costs). However, he only considered the case of infinite service capacity, $Q = \infty$.

In [1] we considered the case with compound Poisson arrivals and finite service capacity, $Q < \infty$. We let the holding costs be even non-linear but omitted totally the service costs. We also needed an additional (technical) assumption that the size of an arriving batch is bounded by some constant $M < \infty$. In the case of linear holding costs our results in [1] imply the following two facts:

(i) The operating policy π_∞ that leaves all the customers unserved is never optimal.
(ii) An optimal operating policy belongs to the class of queue length threshold policies π_x with threshold $x \le Q$.

It is clear that the former result is due to our assumption to omit the service costs.

In this paper we partly generalize the results of [3] and [1]. So we assume compound Poisson arrivals. We restrict ourself to the case of linear holding costs, but (as a generalization to [3]) let the service capacity be finite. In addition to holding costs, we also consider the costs due to serving customers (as a generalization to [1]). Under the same additional assumption as in [1], we will find sufficient conditions for the following two cases:

(i) The operating policy π_∞ that leaves all the customers unserved is optimal (under Condition C1).
(ii) An optimal operating policy belongs to the class of queue length threshold policies π_x with threshold $x \le Q$ (under Conditions C2 and C3).

The rest of the paper is organized as follows. In Section 2 we present the model and the main results (including conditions C1, C2 and C3). In Section 3 we prove the claim presented in case (i) above. Case (ii) is proved step-by-step in Sections 4 – 8. First, in Section 4, we prove that it is optimal to initiate the service infinitely many times. In Section 5 we introduce the so called Q-policies and show that it is sufficient to consider such policies when seeking an optimal policy. Some important properties of these Q-policies are presented in Section 6. In Section 7, we introduce the so called stationary Q-policies and find an optimal policy among these stationary Q-policies. Finally, in Section 8, this optimal stationary policy is shown to be optimal also among all the Q-policies and, thus, among all the admissible policies.

2. The model and the main results

In this section we first introduce the queueing model. The main results concerning the optimal control of this queueing system are presented at the end of the section in Theorem 2.3.

Consider an $M^X/G(Q)/1$ batch service queueing system with $Q < \infty$. In this model the service capacity is finite and customers arrive in batches, the sizes of which are independent and identically distributed. Let β_n denote the size of the nth arriving batch. As in [1], we make the following assumption.

Assumption 2.1. *We assume that there is $M < \infty$ such that $P\{\beta_1 \le M\} = 1$.*

The batches β_n arrive according to a Poisson process, the intensity of which is denoted by λ. Let $(A(t))_{t\geq 0}$ denote the customer arrival process with $A(0) = 0$. Thus,

$$A(t) = \sum_{n=1}^{\infty} \beta_n 1_{\{\tau_n \leq t\}},$$

where τ_n denotes the arrival time of the nth batch.

Customers are served in batches, the sizes of which are not greater than Q. The service times S_n are assumed to be strictly positive, independent and identically distributed with a finite mean $E[S_1] < \infty$. In particular, they are assumed to be independent of the service batches. The following assumption implies that, for example, the system with the usual operating policy π_1 (i.e. the queue length threshold policy with threshold $x = 1$) is stable.

Assumption 2.2. *We assume that* $\lambda E[\beta_1]E[S_1] < Q$.

The first service starts at time $T_0 = 0$. The size of the first service batch, B_0, need not be specified, since, according to our assumptions, the first service time S_1 is independent of B_0 and, as we will assume later, only those customers that are waiting (but not yet in service) cause some costs. The number of those customers that remain in the queue of waiting customers at time 0 is denoted by $X(0)$. We assume that $X(0) < M$. By this way, the starting time 0 looks like a non-trivial service epoch (which will be defined in Section 6). The conditional probability measure that takes x as the initial queue length ($X(0) = x$) is denoted by P_x. The corresponding conditional expectation operator is denoted by E_x.

An *operating policy* $\pi = ((T_n), (B_n))$ is defined by giving the service epochs T_n, $n \in \{0, 1, \ldots\}$, and the service batches B_n, $n \in \{1, 2, \ldots\}$. It is required that $T_0 = 0$ and $T_n \geq T_{n-1} + S_n$ for $n \geq 1$. As regards the service batches, it is required that $B_n \leq Q$ and $\sum_{k=1}^{n} B_k \leq X(0) + A(T_n)$ for all n. When needed, a more complete notation, $\pi = ((T_n^\pi), (B_n^\pi))$, is used.

If the last service is initiated at time $T_{n_0}^\pi$, we denote $T_n^\pi = \infty$ and $B_n^\pi = 0$ for all $n > n_0$. Let π_∞ denote such a policy that leaves all the customers unserved (the *non-serving* policy). Then $T_n^{\pi_\infty} = \infty$ and $B_n^{\pi_\infty} = 0$ for all $n > 0$.

An operating policy is said to be *admissible* if the decisions are based on the current and past information only. More precisely, T_n^π shall be a stopping time with respect to the history $\mathcal{F}_n$ generated by the initial queue length $X(0)$, the arrival process A and the service times $S_1, \ldots, S_n$. In addition, B_n^π shall be measurable with respect to the corresponding stopped σ-algebra $\mathcal{F}_n(T_n^\pi)$. The family of admissible operating policies is denoted by Π.

With each policy $\pi \in \Pi$, we associate a *queue length process* $(X^\pi(t))_{t\geq 0}$. The queue length at time t is defined by

$$X^\pi(t) = X(0) + A(t) - \sum_{n=1}^{\infty} B_n^\pi 1_{\{T_n^\pi \leq t\}}.$$

Note that $X^\pi(t)$ stands for the number of *waiting* customers (excluding the customers in service) at time t and $X(0)$ is the initial queue length common to all policies π.

Let $x \geq 0$. The *queue length threshold policy* $\pi_x = ((T_n), (B_n))$ with threshold x is formally defined by setting $T_0 = 0$ and, for $n \in \{1, 2, \ldots\}$,

- $T_n = \inf\{t \geq T_{n-1} + S_n \mid X_n(t) \geq x\}$ and
- $B_n = \min\{X_n(T_n), Q\}$.

Here X_n denotes the (partial) queue length process that takes into account the services up to time T_{n-1},

$$X_n(t) = X(0) + A(t) - \sum_{m=1}^{n-1} B_m 1_{\{T_m \leq t\}}.$$

The resulting policy is clearly admissible.

Holding costs (for policy π) are assumed to accumulate continuously at rate $h(X^\pi(t))$, where $h(0) = 0$ and $h(x)$ is non-decreasing as a function of x. Note that the cost rate process $h(X^\pi(t))$ is thus non-decreasing within service intervals $[T_{n-1}^\pi, T_n^\pi)$. In addition to the holding costs, at every service epoch T_n^π a service cost of amount of $K + cB_n^\pi$ is incurred, where $K \geq 0$ and $c \geq 0$ are constant.

