Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms

Probability, Networks and Algorithms

PNA

A unifying conservation law for single server queues

U. Ayesta

REPORT PNA-E0508 SEPTEMBER 2005

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

A unifying conservation law for single server queues

ABSTRACT
In this paper we develop a conservation law for a work conserving multi-class *GI/GI/*1 queue operating under a general scheduling discipline. In the context of single-class queues, conservation laws have been obtained for both non-anticipating and anticipating disciplines with general service time distributions. In the context of multi-class queues, conservation laws have been previously obtained for (i) non-anticipating disciplines and exponential service time distribution and (ii) non-preemptive disciplines and general service time distribution. The conservation law we develop generalizes already existing conservation laws, and includes in particular popular multi-class time-sharing disciplines such as Discriminatory Processor Sharing (DPS) and Generalized Processor Sharing (GPS). In the literature, the conservation laws for single-class and multi-class queues are presented as if they were different in nature. The conservation law we develop includes existing conservation laws as special cases.

2000 Mathematics Subject Classification: 60K25
Keywords and Phrases: Conservation law, M/G/1

A unifying conservation law for single server queues

Urtzi Ayesta
CWI. Amsterdam
e-mail: ayesta@cwi.nl

September 19, 2005

Abstract

In this paper we develop a conservation law for a work conserving multi-class $GI/GI/1$ queue operating under a general scheduling discipline. In the context of single-class queues, conservation laws have been obtained for both non-anticipating and anticipating disciplines with general service time distributions. In the context of multi-class queues, conservation laws have been previously obtained for (i) non-anticipating disciplines and exponential service time distribution and (ii) non-preemptive disciplines and general service time distribution. The conservation law we develop generalizes already existing conservation laws, and includes in particular popular multi-class time-sharing disciplines such as Discriminatory Processor Sharing (DPS) and Generalized Processor Sharing (GPS). In the literature, the conservation laws for single-class and multi-class queues are presented as if they were different in nature. The conservation law we develop includes existing conservation laws as special cases.

1 Introduction

The so-called *work conservation* property is fundamental to single-server (multi-class) systems. Let us consider a single-server queue with M job classes. Let $U_k(t)$ be the unfinished work at time t of class-k jobs, $k = 1, \ldots, M$, and let $U(t) = \sum_{j=1}^{M} U_j(t)$ denote the total unfinished work in the system. The unfinished work in the system $U(t)$ is a function that has vertical jumps at arrival epochs equal in size to the corresponding service requirements of the customer and remains constant when it hits the horizontal axis. When $U(t) > 0$, the unfinished work drains with a rate equal to the service rate. We say that the scheduling discipline is work-conserving if $U(t)$ decreases at rate $1(sec/sec)$ whenever $U(t) > 0$. Then, a sample path argument shows that for any work-conserving discipline, the unfinished work in the system is the same, regardless of the scheduling discipline being deployed.

In this paper we focus on conservation laws for the time average unfinished work. We refer to Green and Stidham [13] and Sigman [26] for the derivation and application of sample-path conservation laws. Let $\overline{U}_j$ denote the time average unfinished work of class-j, $j = 1, \ldots, M$. The general conservation law described above states that for work conserving systems

$$\sum_{i=1}^{M} \overline{U}_j = \overline{U}, \tag{1}$$

where $\overline{U}$ is the total time average unfinished work. The constant $\overline{U}$ depends only on the inter-arrival and service time distributions, and its value is independent of how the server's capacity is shared among the jobs of the various classes.

The work conservation property has led to the development of the so-called work conservation laws. In the case of a single-class queue, Kleinrock [17, Section 4] proved that the expected conditional response time must satisfy an integral equation. Kleinrock's original result was obtained for the subset of non-anticipating scheduling disciplines. A scheduling discipline is said to be non-anticipating if the share of service a job gets is independent of its remaining service time. O'Donovan [20] generalized this result by deriving a conservation law for the set of scheduling disciplines that are anticipating, that is, disciplines that may use information on the remaining service time when deciding which job will be served.

For the multi-class case, Kleinrock [17, Section 3.4] found that the work-conserving property allows to obtain a linear relation between the mean unconditional response time of the various classes. Such a linear relation has been obtained for (i) non-anticipating disciplines and exponential service time distribution [7] and (ii) non-preemptive disciplines and general service time distribution [17].

