9

LRD Isolation in Generalized Processor Sharing

This chapter introduces the application of stochastic network calculus to analysis of queuing processes in generalized processor sharing (GPS) and packet-based GPS (PGPS) systems with long-range dependent (LRD) traffic inputs.

The GPS discipline is a widely studied non-FIFO scheduling discipline [84], due to its attractive characteristics. One is that each backlogged session is guaranteed a minimum service rate in GPS. This ensures that the misbehavior of other flows has a limited effect on an individual session, and provides the foundation of isolation between sessions. Achieving isolation further enables GPS to guarantee differentiated QoS for individual sessions. Another attractive characteristic of GPS is that it is work-conserving and any excess service rate can be redistributed among backlogged flows. The second characteristic enables GPS to obtain a statistical multiplexing gain between input flows. Because of these two characteristics, GPS is deemed an ideal scheduling discipline that meets the following two requirements. One is to provide isolation between flows, where isolation means that the queuing process behaves no worse than its single server queue (SSQ) process with a comparable service rate. This guarantees that the scheduling discipline is able to protect an individual flow against misbehavior from other flows. The other is to realize a statistical multiplexing gain. This suggests that a flow can utilize excess service rate allocated to other flows. When GPS is extended to packet-switched networks, it is usually referred to as weighted fair queuing (WFQ)[33] or packet-based GPS (PGPS) [113].

Long-range dependent (LRD) traffic is an important class of traffic in modern-day networks because long-range dependency is exhibited in many types of networks and network traffic, such as Ethernet [94], WWW [27], compressed video traffic [9, 62], and TCP traffic [48]. Since LRD traffic has burstiness extending over various timescales, a Weibull bound rather than a conventional exponential bound is usually associated with LRD traffic's single-server queue process [55].

By applying the results from previous chapters, this chapter studies the queuing behavior of a single-server system, where the GPS discipline is

Y. Jiang, Y. Liu, *Stochastic Network Calculus*,
DOI: 10.1007/978-1-84800-127-5_9,

adopted and the inputs are LRD flows. In particular, we derive two bounds for individual queuing processes in the GPS system. Complimenting the analysis on GPS in Section 7.3.4, these bounds provide valuable insights into the isolation between multiple GPS sessions. More specifically, it is shown that the index parameter in the upper bound of one LRD flow, in addition to the decay rate and the asymptotic constant, may be affected by other LRD flows. In addition, a necessary and sufficient condition for a flow being guaranteed to be *LRD-isolated* from other flows is derived. Based on this condition, a technique that can be used to quickly check if a flow can be guaranteed to be LRD-isolated from other flows with a given GPS service weight assignment is introduced. When some flows have already been assigned contract weights according to some service level agreement (SLA) that cannot be changed, the introduced technique can also be used to determine the minimum contract weight to be assigned to the flow in order for it to be guaranteed to be LRD-isolated from other flows.

9.1 Introduction

In this section, we briefly review the fundamental knowledge on the GPS scheduling discipline and the queue length distribution of LRD traffic. Their relations with that arrival curve and service curve are also shown to facilitate the analysis based on results from the earlier chapters.

9.1.1 GPS Fundamentals

Generalized processor sharing (GPS) is a scheduling discipline defined under the assumption that sources are described by fluid models [143]. Consider a GPS server with rate γ serving N sessions. Let each session i be assigned a weight parameter that is a fixed real-valued positive number ϕ_i. The set $\{\phi_1, \phi_2, ..., \phi_N\}$ thus represents the GPS assignment. Each session is assumed to have its dedicated queue. The N sessions share the server in the following way [113], as also introduced in Section 7.3.4:

- It is work-conserving; i.e., as long as there are packets backlogged in any of the GPS queues, the server is never idle.
- The excess service rate, if any, is redistributed among the backlogged sessions in proportion to their weight parameters.
- Let $S_k(s,t)$ denote the amount of traffic served in the time interval $[s,t]$ for session k. If session i is backlogged in the system during the entire interval $[s,t]$ (i.e., there is always traffic queued for session i), then

$$\frac{S_i(s,t)}{S_j(s,t)} \geq \frac{\phi_i}{\phi_j}, \quad j = 1, 2, ..., N. \tag{9.1}$$