Let $\alpha > 0$ be fixed. The *discounted cost process* $(D^\pi(t))_{t\geq 0}$ for each $\pi \in \Pi$ is defined by

$$D^\pi(t) = \int_0^t e^{-\alpha u} h(X^\pi(u))\, du + \sum_{n=1}^{\infty} e^{-\alpha T_n^\pi}(K + cB_n^\pi) 1_{\{T_n^\pi \leq t\}}.$$

So, $D^\pi(t)$ takes into account all the costs until time t. Since $D^\pi(t)$ is non-decreasing, the limit

$$D^\pi = \lim_{t\to\infty} D^\pi(t) = \int_0^\infty e^{-\alpha u} h(X^\pi(u))\, du + \sum_{n=1}^{\infty} e^{-\alpha T_n^\pi}(K + cB_n^\pi)$$

is well defined. For each $\pi \in \Pi$, the (expected) *discounted cost* with respect to the initial queue length $x < M$ is defined as follows:

$$V^\pi(x) = \lim_{t\to\infty} E_x[D^\pi(t)] = E_x[\int_0^\infty e^{-\alpha u} h(X^\pi(u))\, du + \sum_{n=1}^{\infty} e^{-\alpha T_n^\pi}(K + cB_n^\pi)]. \quad (2.1)$$

In addition, denote, for any $x < M$,

$$V^*(x) = \inf\{V^\pi(x) \mid \pi \in \Pi\}. \quad (2.2)$$

It is possible to show that the part of the discounted cost due to serving customers,

$$V_S^\pi(x) = E_x[\sum_{n=1}^{\infty} e^{-\alpha T_n^\pi}(K + cB_n^\pi)],$$

is finite for all π. Namely, by taking into account the facts that, for all π and n,

$$T_n^\pi \geq S_1 + \ldots + S_n \quad \text{and} \quad B_n^\pi \leq Q,$$

the discounted serving cost $V_S^\pi(x)$ can be upperbounded by a geometric (and, thus, finite) sum:

$$V_S^\pi(x) \leq \sum_{n=1}^{\infty} E[e^{-\alpha S_1}]^n (K + cQ) = M_S < \infty.$$

As regards the other part of the discounted cost,

$$V_H^\pi(x) = E_x[\int_0^\infty e^{-\alpha u} h(X^\pi(u))\, du],$$

which is due to holding customers, it can be finite or infinite, depending on our operating policy π and cost function $h(x)$. However, if $h(x) = hx$ with some $h > 0$, then we have, for all π and x,

$$V_H^\pi(x) < \infty,$$

implying also that, for all π and x,

$$V^\pi(x) < \infty \quad \text{and} \quad V^*(x) < \infty.$$

This can be proved by considering the non-serving policy π_∞. It is an easy exercise to calculate the discounted cost for this policy:

$$V^{\pi_\infty}(x) = V_H^{\pi_\infty}(x) = \frac{hx}{\alpha} + \frac{h\lambda E[\beta_1]}{\alpha^2} < \infty.$$

Note that this result is even independent of our assumptions 2.1 and 2.2 (as long as $E[\beta_1] < \infty$). On the other hand, since $X^{\pi_\infty}(t) \geq X^\pi(t)$ for any π and t, holding costs are greatest for π_∞. Thus, for all π and x,

$$V_H^\pi(x) \leq V_H^{\pi_\infty}(x) < \infty.$$

Let then

$$x^* = 1 + \max\{x \in \{-1, 0, \ldots, Q-1\} \mid \zeta^*(x) < z^*\}, \tag{2.3}$$

where

$$z^* = \alpha(K + V^*(0))$$

and

$$\zeta^*(x) = \begin{cases} -\infty, & x = -1, \\ h(x) + \lambda c(E[(x+\beta_1) \wedge Q] - x) + & \\ \quad \lambda(E[V^*((x+\beta_1-Q)^+)] - V^*(0)) - \alpha c x, & x \in \{0, 1, \ldots, Q-1\}. \end{cases}$$

Here we have used notation: $x \wedge Q = \min\{x, Q\}$ and $(x-Q)^+ = \max\{x-Q, 0\}$. We further note that $x^* \in \{0, 1, \ldots, Q\}$. In the special case of ordinary Poisson arrivals, we have

$$\zeta^*(x) = \begin{cases} -\infty, & x = -1, \\ h(x) + \lambda c - \alpha c x, & x \in \{0, 1, \ldots, Q-1\}. \end{cases}$$

Now we are ready to present our main results. After presenting Conditions C1, C2 and C3, we first consider the general case, i.e. non-linear cost rate functions $h(x)$, in Theorem 2.3. Then, in Corollary 2.4, we assume that the holding cost rate function $h(x)$ is linear.

Condition C1: $h(x+1) - h(x) \leq \alpha(c + \frac{K}{Q})$ for all $x \geq 0$.

Condition C2: $h(x+1) - h(x) \geq \alpha(c + \frac{K}{Q})$ for all $x \geq 0$.

Condition C3: $V^*(x) < \infty$ for all $x < M$.

Note that, under Condition C1, $h(x) \leq \alpha(c + \frac{K}{Q})x$ for all x. Thus, $V^{\pi_\infty}(x) < \infty$ for all x, implying also that $V^*(x) < \infty$ for all x. Note further that Conditions C1 and C2 are complementary only when the cost rate function $h(x)$ is linear. These conditions are slightly different from those presented in [2] and [3]. The difference is due to the fact that we excluded the customers in the service when considering the holding costs, whereas Deb and Serfozo included them.

Theorem 2.3. *Let $h(x)$ be any non-decreasing function with $h(0) = 0$.*

(i) *Under Condition C1, the non-serving policy π_∞ is discounted cost optimal.*
(ii) *Under Conditions C2 and C3, the queue length threshold policy π_{x^*} is discounted cost optimal.*

Corollary 2.4. *Assume that $h(x) = hx$ with some $h > 0$. Then there are two possibilities:*

(i) *If $h \le \alpha(c + \frac{K}{Q})$, then the non-serving policy π_∞ is discounted cost optimal.*
(ii) *If $h \ge \alpha(c + \frac{K}{Q})$, then the queue length threshold policy π_{x^*} is discounted cost optimal.*

Finally, by omitting the service costs (i.e. $K = c = 0$), we get the following result, which is in line with our former results presented in [1]. In this case, we have $z^* = \alpha V^*(0)$ and

$$\zeta^*(x) = \begin{cases} -\infty, & x = -1, \\ h(x) + \lambda(E[V^*((x+\beta_1 - Q)^+)] - V^*(0)), & x \in \{0, 1, \dots, Q-1\}. \end{cases}$$

Corollary 2.5. *Let $K = c = 0$. If $V^*(x) < \infty$ for all $x < M$, then the queue length threshold policy π_{x^*} is discounted cost optimal.*

3. Optimal policy: do not serve at all!

In this section we prove the first part (i) of Theorem 2.3. Therefore, we assume that Condition C1 is valid. We will prove that, under this assumption, the non-serving policy π_∞ is discounted cost optimal even pathwise.

Proposition 3.1. *$D^{\pi_\infty} \le D^\pi$ for all $\pi \in \Pi$.*

Proof. Let $\pi \in \Pi$. Denote (here) briefly: $T_n^\pi = t_n$ and $B_n^\pi = b_n$ for all n. Now, by C1, we have

$$\begin{aligned} D^{\pi_\infty} - D^\pi &\le \sum_{n=1}^{\infty} \left(\int_{t_n}^{\infty} e^{-\alpha u}\, \alpha(c + \tfrac{K}{Q}) b_n \, du - e^{-\alpha t_n}(K + c b_n) \right) \\ &= \sum_{n=1}^{\infty} e^{-\alpha t_n} \left(c b_n + \tfrac{K}{Q} b_n - K - c b_n \right) \le 0. \end{aligned}$$

The last inequality above follows from the fact that $Q \ge b_n$ for all n. □

The first part (i) of Theorem 2.3 follows by taking the expectations.

4. Optimal policy: serve infinitely many times

In this section we start the proof of the second part (ii) of Theorem 2.3. Therefore, we assume that Condition C2 is valid. We will show that, under this assumption, it is sufficient to consider such operating policies that initiate the service infinitely many times.