The application of work conservation laws in multi-class systems have proven extremely successful in the design of optimal control policies of multi-class systems. In a seminal work Coffman and Mitrani [7] showed that in the case of exponential service time distribution, the vector of expected unconditional response times is achievable by a scheduling policy *if and only if* it satisfies a set of linear constraints. Then one can optimize a linear objective function of the performance measures over the achievable region to determine the optimal policy. This approach has led to the development of the so-called Achievable Region approach. which been proven useful in a wide variety of control problems, scheduling in multi-class queues, indexable systems and restless bandits. For instance, the well known optimality of the $c\mu$ rule can be readily obtained within the framework of the achievable region. We refer to Federgruen and Groenevelt [10], Shantikumar and Yao [25], Dacre, Glazebrook and Niño-Mora [8] and Green and Stidham [13] for more detailed discussions. The application of conservation laws in single class systems are sparser. In [18] the authors applied the conservation law in order to determine the optimum scheduling policy that minimizes a cost function that depends on the expected conditional response time and required service time. Aalto *et al.* [1] developed a framework based on the work conservation law in order to compare the expected unconditional response time of two scheduling disciplines.

For more information on conservation laws we refer to the textbooks: Gelenbe and Mitrani [12, Chapter 6], Heyman and Sobel [14, Sections 11.4-5], Wolff [29, Chapter 10] and Baccelli and Brémaud [3, Section 3.2]. See also Miyazawa [19] and references therein for more details on the derivation and applications of conservation laws. Conservation laws in queues with vacations are discussed in Boxma [5].

The derivation of conservation laws for the single-class and multi-class systems are obtained by quite different approaches (more details in Section 3). In this paper, we derive a conservation law for a multi-class $GI/GI/1$ queue with a general work conserving scheduling discipline and general required service time distribution. Provided the scheduling discipline is work conserving, that is equation (1) holds, the scheduling discipline may be anticipating or non-anticipating, preemptive or non-preemptive. We will further show that already existing conservation laws for multi-class and single-class queues can be obtained as particular cases of our more general work conservation law. Thus, the conservation law developed somewhat provides a unified view on the conservation laws for single server queues. It is worthwhile to note that for the case of single-server single-class queue and anticipating service discipline, we obtain an alternative (equivalent) expression of the conservation law that O'Donovan [20] already developed.

The remainder of the paper is organized as follows. In Section 2 we introduce the notation and main assumptions that will be used throughout the paper. In Section 3 we review the existing conservation laws for single-class and multi-class systems. In Section 4 we state the main result of this paper, i.e., a work conservation law for a multi-class queue, with general work conserving scheduling discipline and general service time distribution. In Section 5 we show that previously obtained conservation laws can be readily derived as special cases.

2 Notation and Assumptions

Throughout the paper we consider a $GI/GI/1$ queue operating under a work conserving discipline. Hence, the unfinished work in the system at time t, $U(t)$, is independent of the discipline being deployed.

2.1 Single-class

Let $F(x)$, for all $x \geq 0$, denote the required service time distribution, and let $\overline{F}(x) = 1 - F(x)$, for all $x \geq 0$, denote the complementary distribution function. Let $E[X]$ and $E[X^2]$ be the first and second moments of the required service time distribution. Let λ denote the mean arrival rate. We assume the queue is stable, i.e., $\rho = \lambda E[X] < 1$. Let $T(x)$ be the expected conditional response time of a job with required service time x, that is the expected elapsed time between the arrival and departure of a job that requires x units of service. In addition, for the analysis of anticipating disciplines, it will be necessary to define $T(u;x)$, the expected conditional time that a job with total required service time x spends in the system in order to get served u units of service, $u \leq x$. In particular, when $u = x$, $T(x;x) = T(x)$ denotes the regular expected conditional response time. We note that for the set of disciplines that are non-anticipating, $T(u;x) = T(u)$, for all $u \leq x$.

We assume that the expected unfinished work in the system is finite, that is, $\overline{U} < \infty$. In the case of Poisson arrivals, the expected unfinished work in the system can be obtained by the well-known Pollaczek-Khinchin formula,

$$\overline{U} = \overline{U}^{FCFS} = \frac{\lambda E[X^2]}{2(1-\rho)}, \tag{2}$$

which was derived under the assumption that the required service time distribution is finite, i.e., $E[X^2] < \infty$.

The set of non-anticipating, work-conserving scheduling disciplines includes among others FCFS, Processor-Sharing (PS), Foreground-Background (FB)[1] and Last Come First Served (LCFS). Important examples of anticipating disciplines are the Shortest Remaining Processing Time (SRPT) [24, 22] and the Fair Sojourn Protocol (FSP) [11]. The FSP discipline can be seen as an improved version of PS obtained when the server knows the remaining required service time of the jobs present in the system.

2.2 Multi-class

Let λ_j denote as before the mean arrival rate of class-j jobs. We assume that the compound job arrival process is renewal. As a direct consequence, the system regenerates itself at the beginning of each busy period. The service time distribution of class-j is denoted by $F_j(\cdot)$, and let $\overline{F}_j(\cdot)$ its complementary distribution. Let $E[X_j]$ and $E[X_j^2]$ denote the first and second moments of the required service time distribution, $j = 1, \ldots, M$. The load of class-j is given by $\rho_j = \lambda_j E[X_j]$, and the total load is $\rho = \sum_{i=1}^{M} \rho_i$. We assume to be in the stable regime, i.e., $\rho < 1$.