Let $\pi \in \Pi$ such that $T_n^\pi = \infty$ for some $n \in \{1, 2, \dots\}$, and denote

$$n_0 = \sup\{n > 0 \mid T_n^\pi < \infty\}.$$

Define a modified policy $\tilde{\pi}$ by setting $T_0^{\tilde{\pi}} = 0$, and, for $n \in \{1, 2, \dots, n_0\}$,

- $T_n^{\tilde{\pi}} = T_n^\pi$ and
- $B_n^{\tilde{\pi}} = B_n^\pi$,

and, for $n \in \{n_0 + 1, n_0 + 2, \dots\}$,

- $T_n^{\tilde{\pi}} = \inf\{t \geq T_{n-1}^{\tilde{\pi}} + S_n \mid X_n^{\tilde{\pi}}(t) \geq Q\}$ and
- $B_n^{\tilde{\pi}} = Q$.

Thus, $\tilde{\pi}$ is identical to π up to the service epoch $T_{n_0}^\pi$, but changes thereafter to the queue length threshold rule (with threshold Q). It is clear from the construction that the resulting policy is admissible, i.e. $\tilde{\pi} \in \Pi$.

Proposition 4.1. *$D^{\tilde{\pi}} \leq D^\pi$.*

Proof. Denote (here) briefly: $T_n^{\tilde{\pi}} = t_n$. Now, by C2, we have

$$\begin{aligned} D^{\tilde{\pi}} - D^\pi &\leq \sum_{n=n_0+1}^{\infty} \left(e^{-\alpha t_n}(K + cQ) - \int_{t_n}^{\infty} e^{-\alpha u}\, \alpha(c + \tfrac{K}{Q})Q\, du \right) \\ &= \sum_{n=n_0+1}^{\infty} e^{-\alpha t_n}\,(K + cQ - cQ - \tfrac{K}{Q}\, Q) = 0, \end{aligned}$$

which proves the claim. □

Proposition 4.1 tells that, under Condition C2, we can restrict ourself to such operating policies $\pi \in \Pi$ for which $T_n^\pi < \infty$ for all n.

5. Optimal policy: a Q-policy

In this section we continue the proof of the second part (ii) of Theorem 2.3. Therefore, we again assume that Condition C2 is valid. We first recall the definition of a Q-policy from [1], and then show that it is sufficient to consider only such operating policies when seeking an optimal policy.

Definition 5.1. *An operating policy $\pi \in \Pi$ is said to be a Q-*policy *if, for all $n \in \{1, 2, \dots\}$,*

- *$T_n^\pi \leq \inf\{t \geq T_{n-1}^\pi + S_n \mid X_n^\pi(t) \geq Q\}$ and*
- *$B_n^\pi = \min\{X_n^\pi(T_n^\pi), Q\}$.*

The class of such policies is denoted by Π^Q.

Note that these policies apply the following two principles: (i) after a service completion, a new service starts at the latest when the number of waiting customers reaches or exceeds the service capacity Q, and (ii) whenever a service of a batch is initiated, it includes as many customers as possible. It is also clear that, for example, all the queue length threshold policies π_x with threshold $x \leq Q$ belong to this class. These are called the *queue length threshold Q-policies.*

Let $\pi \in \Pi$ such that $T_n^\pi < \infty$ for all n. Our purpose is now to construct a Q-policy, which is (even pathwise) at least as good as π. This is done in two phases.

Define first a modified policy π^q by setting $T_0^{\pi^q} = 0$, and, for $n \in \{1, 2, \dots\}$,

- $T_n^{\pi^q} = T_n^{\pi}$ and
- $B_n^{\pi^q} = \max\{X_n^{\pi^q}(T_n^{\pi^q}), Q\}$.

Thus, π^q is identical to $\tilde{\pi}$ as regards the service epochs but serves always as many customers as possible. It is clear from the construction that the resulting policies are admissible, i.e. $\pi^q \in \Pi$. In addition, it is easy to see that

$$\sum_{k=1}^{n} B_k^{\pi^q} \geq \sum_{k=1}^{n} B_k^{\pi}$$

for all n, implying (due to simultaneous service epochs) that $X^{\pi^q}(t) \leq X^{\pi}(t)$ for all $t \geq 0$.

Proposition 5.2. $D^{\pi^q} \leq D^{\pi}$.

Proof. Denote (here) briefly: $T_n^{\pi} = t_n$ and $B_n^{\pi^q} - B_n^{\pi} = \delta_n$ for all n. Since

$$X^{\pi}(t) = X^{\pi^q}(t) + \sum_{k=1}^{n} \delta_k$$

for all $t \in [t_n, t_{n+1})$, we have, by C2,

$$\begin{aligned}
D^{\pi^q} - D^{\pi} &\leq \sum_{n=1}^{\infty} \left(e^{-\alpha t_n} c\delta_n - \int_{t_n}^{t_{n+1}} e^{-\alpha u}\, \alpha c \sum_{k=1}^{n} \delta_k \, du \right) \\
&= \sum_{n=1}^{\infty} e^{-\alpha t_n} c\delta_n - \sum_{n=1}^{\infty} (e^{-\alpha t_n} - e^{-\alpha t_{n+1}})\, c \sum_{k=1}^{n} \delta_k \\
&= \sum_{n=1}^{\infty} e^{-\alpha t_n} c\delta_n - \sum_{k=1}^{\infty} c\delta_k \sum_{n=k}^{\infty} (e^{-\alpha t_n} - e^{-\alpha t_{n+1}}) \\
&= \sum_{n=1}^{\infty} e^{-\alpha t_n} c\delta_n - \sum_{k=1}^{\infty} e^{-\alpha t_k} c\delta_k = 0,
\end{aligned}$$

which proves the claim. □

Note that the policy π^q is not necessarily a Q-policy: it applies the second principle (ii) but not the first one (i). Therefore, something more is needed.

Define now another modified policy π^Q by setting $T_0^{\pi^Q} = 0$, and, for $n \in \{1, 2, \dots\}$,

- $T_n^{\pi^Q} = \min\{T_n^{\pi^q}, \inf\{t \geq T_{n-1}^{\pi^Q} + S_n \mid X_n^{\pi^Q} \geq Q\}\}$ and
- $B_n^{\pi^Q} = \max\{X_n^{\pi^Q}(T_n^{\pi^Q}), Q\}$.

Thus, π^Q is such a modification of π^q that applies principle (i). It is clear from the construction that the resulting policy is admissible, i.e. $\pi^Q \in \Pi$. In addition, it is easy to see that

$$B_n^{\pi^Q} = B_n^{\pi^q}$$

for all n, implying (since $T_n^{\pi^Q} \leq T_n^{\pi^q}$) that $X^{\pi^Q}(t) \leq X^{\pi^q}(t)$ for all $t \geq 0$. Note further that

$$B_n^{\pi^Q} < Q \quad \Rightarrow \quad T_n^{\pi^Q} = T_n^{\pi^q}. \tag{5.1}$$

Proposition 5.3. $D^{\pi^Q} \leq D^{\pi^q}$.

Proof. Denote (here) briefly: $T_n^{\pi^Q} = t_n$, $T_n^{\pi^q} = t'_n$ and $B_n^{\pi^Q} = B_n^{\pi^q} = b_n$ for all n. Now, first by C2 and then by (5.1),

$$\begin{aligned} D^{\pi^Q} - D^{\pi^q} &\leq \sum_{n=1}^{\infty}\left((e^{-\alpha t_n} - e^{-\alpha t'_n})(K + cb_n) - \int_{t_n}^{t'_n} e^{-\alpha u}\,\alpha(c + \tfrac{K}{Q})b_n\,du\right) \\ &\overset{(5.1)}{=} \sum_{n=1}^{\infty}\left((e^{-\alpha t_n} - e^{-\alpha t'_n})(K + cQ) - \int_{t_n}^{t'_n} e^{-\alpha u}\,\alpha(c + \tfrac{K}{Q})Q\,du\right) \\ &= \sum_{n=1}^{\infty}(e^{-\alpha t_n} - e^{-\alpha t'_n})(K + cQ - cQ - \tfrac{K}{Q}Q) = 0, \end{aligned}$$

which proves the claim. □

As a corollary of the previous two propositions we get the following theorem.

Theorem 5.4. $\inf\{V^\pi(x) \mid \pi \in \Pi^Q\} = \inf\{V^\pi(x) \mid \pi \in \Pi\}$ *for all* $x < M$.

Theorem 5.4 tells that, under Condition C2, we can restrict ourself to the class of Q-policies, Π^Q.

6. Some properties of the Q-policies

In this section we recall from [1] some important properties of the Q-policies. They are needed in the next section, where we continue the proof of Theorem 2.3. All the proofs of the results presented in this section can be found in [1].

Let $\pi \in \Pi^Q$. The principles (i) and (ii) mentioned in the previous section imply that the first non-trivial decisions are made only after such a service completion that leaves less than Q customers waiting (otherwise a new service is initiated immediately with the maximum batch size Q). Inspired by this fact, we define an increasing sequence $(N_k^\pi)_{k=0}^\infty$ of integer-valued random variables by setting $N_0^\pi = 0$ and, for $k \in \{1, 2, \ldots\}$,

$$N_k^\pi = \inf\{n > N_{k-1}^\pi \mid X_n^\pi(T_{n-1}^\pi + S_n) < Q\}. \tag{6.1}$$

The random variables N_k^π are called the *indices of non-trivial service epochs.* We recall from [1] that $E_x[N_k^\pi] < \infty$ for all $x < M$ and $k \geq 0$. The *non-trivial service epochs* $\tilde{T}_k^\pi$ are defined by setting

$$\tilde{T}_k^\pi = T_{N_k^\pi}^\pi.$$

The queue length at the non-trivial service epoch $\tilde{T}_k^\pi$ is denoted by ξ_k^π,

$$\xi_k^\pi = X^\pi(\tilde{T}_k^\pi).$$

We recall from [1] that $P_x\{\xi_k^\pi < M\} = 1$ for all $x < M$ and $k \geq 0$. In addition, we define random variables $\tilde{B}_k^\pi$ and $\tilde{S}_k^\pi$ by setting

$$\tilde{B}_k^\pi = B_{N_k^\pi}^\pi \quad \text{and} \quad \tilde{S}_k^\pi = S_{N_{k-1}^\pi + 1} + \ldots + S_{N_k^\pi}.$$

Since N_1^π is common to all $\pi \in \Pi^Q$, we may write $N_1^\pi = N_1$ and $\tilde{S}_1^\pi = \tilde{S}_1$. Furthermore, we have $T_n^\pi = S_{1,n}$ and $B_n^\pi = Q$ for all $\pi \in \Pi^Q$ and $n < N_1$. It follows that the (partial) queue length process $\tilde{X}_1 = X_{N_1}^\pi$ is common to all $\pi \in \Pi^Q$. Note further that

$$\tilde{B}_1^\pi = \min\{\tilde{X}_1(\tilde{T}_1^\pi), Q\} \quad \text{and} \quad \xi_1^\pi = \max\{\tilde{X}_1(\tilde{T}_1^\pi) - Q, 0\}. \tag{6.2}$$

Consider now the common queue length process $\tilde{X}_1$ during the interval from $\tilde{S}_1$ to the first non-trivial service epoch. First we define the process $X^Q = (X^Q(t))_{t\geq 0}$ by setting

$$X^Q(t) = \tilde{X}_1(\min\{\tilde{S}_1 + t, \tilde{T}_1^{\pi Q}\}). \tag{6.3}$$

So, X^Q is a "truncated" compound Poisson process with initial value $\tilde{X}_1(\tilde{S}_1)$. It is adapted to the history $\mathcal{F}^Q = (\mathcal{F}^Q(t))_{t\geq 0}$, where

$$\mathcal{F}^Q(t) = \tilde{\mathcal{F}}_1(\tilde{S}_1 + t).$$

Here $\tilde{\mathcal{F}}_1$ denotes the history generated by the initial queue length, the arrival process and the service times $S_1, \ldots, S_{N_1}$. Note that X^Q is a strong Markov process with integer-valued, non-decreasing and right-continuous paths. The conditional distribution of X^Q with initial value $x \in \mathbb{Z}_+$ is denoted by P_x^Q, and the corresponding expectation operator by E_x^Q.

Let B denote the Banach space of real-valued and bounded functions $f\colon \mathbb{Z}_+ \to \mathbb{R}$ with the usual supremum norm

$$||f|| = \sup\{|f(x)| \mid x \in \mathbb{Z}_+\}.$$

Define an operator $\mathcal{A}^Q\colon B \to B$ by setting, for all $f \in B$ and $x \in \mathbb{Z}_+$,

$$\mathcal{A}^Q f(x) = \lim_{h\to 0^+} \tfrac{1}{h}(E_x^Q[f(X^Q(h))] - f(x)). \tag{6.4}$$

It is called the *infinitesimal operator* of the Markov process X^Q. A straightforward calculation reveals that

$$\mathcal{A}^Q f(x) = \begin{cases} \lambda(E[f(x+\beta_1)] - f(x)), & x < Q, \\ 0, & x \geq Q. \end{cases} \tag{6.5}$$

Due to Dynkin's formula [4], we have the following important result (Proposition 4.2 in [1]).

Proposition 6.1. *For any $\pi \in \Pi^Q$, $f \in B$ and $x < M$,*

$$E_x[e^{-\alpha\tilde{T}_1^\pi} f(\tilde{X}_1(\tilde{T}_1^\pi))] = E_x[e^{-\alpha\tilde{S}_1} f(\tilde{X}_1(\tilde{S}_1)) + \int_{\tilde{S}_1}^{\tilde{T}_1^\pi} e^{-\alpha t}(\mathcal{A}^Q f(\tilde{X}_1(t)) - \alpha f(\tilde{X}_1(t)))\, dt].$$

7. Stationary Q-policies

In this section we continue the proof of the second part (ii) of Theorem 2.3. Now we assume that Conditions C2 *and* C3 are valid. We first recall the definition of a stationary Q-policy from [1]. Then we introduce some operators needed later on. The basic operator $\mathcal{T}^\pi$ defined in (7.3) is associated with the discounted costs up to the first non-trivial service epoch of a Q-policy π and the value of the given function v at the end of this interval. When applied to the discounted cost function of a stationary Q-policyπ, we get $\mathcal{T}^\pi V^\pi = V^\pi$ (Proposition 7.2). By taking the pointwise infimum over all Q-policies, we obtain another operator $\mathcal{T}$. As shown in Proposition 7.4, for any function v increasing at least with rate c, there is a queue length threshold Q-policy π^v such that $\mathcal{T}v = \mathcal{T}^{\pi^v} v$. In Proposition 7.9 we show that this operator, when limited to a certain subset of non-decreasing functions, has a unique fixed point w. Thus, $\mathcal{T}^{\pi^w} w = w$. Finally, in Theorem 7.11, we show that the queue length threshold Q-policy π^w is optimal among these stationary Q-policies.

Definition 7.1. *A Q-policy* $\pi \in \Pi^Q$ *is said to be* stationary *if* $(\xi_k^\pi, \tilde{T}_k^\pi)_{k=0}^\infty$ *is a Markov renewal sequence and the process* X^π *is semi-regenerative with respect to this Markov renewal sequence. The family of stationary Q-policies is denoted by* Π^{QS}.