Let $T_j(x)$ be the expected conditional response time of a class-j job with service time requirement x. In addition, and bearing in mind the analysis of anticipating disciplines, let $T_j(u;x)$ denote the expected conditional time that a class-j job with total required service time x spends in the system in order to get served u units of service, $u \leq x$. In particular, when $u = x$, $T_j(x;x) = T_j(x)$, $j = 1, \ldots, M$, denotes the regular expected conditional response time. We note that for the set of disciplines that are non-anticipating, $T_j(u;x) = T_j(u)$, for all $u \leq x$.

In the analysis we make the following assumption.

Assumption 1 *The function $T_j(u;x)$, $j = 1, \ldots, M$, has a continuous partial derivative with respect to x.*

In Section 4 we argue that Assumption 1 is not very restrictive.

We denote by $\overline{T}_j$ the expected unconditional response time of class-j jobs, that is,

$$\overline{T}_j = \int_0^\infty T_j(x) dF_j(x).$$

Similar to the single-class case, the value of the total unfinished work in the system $\overline{U}$ is independent of the scheduling discipline and its value is given by the Pollaczek-Khinchin formula,

$$\overline{U} = \overline{U}^{FCFS} = \frac{\sum_{i=1}^{M} \lambda_i E[X_i^2]}{2(1-\rho)}, \tag{3}$$

where it is required that the second moment of the various classes are finite, that is, $E[X_j^2] < \infty$, $j = 1, \ldots, M$.

In the context of multi-class systems, the most popular disciplines are non-anticipating, for instance Generalized Processor Sharing (GPS) [21, 27], Discriminatory Processor Sharing (DPS) [16, 9] or Priority Disciplines are all non-anticipating disciplines.

[1] Also known as Least Attained Service (LAS)

3 Review of Conservation Laws

In this section we review already existing work conservation laws.

3.1 Single-class

Kleinrock [17, Section4] derived a conservation law that applies to the set of non-anticipating, work-conserving scheduling disciplines. The original result by Kleinrock assumed Poisson arrivals, later its result was generalized to ergodic arrivals [20, 3].

Proposition 1 [17, 20, 3] *In an ergodic stable queue, under any work conserving and non-anticipating scheduling discipline, the expected conditional response time satisfies*

$$\lambda \int_0^\infty T(x)\overline{F}(x)dx = \overline{U}. \tag{4}$$

As it has been already mentioned, in the case of Poisson arrivals the value of $\overline{U}$ is given by equation (2).

Proposition 1 places an integral constraint on the expected conditional response time. For example, let π_1 and π_2 be two non-anticipating scheduling disciplines and let us denote the expected conditional response time with each discipline as $T^{\pi_1}(x)$ and $T^{\pi_2}(x)$, respectively. Let us assume further that $T^{\pi_1}(x) \leq T^{\pi_2}(x)$ for all $x \in [0, x_0)$ and $T^{\pi_1}(x) = T^{\pi_2}(x)$ for all $x \in [x_1, \infty)$, where $x_1 > x_0$. Then, since $\overline{F}(x)$ is a non-increasing function, there exists some subinterval $S \subseteq [x_0, x_1)$ such that for all $x \in S$, $T^{\pi_1}(x) \geq T^{\pi_2}(x)$. Thus the fact that π_1 outperforms π_2 in the interval $[0, x_0)$, comes at the cost of π_2 outperforming π_1 in some other interval S.

Equation (4) can be interpreted as a generalization of Little's law $H = \lambda G$ (see Brumelle [6]) or Palm inversion formula (see [3]). Informally, the equation $H = \lambda G$ states that the time average of a stochastic process (H) is equal to the arrival rate (λ) times the average contribution of each job to the process (G).

In the case of equation (4), the time average here corresponds to $\overline{U}$ and the arrival rate is λ. The shaded area in Figure 1 represents the cumulative contribution over time to the unfinished work of a job with required service time x, arrival time a^i and departure time d^i. Under the assumptions of Proposition 1, the expected cumulative burden of a job to the unfinished work (average value of the shaded region) is $\int_0^\infty T(x)\overline{F}(x)dx$

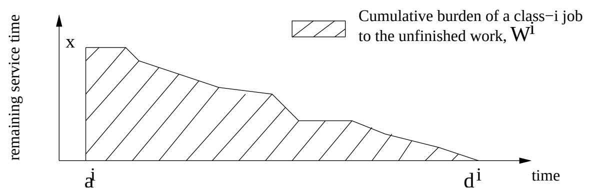

Figure 1: Contribution over time of a job on the unfinished work

For the sake of completeness, we present the following derivation based on [2] and Heyman and Sobel [14, p.426]. For similar derivations see [29, Section 10.4] or [3, p.162].