It is clear that all the queue length threshold Q-policies are stationary.

Let $\mathcal{E}_M = \{0, 1, \ldots, M-1\}$. By B_M we denote the space of real-valued functions $v\colon \mathcal{E}_M \to \mathbb{R}$. Since $\mathcal{E}_M$ is finite, B_M is trivially Banach with norm

$$||v||_M = \sup\{|v(x)| \mid x \in \mathcal{E}_M\}.$$

Denote, for any $t \geq 0$,

$$\tilde{D}_1(t) = \int_0^t e^{-\alpha u} h(\tilde{X}_1(u))\,du + \sum_{n=1}^{N_1-1} e^{-\alpha S_{1,n}}(K + cQ)1_{\{S_{1,n} \leq t\}}. \tag{7.1}$$

It follows from Condition C3 and results of Section 5 that there is $\pi \in \Pi^Q$ such that $V^\pi(x) < \infty$ for all $x \in \mathcal{E}_M$. Thus,

$$E_x[D_1(\tilde{S}_1)] \leq V^\pi(x) < \infty.$$

After this observation, it is not so hard to prove that $E_x[D_1(\tilde{T}_1^{\pi Q})] < \infty$ implying that even for all $\pi \in \Pi^Q$

$$E_x[D_1(\tilde{T}_1^\pi)] < \infty. \tag{7.2}$$

For each $\pi \in \Pi^Q$, we now define an operator $\mathcal{T}^\pi\colon B_M \to B_M$ by setting, for all $v \in B_M$ and $x \in \mathcal{E}_M$,

$$\mathcal{T}^\pi v(x) = E_x[\tilde{D}_1(\tilde{T}_1^\pi) + e^{-\alpha \tilde{T}_1^\pi}(K + c\tilde{B}_1^\pi + v(\xi_1^\pi))]. \tag{7.3}$$

It follows immediately from (7.3) that, for any $\pi \in \Pi^Q$ and $u, v \in B_M$,

$$u \leq v \quad \Rightarrow \quad \mathcal{T}^\pi u \leq \mathcal{T}^\pi v. \tag{7.4}$$

Proposition 7.2. *Let* $\pi \in \Pi^{QS}$. *Then*

(i) $V^\pi \in B_M$ *and* $V^\pi = \mathcal{T}^\pi V^\pi$,

(ii) $\lim_{k\to\infty}(\mathcal{T}^\pi)^k v = V^\pi$ *for all* $v \in B_M$.

Proof. First, prove (ii) as Proposition 6.3(ii) in [1]. This can be done independent of the first claim (i). The fact that $V^\pi \in B_M$ follows now easily from (ii) and (7.2). After this, prove the latter part of (i) as Proposition 6.3(i) in [1]. □

For each $v \in B_M$, we define $v^Q \in B$ by setting

$$v^Q(x) = \begin{cases} K + c(x \wedge Q) + v((x-Q)^+), & x < Q + M, \\ 0, & x \geq Q + M, \end{cases} \tag{7.5}$$

where $x \wedge Q = \min\{x, Q\}$ and $(x - Q)^+ = \max\{x - Q, 0\}$. Note that, by (6.5),

$$\mathcal{A}^Q v^Q(x) = \begin{cases} \lambda c(E[(x+\beta_1) \wedge Q] - x) + \\ \quad \lambda(E[v((x+\beta_1-Q)^+)] - v(0)), & x < Q, \\ 0, & x \geq Q. \end{cases} \tag{7.6}$$

Now, by (6.2) and (7.5), we get the following representation for the basic operator $\mathcal{T}^\pi$:

$$\mathcal{T}^\pi v(x) = E_x[\tilde{D}_1(\tilde{T}_1^\pi) + e^{-\alpha \tilde{T}_1^\pi} v^Q(\tilde{X}_1(\tilde{T}_1^\pi))]. \tag{7.7}$$

With each $v \in B_M$, we also associate a real number z^v and a real-valued function $\zeta^v(x)$ by setting

$$z^v = \alpha(K + v(0)) \tag{7.8}$$

and

$$\zeta^v(x) = \begin{cases} -\infty, & x = -1, \\ h(x) + \mathcal{A}^Q v^Q(x) - \alpha c x, & x \in \{0, 1, \dots, Q-1\}. \end{cases} \tag{7.9}$$

By applying now Proposition 6.1 to function $v^Q(x)$ and recalling that $\tilde{X}_1(t) < Q$ for all $t \in [\tilde{S}_1, \tilde{T}_1^\pi)$, we get still another representation for the basic operator $\mathcal{T}^\pi$:

$$\mathcal{T}^\pi v(x) = E_x[\tilde{D}_1(\tilde{S}_1) + e^{-\alpha \tilde{S}_1}(K + c\tilde{X}_1(\tilde{S}_1) + v(0)) + \gamma^{\pi,v}] \tag{7.10}$$

where the first two terms on the right hand side are independent of π and

$$\gamma^{\pi,v} = \int_{\tilde{S}_1}^{\tilde{T}_1^\pi} e^{-\alpha t}(\zeta^v(\tilde{X}_1(t)) - z^v)\, dt.$$

Let then

$$I_M = \{v \in B_M \mid v(x) - v(x-1) \geq c \quad \text{for all } x \in \{1, 2, \dots, M-1\}\}.$$

Proposition 7.3. *Function ζ^v is non-decreasing for all $v \in I_M$.*

Proof. Let $v \in I_M$ and $x \in \{1, 2, \dots, Q-1\}$. Now, by C2,

$$\begin{aligned} \zeta^v(x) - \zeta^v(x-1) &= h(x) - h(x-1) - \alpha c + \lambda c(E[1_{\{x+\beta_1 \leq Q\}}] - 1) + \\ &\quad \lambda(E[(v(x+\beta_1-Q) - v(x+\beta_1-1-Q))1_{\{x+\beta_1>Q\}}]) \\ &\geq \alpha(c + \tfrac{K}{Q}) - \alpha c + \lambda c(E[1_{\{x+\beta_1 \leq Q\}}] - 1) + \\ &\quad \lambda(E[(v(x+\beta_1-Q) - v(x+\beta_1-1-Q))1_{\{x+\beta_1>Q\}}]) \\ &= \alpha\tfrac{K}{Q} + \lambda c(E[1_{\{x+\beta_1 \leq Q\}}] - 1) + \\ &\quad \lambda(E[(v(x+\beta_1-Q) - v(x+\beta_1-1-Q))1_{\{x+\beta_1>Q\}}]). \end{aligned}$$

Since $v \in I_M$, we get the following inequality:

$$\begin{aligned} \zeta^v(x) - \zeta^v(x-1) &\geq \alpha\tfrac{K}{Q} + \lambda c(E[1_{\{x+\beta_1 \leq Q\}}] - 1) + \lambda(E[c1_{\{x+\beta_1>Q\}}]) \\ &= \alpha\tfrac{K}{Q} - \lambda c(1 - P\{x+\beta_1 \leq Q\}) + \lambda c P\{x+\beta_1 > Q\} \\ &= \alpha\tfrac{K}{Q} \geq 0, \end{aligned}$$

which proves the claim. □

With each $v \in I_M$, we associate an integer-valued threshold x^v by setting

$$x^v = 1 + \max\{x \in \{-1, 0, \dots, Q-1\} \mid \zeta^v(x) < z^v\}. \tag{7.11}$$

Note that $x^v \in \{0, 1, \dots, Q\}$. Finally, for each $v \in I_M$, let π^v denote the queue length threshold policy with threshold x^v,

$$\pi^v = \pi_{x^v}. \tag{7.12}$$

Proposition 7.4. *For any $v \in I_M$ and $x \in \mathcal{E}_M$,*

$$\mathcal{T}^{\pi^v} v(x) = \inf\{\mathcal{T}^\pi v(x) \mid \pi \in \Pi^Q\}.$$

Proof. By Proposition 7.3,

$$\gamma^{\pi^v,v} = \inf\{\gamma^{\pi,v} \mid \pi \in \Pi^Q\}.$$

The claim follows now from (7.10). □

Next we define an operator $\mathcal{T}\colon I_M \to B_M$ by setting, for all $v \in I_M$ and $x \in \mathcal{E}_M$,

$$\mathcal{T}v(x) = \mathcal{T}^{\pi^v}v(x). \tag{7.13}$$

It follows from Proposition 7.4 and formula (7.4) that, if $u, v \in I_M$ such that $u \le v$, then

$$\mathcal{T}u = \mathcal{T}^{\pi^u}u \le \mathcal{T}^{\pi^v}u \le \mathcal{T}^{\pi^v}v = \mathcal{T}v. \tag{7.14}$$

Proposition 7.5. *For each $u, v \in I_M$, there is $0 < d < 1$ such that*

$$||\mathcal{T}u - \mathcal{T}v||_M \le d\,||u - v||_M.$$

Proof. Let $x \in \mathcal{E}_M$. By Proposition 7.4,

$$\begin{aligned} \mathcal{T}u(x) - \mathcal{T}v(x) &= \mathcal{T}^{\pi^u}u(x) - \mathcal{T}^{\pi^v}v(x) \\ &\le \mathcal{T}^{\pi^v}u(x) - \mathcal{T}^{\pi^v}v(x) \\ &= E_x[e^{-\alpha\tilde{T}_1^{\pi^v}}(u(\xi_1^{\pi^v}) - v(\xi_1^{\pi^v}))]. \end{aligned}$$

The claim now easily follows from the facts that $u(\xi_1^{\pi^v}) - v(\xi_1^{\pi^v}) \le ||u - v||_M$ and $\tilde{T}_1^{\pi^v} \ge \tilde{S}_1 \ge S_1 > 0$ (cf. the proof of Proposition 5.3 in [1]). □

Let then

$$I_M^* = \{v \in I_M \mid v \le \mathcal{T}v\}.$$

First we show that I_M^* is non-empty. Let $\ell \in I_M$ be defined by $\ell(x) = cx$. Note that, for all $x < Q + M$,

$$\ell^Q(x) = K + c(x \wedge Q) + c(x - Q)^+ = K + cx.$$

In addition, we have $\mathcal{A}^Q\ell^Q(x) = \lambda c E[\beta_1]$ (i.e. constant) for all $x < Q$.

Proposition 7.6. $\ell \in I_M^*$.

Proof. Let $x \in \mathcal{E}_M$. Denote (here) briefly: $\pi^\ell = \pi$. Now we have to prove that

$$\mathcal{T}^\pi\ell(x) \ge \ell(x).$$

First, for each $n \in \{1, 2, \dots\}$, let

$$B_n^\circ = \min\{X_n^\circ(T_n^\pi), Q\},$$

where X_n° denotes the (partial) queue length process that takes into account only the original $X(0)$ waiting customers (ignoring, thus, all the future arrivals) and the services up to time T_{n-1}^π,

$$X_n^\circ(t) = X(0) - \sum_{m=1}^{n-1} B_m^\circ 1_{\{T_m^\pi \le t\}}.$$

It is clear that, for all n and t,

$$X_n^\circ(t) \le X_n^\pi(t).$$

In addition, $B_n^\circ \le B_n^\pi$ and $X_n^\circ(t)$ is constant during service intervals (T_{n-1}^π, T_n^π) for all n. Finally, we define

$$\xi_1^\circ = \max\{X_{N_1}^\circ(T_1^\pi) - Q, 0\}$$

and note that $\xi_1^\circ \le \xi_1^\pi$. It follows that

$$\begin{aligned}\mathcal{T}^\pi \ell(x) &= E_x[\sum_{n=1}^{N_1}(\int_{T_{n-1}^\pi}^{T_n^\pi} e^{-\alpha t} h(X_n^\pi(t))\,dt + e^{-\alpha T_n^\pi}(K + cB_n^\pi)) + e^{-\alpha T_{N_1}^\pi}\ell(\xi_1^\pi)] \\ &\ge E_x[\sum_{n=1}^{N_1}(\int_{T_{n-1}^\pi}^{T_n^\pi} e^{-\alpha t} h(X_n^\circ(t))\,dt + e^{-\alpha T_n^\pi} cB_n^\circ) + e^{-\alpha T_{N_1}^\pi}\ell(\xi_1^\circ)] \\ &= E_x[\sum_{n=1}^{N_1}((e^{-\alpha T_{n-1}^\pi} - e^{-\alpha T_n^\pi})\frac{1}{\alpha} h(X(0) - \sum_{m=1}^{n-1} B_m^\circ) + e^{-\alpha T_n^\pi} cB_n^\circ) + e^{-\alpha T_{N_1}^\pi} c\xi_1^\circ].\end{aligned}$$

Now, by C2,

$$\begin{aligned}\mathcal{T}^\pi \ell(x) &\ge E_x[\sum_{n=1}^{N_1}((e^{-\alpha T_{n-1}^\pi} - e^{-\alpha T_n^\pi})(cX(0) - \sum_{m=1}^{n-1} cB_m^\circ) + e^{-\alpha T_n^\pi} cB_n^\circ) + e^{-\alpha T_{N_1}^\pi} c\xi_1^\circ] \\ &= E_x[cX(0)(1 - e^{-\alpha T_{N_1}^\pi}) + \sum_{n=1}^{N_1} cB_n^\circ e^{-\alpha T_{N_1}^\pi} + e^{-\alpha T_{N_1}^\pi} c\xi_1^\circ].\end{aligned}$$

By further taking into account that $\xi_1^\circ = X(0) - \sum_{n=1}^{N_1} B_n^\circ$, we finally get

$$\mathcal{T}^\pi \ell(x) \ge E_x[cX(0)] = cx = \ell(x),$$

which completes the proof. □

Proposition 7.7. *$\mathcal{T}v \in I_M^*$ for all $v \in I_M^*$.*

Proof. Let $v \in I_M^*$. The first step is to prove that $\mathcal{T}v \in I_M$. This is rather complicated but, however, can be done along similar lines as in the proof of Proposition 5.4 of [1].

First we introduce some notation. The sample space of our stochastic system consists of pairs $(x, \omega) \in \mathcal{E}_M \times \Omega$, where x refers to the initial queue length and ω gives a sample of the arrival process A and the service times $S_1, S_2, \ldots$. Thus, for any $x \in \mathcal{E}_M$ and $\pi \in \Pi$, such random variables as $T_n^\pi(x)$, $B_n^\pi(x)$ and $X^\pi(t; x)$ are well defined on Ω (meaning the nth service epoch, the nth service batch and the queue length at time t, respectively, in such realizations that start with initial queue length x).

Denote (here) briefly: $\pi^v = \pi$. Now we define a modified policy $\pi' = ((T_n^{\pi'}), (B_n^{\pi'}))$ as follows. Let $T_0^{\pi'} = 0$ and, for $n \in \{1, 2, \ldots\}$ and $x < M - 1$,

- $T_n^{\pi'}(x) = T_n^\pi(x+1)$ and
- $B_n^{\pi'}(x) = \min\{X_n^{\pi'}(T_n^{\pi'}(x); x), Q\}$.