Let W^i,$i = 1, 2, \ldots$, be the cumulative burden of the i-th job on the unfinished work in the system during its complete response time (shaded region in Figure 1). Formally, $W^i := \int_{t=0}^{d^i} R^i(a^i + t)dt$, where a^i is the arrival time, d^i the departure time and $R^i(t)$ the unfinished work of the i-th customer at time t. In particular, $R^i(a^i) = x$ is the total service requirement of this job and $R^i(a^i + d^i) = 0$.

Under the conditions of Proposition 1, the sequence $\{W^n\}_{n=1}^\infty$ is a regenerative process with finite cycle lengths. Hence, the process $\{W^n\}_{n=1}^\infty$ is stationary and ergodic. Let $\tau^i(x) = d^i - a^i$ be the response time of the i-th job which requires x units of service. Let $\tau(y; x)$ denote the time needed for this job to be served y units of service, $y \leq x$. If $E[W]$ denotes the expected cumulative burden of an arbitrary job then

$$\begin{aligned} E[W] &= E[\int_{x=0}^{\infty}\int_{t=0}^{\tau^i(x)} R^i(a^i+t)dtdF(x)] \\ &= E[\int_{x=0}^{\infty}\int_{u=0}^{x} \tau^i(x-u;x)dudF(x)]. \end{aligned}$$

This corresponds to integrating the shaded area in Figure 1 either horizontally (first equation) or vertically (second equation). By a simple change of variables and interchanging the order of the integrals we obtain

$$\begin{aligned} E[W] &= E[\int_{x=0}^{\infty}\int_{u=0}^{x} \tau^i(u;x)dudF(x)] \\ &= \int_{u=0}^{\infty}\int_{x=u}^{\infty} T(u;x)dF(x)du \\ &= \int_{0}^{\infty} T(u)\overline{F}(u)du, \end{aligned} \tag{5}$$

where the last step follows by the non-anticipating property of the scheduling discipline, and thus $T(u;x) = T(u)$, for all $u \leq x$.

O'Donovan [20, Equation (9)] obtained a more general conservation law that is also valid for the set of anticipating disciplines.

Proposition 2 [20, 3] *In an ergodic stable queue, under any work conserving scheduling discipline, the following relation holds:*

$$\lambda \int_{r=0}^{\infty} r \int_{x=r}^{\infty} \partial_r S(r;x)dF(x) = \overline{U}, \tag{6}$$

where $S(r;x) = T(x) - T(x-r;x)$ *and* $\overline{U}$ *is the time-average unfinished work in the system.*

For the set of non-anticipating disciplines, we have that $\partial_r S(r;x) = dT(x-r)$. After some manipulations, it can be shown that equation (6) simplifies to equation (4). O'Donovan obtained Proposition 2 by first deriving an expression for the expected number of jobs with residual service time less than or equal to r, $M(r)$, and by noting that $\int_0^\infty r dM(r) = \overline{U}$.

3.2 Multi-class

We review now the existing conservation laws for multi-class queues. As we mentioned in the introduction, in the context of multi-class queues, conservation laws have been obtained for two particular settings: (i) non-anticipating disciplines and exponential service time distribution and (ii) non-preemptive disciplines and general service time distribution.

3.2.1 Non-anticipating discipline and exponential service time distribution

Let us consider a $GI/M/1$ multi-class queue. We assume that the service time distributions are exponential with mean $E[X_j]$, $j = 1, \ldots, M$.

Proposition 3 [7, 12] *Consider a* $GI/M/1$ *multi-class queue, under any work conserving non-anticipating scheduling discipline. The expected unconditional response times of the various classes satisfy*

$$\sum_{i=1}^{M} \rho_i \overline{T}_j = \overline{U}. \tag{7}$$

Proposition 3 is easily obtained. First note that by Little's law the expected number of jobs in the system is $\lambda_j \overline{T}_j$, $j = 1, \ldots, M$. Due to the non-anticipating feature of the discipline and the memoryless property of the exponential distribution, the expected remaining required service time of each of these jobs is precisely $E[X_j]$, $j = 1, \ldots, M$. Hence, it follows that $\overline{U}_j = \rho_j \overline{T}_j$, $j = 1, \ldots, M$. Substitution in the general work conservation law (equation 1)) leads to the desired result.

3.2.2 Non-preemptive discipline and general service time distribution

Let us now consider a $GI/GI/1$ multi-class queue operating under a non-preemptive discipline. Kleinrock first analyzed this system under the additional assumption of Poisson arrivals [15], and the generalization to general arrival process is due Schrage [23].

Let $\overline{V}_j$ denotes the expected waiting time in the queue for a class-j job, i.e., the elapsed time between the arrival and start being served. Since we assume the discipline is non-preemptive, the expected waiting time in the queue is independent of the required service time. Thus, the expected conditional response time consists of two components: the waiting time in the queue and the service time: we have that for all $x \geq 0$,

$$T_j(x) = x + \overline{V}_j, \tag{8}$$

Now we are in position to state the work conservation law for this system.

Proposition 4 [23, 17] *Consider a $GI/GI/1$ multi-class queue, under any work conserving non-preemptive scheduling discipline. Then the expected waiting times of the various classes satisfy*

$$\sum_{i=1}^{M} \rho_i \overline{V}_i + \overline{V}_0 = \overline{U}, \tag{9}$$

where $\overline{V}_0 = \frac{1}{2}\sum_{i=1}^{M} \lambda_j E[X_j^2]$.

We sketch the proof. By Little's law and since the scheduling discipline is non-preemptive, we have that the expected number of class-j jobs waiting to get service is $\lambda_j \overline{V}_j$, $j = 1, \ldots, M$. Each of these jobs contributes on average $E[X_j]$ units to the unfinished work. Hence, the contribution to the unfinished work in the system due to the class-j jobs that are waiting for service is $\rho_j \overline{V}_j$, $j = 1, \ldots, M$. Thus, the general conservation law can be expressed as $\sum_{i=1}^{M} \rho_i \overline{V}_i + \overline{V}_0 = \overline{U}$, where $\overline{V}_0$ denotes the expected residual unfinished work of the job in service. In order to calculate $\overline{V}_0$ we can use a renewal-theory argument. Conditioned on the fact that queue is busy with a class-j job, what happens with probability ρ_j, the sequence of required service times form a renewal process. Hence, the expected residual unfinished work can be interpreted as the *equilibrium forward recurrence time* from an arbitrary point to the next renewal epoch, which is precisely $\frac{E[X_j^2]}{2E[X_j]}$. Unconditioning it follows that

$$\begin{aligned} \overline{V}_0 &= \sum_{j=1}^{M} \rho_j \frac{E[X_j^2]}{2E[X_j]} \\ &= \frac{1}{2}\sum_{i=1}^{M} \lambda_j E[X_j^2]. \end{aligned}$$

Similar to the previous cases, the value of the constant $\overline{U}$ is independent of the scheduling discipline. For the case of Poisson arrivals, its value is given by equation (3) and equation (9) simplifies to

$$\sum_{i=1}^{M} \rho_i \overline{V}_i = \rho \frac{\sum_{i=1}^{M} \lambda_i E[X_i^2]}{2(1-\rho)}.$$

4 Unified Conservation Law

In the previous section, we have seen that the derivation of the various conservation laws are somewhat different in nature. This difference is more obvious when comparing the cases of single-class and multi-classes. In this section we state the main result of this paper. In Theorem 1 we develop a conservation law for a $GI/GI/1$ multi-class system operating under a general scheduling discipline and with a general service time distribution. The key idea of the proof consists on realizing that, provided the compound arrival process is renewal, we can apply the Palm inversion formula to each of the classes, and thus we obtain $\overline{U_j} = \lambda_j E[W_j]$, where $\overline{U_j}$ is the time average unfinished class-j work in the system and $E[W_j]$ denotes the expected cumulative burden of a class-j job on the unfinished work.

Invoking now the general conservation law, we obtain the expression $\overline{U} = \sum_{j=1}^{M} \lambda_j E[W_j]$. In spite of the apparent simplicity of the derivation, this result generalizes already existing laws (the only requirement is that the discipline is work conserving). In addition, this expression unifies the different conservation laws already derived for single-server queues, and in fact, in Section 5 we will show that previous laws can be obtained as particular cases of the one stated here.

Theorem 1 *Consider a $GI/GI/1$ multi-class queue under a work-conserving scheduling discipline. Then, provided that Assumption 1 is satisfied, the expected conditional response time of the various classes satisfy:*

$$\sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} \overline{F}_j(x) \left(T_j(x) + \int_{u=0}^{x} \frac{\partial T_j(u;x)}{\partial x} du \right) dx = \overline{U}. \tag{10}$$

If in addition we assume that class-j, $j = 1, \ldots, M$, jobs arrive according to a Poisson process, then

$$\overline{U} = \frac{\sum_{j=1}^{M} \lambda_j E[X_j^2]}{2(1-\rho)}.$$

Proof. We consider a class j, $j = 1, \ldots, M$. Let W_j^i, $i = 1, 2, \ldots$, be the contribution of the i-th class-j job to the unfinished work of the system. Since $\rho < 1$, the busy period has a finite length *with probability 1.* Furthermore, since the superposed arrival process is a renewal process, the begin points of the consecutive busy periods constitute regeneration points and, as a consequence, the sequence $\{W_j^n\}_{n=1}^{\infty}$ is a regenerative process with finite cycle lengths. Hence, the process $\{W_j^n\}_{n=1}^{\infty}$ is stationary and ergodic.