Here $X_n^{\pi'}$ denotes the nth partial queue length process generated by π'. Finally, let $T_n^{\pi'}(M-1) = T_n^\pi(M-1)$ and $B_n^{\pi'}(M-1) = B_n^\pi(M-1)$ for $n \in \{1, 2, \ldots\}$. The resulting policy π' is clearly a Q-policy.

Now, let $x < M - 1$. Our purpose is to show that

$$\mathcal{T}^\pi v(x+1) \ge \mathcal{T}^\pi v(x) + c.$$

However, by Proposition 7.4, $\mathcal{T}^{\pi}v(x) \leq \mathcal{T}^{\pi'}v(x)$. Thus, it is sufficient to prove that

$$\mathcal{T}^{\pi}v(x+1) \geq \mathcal{T}^{\pi'}v(x) + c. \tag{7.15}$$

This will be done next.

Denote (here) by A the following event on Ω:

$$A = \{\tilde{X}_1(\tilde{S}_1(x); x) < Q-1\}.$$

So, in this set with initial queue length x, the queue length at the completion time of the $N_1(x)$th service is less than $Q-1$. But this implies that, with initial queue length $x+1$, the queue length at the same time is less than Q. Thus, $N_1(x) = N_1(x+1)$, implying that

$$\tilde{S}_1(x) = \tilde{S}_1(x+1) \quad \text{and} \quad \tilde{T}_1^{\pi'}(x) = \tilde{T}_1^{\pi}(x+1).$$

In addition, for all $t \geq 0$,

$$\tilde{X}_1(t; x+1) = \tilde{X}_1(t; x) + 1.$$

By (6.2), it follows that

$$\tilde{B}_1^{\pi}(x+1) = \tilde{B}_1^{\pi'}(x) + 1_{\{\tilde{X}_1(\tilde{T}_1^{\pi'}(x);x)<Q\}} \quad \text{and} \quad \xi_1^{\pi}(x+1) = \xi_1^{\pi'}(x) + 1_{\{\tilde{X}_1(\tilde{T}_1^{\pi'}(x);x)\geq Q\}}.$$

Since $v \in I_M$, the latter equality implies that

$$v(\xi_1^{\pi}(x+1)) \geq v(\xi_1^{\pi'}(x)) + c\,1_{\{\tilde{X}_1(\tilde{T}_1^{\pi'}(x);x)\geq Q\}}.$$

Further, by C2,

$$\begin{aligned}
&\tilde{D}_1(\tilde{T}_1^{\pi}(x+1); x+1) \\
&\quad = \tilde{D}_1(\tilde{T}_1^{\pi'}(x); x) + \int_0^{\tilde{T}_1^{\pi'}(x)} e^{-\alpha t}(h(\tilde{X}_1(t;x)+1) - h(\tilde{X}_1(t;x)))\,dt \\
&\quad \geq \tilde{D}_1(\tilde{T}_1^{\pi'}(x); x) + \int_0^{\tilde{T}_1^{\pi'}(x)} e^{-\alpha t}\alpha c\,dt \\
&\quad = \tilde{D}_1(\tilde{T}_1^{\pi'}(x); x) + c(1 - e^{-\alpha \tilde{T}_1^{\pi'}(x)}).
\end{aligned}$$

Thus, we have

$$\begin{aligned}
&E_{x+1}[\tilde{D}_1(\tilde{T}_1^{\pi}) + e^{-\alpha \tilde{T}_1^{\pi}}(K + c\tilde{B}_1^{\pi} + v(\xi_1^{\pi})); A] \\
&\quad \geq E_x[\tilde{D}_1(\tilde{T}_1^{\pi'}) + c(1 - e^{-\alpha \tilde{T}_1^{\pi'}}) + \\
&\qquad e^{-\alpha \tilde{T}_1^{\pi'}}(K + c\tilde{B}_1^{\pi'} + c\,1_{\{\tilde{X}_1(\tilde{T}_1^{\pi'})<Q\}} + v(\xi_1^{\pi'}) + c\,1_{\{\tilde{X}_1(\tilde{T}_1^{\pi'})\geq Q\}}); A] \\
&\quad = E_x[\tilde{D}_1(\tilde{T}_1^{\pi'}) + e^{-\alpha \tilde{T}_1^{\pi'}}(K + c\tilde{B}_1^{\pi'} + v(\xi_1^{\pi'})); A] + cP_x(A).
\end{aligned} \tag{7.16}$$

Consider then the complementary set

$$A^c = \{\tilde{X}_1(\tilde{S}_1(x); x) = Q-1\}.$$

Let $\hat{S}_1$ be such a random variable defined on $\mathcal{E}_M \times \Omega$ that $\hat{S}_1(x+1) = \tilde{S}_1(x)$. Note that, in A^c, with initial queue length x, the queue length at the completion time of the $N_1(x)$th service is equal to $Q-1$; but with initial queue length $x+1$, the queue length at the same time is equal to Q. Thus, $N_1(x) < N_1(x+1)$, implying that

$$\tilde{T}_1^{\pi'}(x) = T_{N_1(x)}^{\pi'}(x) = T_{N_1(x)}^{\pi}(x+1) = \hat{S}_1(x+1) = \tilde{S}_1(x).$$

Thus,

$$\tilde{B}_1^{\pi'}(x) = Q - 1 \quad \text{and} \quad \xi_1^{\pi'}(x) = 0.$$

In addition, note that, for all $t \le \hat{S}_1(x+1) = \tilde{S}_1(x)$,

$$\tilde{X}_1(t; x+1) = \tilde{X}_1(t; x) + 1.$$

By further taking into account C2, it follows that

$$\begin{aligned} &\tilde{D}_1(\hat{S}_1(x+1); x+1) \\ &\quad = \tilde{D}_1(\tilde{S}_1(x); x) + \int_0^{\tilde{S}_1(x)} e^{-\alpha t}(h(\tilde{X}_1(t;x)+1) - h(\tilde{X}_1(t;x)))\,dt + \\ &\qquad e^{-\alpha \tilde{S}_1(x)}(K + cQ) \\ &\quad \ge \tilde{D}_1(\tilde{S}_1(x); x) + \int_0^{\tilde{S}_1(x)} e^{-\alpha t}\alpha c\,dt + e^{-\alpha \tilde{S}_1(x)}(K + cQ) \\ &\quad = \tilde{D}_1(\tilde{T}_1^{\pi'}(x); x) + c\,(1 - e^{-\alpha \tilde{T}_1^{\pi'}(x)}) + e^{-\alpha \tilde{T}_1^{\pi'}(x)}(K + c\tilde{B}_1^{\pi'} + c). \end{aligned}$$

Still we observe that

$$E_{x+1}[\tilde{D}_1(\tilde{T}_1^{\pi}) + e^{-\alpha \tilde{T}_1^{\pi}}(K + c\tilde{B}_1^{\pi} + v(\xi_1^{\pi})); A^c] = E_{x+1}[\tilde{D}_1(\hat{S}_1) + e^{-\alpha \hat{S}_1}\mathcal{T}v(0); A^c].$$

Thus, we have

$$\begin{aligned} &E_{x+1}[\tilde{D}_1(\tilde{T}_1^{\pi}) + e^{-\alpha \tilde{T}_1^{\pi}}(K + c\tilde{B}_1^{\pi} + v(\xi_1^{\pi})); A^c] \\ &\quad \ge E_x[\tilde{D}_1(\tilde{T}_1^{\pi'}) + c(1 - e^{-\alpha \tilde{T}_1^{\pi'}}) + e^{-\alpha \tilde{T}_1^{\pi'}}(K + c\tilde{B}_1^{\pi'} + c + \mathcal{T}v(0)); A^c] \\ &\quad \ge E_x[\tilde{D}_1(\tilde{T}_1^{\pi'}) + c(1 - e^{-\alpha \tilde{T}_1^{\pi'}}) + e^{-\alpha \tilde{T}_1^{\pi'}}(K + c\tilde{B}_1^{\pi'} + c + v(0)); A^c] \\ &\quad = E_x[\tilde{D}_1(\tilde{T}_1^{\pi'}) + e^{-\alpha \tilde{T}_1^{\pi'}}(K + c\tilde{B}_1^{\pi'} + v(\xi_1^{\pi'})); A^c] + cP_x(A^c). \end{aligned} \tag{7.17}$$

The latter inequality above follows from the fact that $v \in I_M^*$. By summing up (7.16) and (7.17), we obtain (7.15). Thus, $\mathcal{T}v \in I_M$.