Applying the Palm inversion formula[6, 3] to the unfinished work of class j, we obtain $\overline{U_j} = \lambda_j E[W_j]$. We derive now the value of $E[W_j]$. Let $\tau_j^i(u;x)$ denote the amount of time that the i-th class-j job spends in the system until it gets $u \leq x$ units of service. In particular we note that $\tau_j^i(x;x)$ is equal to the response time of the i-th class-j job. Note that $E[\tau_j^i(u;x)] = T_j(u;x)$

Then, similar to equation (5) it follows

$$\begin{aligned} E[W_j] &= E[\int_{x=0}^{\infty} \int_{u=0}^{x} \tau_j^i(u;x) du dF_j(x)] \\ &= \int_{x=0}^{\infty} \int_{u=0}^{x} T_j(u;x) du dF_j(x) \end{aligned}$$

Since by Assumption 1 the function $T_j(u;x)$ has a continuous partial derivative with respect to x, integrating by parts we get

$$E[W_j] = -\overline{F}_j(x) \int_{u=0}^{x} T_j(u;x) du \Big|_{x=0}^{\infty} + \int_{x=0}^{\infty} \overline{F}_j(x) \left(T_j(x;x) + \int_{u=0}^{x} \frac{\partial T_j(u,x)}{\partial x} du \right) dx. \tag{11}$$

The second moment of the service time distribution satisfies

$$\int_0^{\infty} x^2 dF_j(x) = \int_0^{\infty} 2x\overline{F}_j(x) dx + \lim_{x\to\infty} x^2 \overline{F}_j(x).$$

Since the service time distribution has a finite second moment it implies that $\lim_{x\to\infty} x^2 \overline{F}_j(x) = 0$. Let $B(y)$ be the expected length of the busy period initiated by a job of size y. Then it follows that

$$T_j(u;x) \leq T_j(x;x) \leq B(x + \overline{U}) = \frac{x + \overline{U}}{1-\rho}.$$

Therefore,

$$\begin{aligned} \lim_{x\to\infty} \overline{F}_j(x) \int_{u=0}^{x} T_j(u;x) du &\leq \lim_{x\to\infty} \overline{F}_j(x) x T_j(x) \\ &\leq \lim_{x\to\infty} \overline{F}_j(x) \frac{x^2 + \overline{U}x}{1-\rho} \\ &= 0. \end{aligned}$$

Thus, equation (11) becomes

$$E[W_j] = \int_{x=0}^{\infty} \overline{F}_j(x) \left(T_j(x) + \int_{u=0}^{x} \frac{\partial T_j(u;x)}{\partial x} du \right) dx.$$

Now the result follows after summing over all the classes and invoking the general conservation law (equation (1)).

When the arrival processes of the various classes are Poisson, the time average unfinished work is given by equation (3) and the result follows. ■

The fact that $T_j(u;x)$ requires a continuous derivative with respect to x does not seem to be very restrictive. For instance, for all non-anticipating disciplines $T_j(u;x) = T_j(u)$ for all $u \leq x$, and hence $\frac{\partial T_j(u;x)}{\partial x} = 0$, for all $x \geq u$. We emphasize that the previous argument holds even when $T_j(x)$ is not a continuous function and contains jumps. This would be the case for example in threshold based disciplines such as Multi Level Processor Sharing (MLPS) disciplines [17, Section 4.7].

The most popular anticipating discipline (SRPT) as well as the recent FSP have been studied solely for single-class systems. For the case of SRPT, an expression for $T(u;x)$ was provided in [20]

$$T^{SRPT}(u;x) = \int_{z=x-u}^{x} \frac{1}{1 - \lambda \int_0^{z^-} t dF(t)} dz + \frac{\lambda \int_0^x t\overline{F}(t) dt}{(1 - \lambda \int_0^{x^-} t dF(t))^2}. \tag{12}$$

Taking the derivative with respect to x, it is easy to see that a sufficient condition for $T^{SRPT}(u;x)$ to have a continuous derivative with respect to x is that the required service time distribution is a continuous function from the right.

Besides the appealing features of FSP, its analysis seems to be difficult. We conjecture but cannot prove that $T^{FSP}(u;x)$ has a continuous derivative with respect to x under a condition similar to the one required for SRPT.

As we have discussed in Section 2.2, in the context of multi-class queue, the most important disciplines are non-anticipating. Therefore, in the following corollary we specialize Theorem 1 to this set of disciplines.

Corollary 1 *In addition to the conditions of Theorem 1, assume that the scheduling discipline is non-anticipating. Then*

$$\sum_{j=1}^{M} \lambda_j \int_0^{\infty} \overline{F}_j(x) T_j(x) dx = \overline{U}. \tag{13}$$

If the arrival processes of the various classes are Poisson, then

$$\overline{U} = \frac{\sum_{j=1}^{M} \lambda_j E[X_j^2]}{2(1-\rho)}.$$

Proof. The proof follows readily from Theorem 1 after noting that for all $0 \leq u \leq x$,

$$\frac{\partial T_j(u;x)}{\partial x} = 0.$$

■

Corollary 1 was used in [2] in order to show that in the case of DPS, the expected conditional response time of each class has an asymptote with slope $1/(1-\rho)$ when $x \to \infty$. In addition a closed form expression of the bias was provided.