It remains to prove that $\mathcal{T}v \le \mathcal{T}(\mathcal{T}v)$. This follows from (7.14), since $v, \mathcal{T}v \in I_M$ and $v \le \mathcal{T}v$. □

Proposition 7.8. *I_M^* is complete.*

Proof. Let (v_n) be Cauchy in I_M^* and $x < M - 1$. Since B_M is complete, there is $v \in B_M$ such that $v_n \to v$. It follows that $v_n(x) \to v(x)$. Now, since $v_n(x+1) - v_n(x) \ge c$, we also have $v(x+1) - v(x) \ge c$. Thus, $v \in I_M$. It remains to prove that $v \le \mathcal{T}v$. However, this can be proved in exactly the same way as in Proposition 5.5 in [1]. □

Note that, since B_M is Banach, it follows from the previous proposition that I_M^* is closed.

Define finally the operator $\mathcal{T}^*\colon I_M^* \to I_M^*$ by setting, for all $v \in I_M^*$,

$$\mathcal{T}^* v = \mathcal{T}v. \tag{7.18}$$

Proposition 7.9. *The operator $\mathcal{T}^*$ has a unique fixed point $w \in I_M^*$. In addition, $\lim_{k\to\infty}(\mathcal{T}^*)^k v = w$ for all $v \in I_M^*$.*

Proof. Since B_M is Banach, $I_M^* \subset B_M$ is closed and $\mathcal{T}^*\colon I_M^* \to I_M^*$ is a strict contraction (Proposition 7.5), the claims follow from Banach's fixed point theorem. □

Consider then the queue length threshold Q-policy $\pi^w = \pi_{x^w}$ associated with the function $w \in I_M^*$ defined above.

Proposition 7.10. $V^{\pi^w} = w$.

Proof. This can be proved as Proposition 6.4 in [1]. □

Theorem 7.11. *The queue length threshold Q-policy $\pi^w = \pi_{x^w}$ is discounted cost optimal among the stationary Q-policies Π^{QS}, i.e., for all $x \in \mathcal{E}_M$,*

$$V^{\pi^w}(x) = \inf\{V^\pi(x) \mid \pi \in \Pi^{QS}\}.$$

Proof. This can be proved as Theorem 7.1 in [1]. For completeness, we present the proof also here.

Let $\pi \in \Pi^{QS}$. By Propositions 7.4 and 7.9,

$$\mathcal{T}^\pi w \geq \mathcal{T}^{\pi^w} w = w.$$

Thus, by (7.4),

$$(\mathcal{T}^\pi)^2 w = \mathcal{T}^\pi(\mathcal{T}^\pi w) \geq \mathcal{T}^\pi w \geq w.$$

By induction, we deduce that $(\mathcal{T}^\pi)^k w \geq w$ for all k. But this implies, by Propositions 7.2(ii) and 7.10, that

$$V^\pi = \lim_{k\to\infty} (\mathcal{T}^\pi)^k w \geq w = V^{\pi^w},$$

which completes the proof. □

8. Optimal policy: a queue length threshold policy

In this section we finalize the proof of the second part (ii) of Theorem 2.3. Therefore we again assume that Condition C2 is valid. Here we prove that the discounted cost optimal policy among the stationary Q-policies found in the previous section is also optimal among all the Q-policies. By Theorem 5.4, this completes the proof of the second part (ii) of Theorem 2.3.

Let $x \in \mathcal{E}_M$ and $\pi \in \Pi^Q$. For each $k \in \{0, 1, \dots\}$, we define a new policy π_k^* by setting $T_0^{\pi_k^*} = 0$, and, for $n \in \{1, 2, \dots, N_k^\pi\}$,

- $T_n^{\pi_k^*} = T_n^\pi$ and
- $B_n^{\pi_k^*} = B_n^\pi$,

and, for $n \in \{N_k^\pi + 1, N_k^\pi + 2, \dots\}$,

- $T_n^{\pi_k^*} = \inf\{t \geq T_{n-1}^{\pi_k^*} + S_n \mid X_n^{\pi_k^*}(t) \geq x^w\}$ and
- $B_n^{\pi_k^*} = \min\{X_n^{\pi_k^*}(T_n^{\pi_k^*}), Q\}$.

So, π_k^* is identical to π up to the non-trivial service epoch $\tilde{T}_k^\pi$ but changes thereafter to the optimal stationary rule. It is clear from the construction that the resulting policy is a Q-policy, i.e. $\pi_k^* \in \Pi^Q$ for all k.

As in [1], we may easily prove that

$$V^{\pi_0^*}(x) \leq V^\pi(x).$$

By further observing that π_0^* does not depend on the original policy π but is, in fact, the optimal stationary Q-policy π^w, we have proved the following theorem.

Theorem 8.1. *The queue length threshold Q-policy $\pi^w = \pi_{x^w}$ is discounted cost optimal among the Q-policies Π^Q, i.e., for all $x \in \mathcal{E}_M$,*

$$V^{\pi^w}(x) = \inf\{V^\pi(x) \mid \pi \in \Pi^Q\}.$$

Theorems 8.1 and 5.4 together imply that the queue length threshold Q-policy $\pi^w = \pi_{x^w}$ is discounted cost optimal among all admissible policies, i.e., for all $x \in \mathcal{E}_M$,

$$V^{\pi^w}(x) = \inf\{V^\pi(x) \mid \pi \in \Pi\}.$$

It follows that

$$w(x) = V^{\pi^w}(x) = \inf\{V^\pi(x) \mid \pi \in \Pi\} = V^*(x),$$

implying that the thresholds x^w and x^* are equal, which completes the proof of the second part (ii) of Theorem 2.3.

Acknowledgements

This work was performed during the author's half-year visit to EURANDOM, Eindhoven. The visit was supported by the Academy of Finland. The author wishes to thank Prof. Onno Boxma for the cooperation during the visit.

References

[1] Aalto S (1998) Optimal control of batch service queues with compound Poisson arrivals and finite service capacity. Math Meth Oper Res 48:317–335

[2] Deb RK, Serfozo RF (1973) Optimal control of batch service queues. Adv Appl Prob 5:340–361

[3] Deb RK (1984) Optimal control of bulk queues with compound Poisson arrivals and batch service. Opsearch 21:227–245

[4] Dynkin, EB (1965) Markov Processes, Volume 1, Springer-Verlag, Berlin

[5] Weiss HJ (1979) Optimal control of batch service queues with nonlinear waiting costs. In: Modeling and Simulation, Vol 10, Part 2, Pittsburgh, pp 305–309

Helsinki University of Technology, Laboratory of Telecommunications Technology, P.O. Box 3000, FIN-02015 HUT, Finland

E-mail address: samuli.aalto@hut.fi