4.1 Performance of anticipating disciplines

The conservation-law derived in Theorem 1 may be useful in evaluating the effect that deploying an anticipating service time discipline has on the mean unconditional response time. Assume the service time distribution is exponential. Then equation (10) becomes

$$\rho\overline{T} = \overline{U} - \lambda \int_{x=0}^{\infty} \overline{F}(x) \int_{u=0}^{x} \frac{\partial T(u;x)}{\partial x} du dx.$$

The second term on the right-hand side captures the effect on the mean unconditional response time that deploying an anticipating scheduling discipline may have. For instance, denote by $\overline{T}^{\pi_1}$ and $\overline{T}^{\pi_2}$ the expected unconditional response time obtained by two arbitrary work-conserving disciplines π_1 and π_2. In addition, assume that π_1 is a non-anticipating discipline. Taking the difference of the unconditional response time obtained by both disciplines we obtain

$$\rho\left(\overline{T}^{\pi_1} - \overline{T}^{\pi_2}\right) = \lambda \int_{x=0}^{\infty} \overline{F}(x) \int_{u=0}^{x} \frac{\partial T^{\pi_2}(u;x)}{\partial x} du dx.$$

Assume now that π_2 is an anticipating discipline that gives preference to short jobs. As a consequence

$$\frac{\partial T^{\pi_2}(u,x)}{\partial x} \geq 0.$$

In the case of SRPT this can be readily verified from equation (12).

It follows that

$$\overline{T}^{\pi_2} \leq \overline{T}^{\pi_1}.$$

It would be interesting to evaluate the reduction of the unconditional response time obtained by disciplines such as SRPT and FSP. Upper and lower bounds on this reduction can be found in [28] and [4]. We plan to investigate whether a precise evaluation of the term $\frac{\partial T(u,x)}{\partial x}$ would allow to obtain more precise bounds.

5 Relation between Conservation Laws

In this section we show that the conservation laws stated in Section 3 can be obtained as a particular case of Theorem 1. This illustrates that Theorem 1 provides a unified view on the conservation laws for single server queues.

5.1 Single-class

We start by considering the single-class case. It is straightforward to derive a work conservation law for a single-class, non-anticipating scheduling discipline, and setting $M = 1$ in Corollary 1 we obtain directly Proposition 1.

In the case of anticipating disciplines, setting $M = 1$ we obtain that

$$\lambda \int_0^{\infty} \overline{F}(x) \left(T(x) + \int_0^x \frac{\partial T(u;x)}{\partial x} du \right) dx = \overline{U}. \tag{14}$$

It can be shown that the above expression is equivalent to the expression obtained by O'Donovan. First, integrating by parts we obtain

$$\lambda \int_{x=0}^{\infty} \int_{u=0}^{x} T(u;x) du dF(x) = \overline{U}.$$

Making the change of variable $u = x - r$, and integrating by parts again over the second term of the integrand we obtain

$$\lambda \int_{x=0}^{\infty} \int_{r=0}^{x} r \partial_r T(x-r;x) dF(x) = \overline{U},$$

where the notation ∂_r denotes a differential with respect to the variable r. Now it suffices to interchange the order of the integrals in order to obtain the result stated in Proposition 2.

5.2 Multi-class

We focus now on a multi-class queue. Assume first that the discipline is non-anticipating and the service time distributions are exponential. From the assumption of exponential service time distribution, it follows that for $j = 1, \ldots, M$,

$$\overline{F}_j(x) dx = E[X_j] dF_j(x). \tag{15}$$

Plugging equation (15) into equation (10) and noting that $\frac{\partial T_j(u;x)}{\partial x} = 0$, for all $j = 1, \ldots, M$, we get

$$\begin{aligned}\overline{U} &= \sum_{j=1}^{M} \lambda_j \int_0^{\infty} T_j(x) dF_j(x) \\ &= \sum_{j=1}^{M} \lambda_j \overline{T}_j,\end{aligned}$$

which is precisely equation (7).

Finally, let us consider a non-preemptive scheduling discipline. The expected conditional response time for a class-j job satisfies, for all $x \geq 0$,

$$T_j(x) = x + \overline{V}_j.$$

Substituting this into equation (10) we obtain

$$\begin{aligned}\overline{U} &= \sum_{j=1}^{M} \lambda_j \int_0^{\infty} T_j(x) \overline{F}_j(x) dx \\ &= \frac{1}{2} \sum_{j=1}^{M} \lambda_j E[X_j^2] + \sum_{j=1}^{M} \lambda_j \overline{V}_j \int_0^{\infty} \overline{F}_j(x) dx \\ &= \overline{V}_0 + \sum_{j=1}^{M} \rho_j \overline{V}_j,\end{aligned}$$

which is equivalent to what Proposition 4 states.

Acknowledgments

The author is thankful to Samuli Aalto, Konstantin Avrachenkov, Sem Borst, Misja Nuyens and Rudesindo Núñez-Queija for reading and providing valuable comments on previous versions of the present paper. The author wishes to thank Rudesindo Núñez-Queija for the derivation of equation (5).

References

[1] S. Aalto, U. Ayesta, and E. Nyberg-Oksanen. Two-level processor sharing scheduling disciplines: Mean delay analysis. In *Proceedings of ACM SIGMETRICS*, pages 97–105, 2004.

[2] K.E. Avrachenkov, U. Ayesta, P. Brown, and R. Núñez-Queija. Discriminatory processor sharing revisited. to appear in Proceedings of INFOCOM, 2005.

[3] F. Baccelli and P. Brémaud. *Elements of queuing theory: Palm martingale calculus and stochastic recurrences.* Springer, 2003.

[4] N. Bansal. On the average sojourn time under M/M/1/SRPT. *Operations Research Letters*, 33(2):195–200, 2005.

[5] O. Boxma. Workloads and waiting times in single-server systems with multiple customer classes. *Queueing Systems*, 5:185–214, 1989.

[6] S.L. Brumelle. On the relation between customer and time average in queues. *Journal of Applied Probability*, 2:508–520, 1971.

[7] E.G. Coffman and I. Mitrani. A characterization of waiting time performance realizable by single-server queues. *Operations Research*, 3(28):810–821, 1980.

[8] M. Dacre, K. Glazebrook, and J. Niño-Mora. The achievable region approach to the optimal control of stochastic systems. *Journal of the Royal Statistical Society. Series B, Methodological*, 61(4):747–791, 1996.

[9] G. Fayolle, I. Mitrani, and R. Iasnongorodski. Sharing a processor among many job classes. *Journal of the ACM*, 27(3):519–532, 1980.

[10] A. Federgruen and H. Groenevelt. Characterization and optimization of achievable performance in general queueing systems. *Operations Research*, 36:733–741, 1988.

[11] E. Friedman and S. Henderson. Fairness and efficiency in processor sharing protocols to minimize sojourn times. *Proceedings of ACM SIGMETRICS*, pages 229–337, 2003.

[12] E. Gelenbe and I. Mitrani. *Analysis and Synthesis of Computer Systems.* London: Academic Press, 1980.

[13] T.C. Green and S. Stidham. Sample-path conservation laws, with application to scheduling queues and fluid systems. *Queueing Systems*, 36:175–199, 2000.

[14] D.P. Heyman and M.J. Sobel. *Stochastic Models in operations research, Volume I: Stochastic Processes and operating characteristics.* McGraw-Hill, 1982.

[15] L. Kleinrock. A conservation law for a wide class of queueing disciplines. *Naval Research Logistics Quarterly*, 12(2):181–192, June 1965.

[16] L. Kleinrock. Time-shared systems: A theoretical treatment. *Journal of the ACM*, 14(2):242–261, 1967.

[17] L. Kleinrock. *Queueing Systems, vol. 2.* John Wiley and Sons, 1976.

[18] L. Kleinrock and A. Nilsson. On optimal scheduling algorithms for time-shared systems. *Journal of the ACM*, 28(3):477–486, 1981.

[19] M. Miyazawa. Rate conservation laws: a survey. *Queueing Systems*, 15:1–58, 1994.

[20] T.M. O'Donovan. Distribution of attained service and residual service in general queueing systems. *Operations Research*, 22:570–575, 1974.

[21] A.K. Parekh and R.G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, 1993.

[22] L.E. Schrage. The queue M/G/1 with feedback to lower priority queues. *Management Science*, 13:466–471, 1967.

[23] L.E. Schrage. An alternative proof of a conservation law for the queue G/G/1. *Operations Research*, 18:185–187, 1970.

[24] L.E. Schrage and L.W. Miller. The queue M/G/1 with the shortest remaining processing time discipline. *Operations Research*, 14:670–684, 1966.

[25] J. Shanthikumar and D. Yao. Multiclass queueing systems: Polymatroidal structure and optimal scheduling control. *Operations Research*, 40(2):293–299, 1992.

[26] K. Sigman. A note on a sample-path rate conservation law and its relationship with H=λ G. *Advances in Applied Probability*, 23:662–665, 1991.

[27] M.J.G. van Uitert. *Generalized Processor Sharing Queues.* PhD thesis, Technische Universiteit Eindhoven, 2003.

[28] A. Wierman, M. Harchol-Balter, and T. Osogami. Nearly insensitive bounds SMART scheduling. In *Proceedings of ACM SIGMETRICS*, 2005.

[29] R.W. Wolff. *Stochastic Modeling and the theory of Queues.* Prentice-Hall, 1989.