From (9.1), it is clear that when session i is backlogged, it is guaranteed a backlog clearing rate (or equivalently a guaranteed service rate) of at least

$$\gamma_i = \frac{\phi_i}{\sum_{j=1}^{N} \phi_j} \gamma. \tag{9.2}$$

As shown in Table 2.1 in Chapter 2, it is easy to verify that the GPS server provides to flow i a deterministic strict service curve $\beta_i(t) = \frac{\phi_i}{\sum_{j=1}^{N} \phi_j} \gamma t$.

Let the arrival process for a stationary GPS session i be $A_i(t)$ with long-term average rate λ_i. Assume $\sum_{i=1}^{N} \lambda_i < \gamma$.

In order to characterize the effect of backlogs from a set of sessions, the concept of the so-called *feasible ordering* [143] of the sessions will be frequently referred to hereafter and is defined based on their arrival rates and GPS service weight parameters as follows [113].

Definition 9.1. *For a given set of input traffic flows in a GPS system whose long-term average rate is λ_i, an ordering is called a feasible ordering among the sessions with respect to $\{\lambda_1, \lambda_2, ..., \lambda_N\}$ and GPS service weight parameters $\{\phi_1, \phi_2, ..., \phi_N\}$ if*

$$\lambda_i < \varphi_i \left(\gamma - \sum_{j=1}^{i-1} \lambda_j \right), \ 1 \le i \le N, \tag{9.3}$$

where $\varphi_i = \frac{\phi_i}{\sum_{j=i}^{N} \phi_j}$ is a constant associated with weight parameters and by convention, $\sum_{j=1}^{i-1} \lambda_j = 0$ when $i = 1$.

Note that one of the results of feasible ordering is that for a given set of input traffic flows in a GPS server with $\sum_{i=1}^{N} \lambda_i < \gamma$ there always exists at least one feasible ordering that satisfies (9.3) after being relabeled (e.g., see [143]).

Also note that the right-hand side of (9.3) can be considered as the service rate available to flow i. It is clear, by definition, that those flows ordered earlier than flow i will affect the service rate available to flow i. However, they will not affect the index parameter of the queuing process of a heavier-tailed LRD flow i, as will be explained later in more detail.

9.1.2 LRD Traffic Characterization

LRD traffic is often characterized by heavy traffic bursts that extend over a wide range of timescales [114] [132]. The LRD traffic backlog, buffered within a singe-server queue (SSQ), often possesses a tail distribution that decays slower than that of traditional (e.g., Poisson) traffic. More specifically, the queue length distribution of traditional traffic obeys a certain exponential form. For the case of LRD traffic, the Weibull distribution has been used to characterize the slower decaying SSQ distribution [9] [37] [116].

The queue length distribution, which is Weibull bounded (WB), is defined as follows [109].

Definition 9.2. *A stochastic SSQ process, denoted by* $W^{SSQ,\gamma}(t)$*, where* γ *is the service rate of the queue, is WB(*C*,* η*,* υ*) with parameters* $C(>0)$ *denoting the asymptotic constant,* $\eta(>0)$ *denoting the decay rate, and* $(0<)\upsilon(\leq 1)$ *denoting the index parameter, if it satisfies*

$$P\{W^{SSQ,\gamma}(t) > w\} < Ce^{-\eta w^{\upsilon}} \tag{9.4}$$

for all $w \geq 0$ *and all* $t \geq 0$*.*

In Definition 9.2, the quantity $P\{W^{SSQ,\gamma}(t) > w\}$ essentially represents the probability that the backlog of the SSQ with service rate γ will exceed a certain queue size w. In other words, $P\{W^{SSQ,\gamma}(t) > w\}$ represents the queue length distribution of the SSQ. In addition, the decay rate η increases with γ because when the service rate increases, the likelihood that the queue length exceeds w will decrease. Also, the index parameter υ can be further expressed in terms of the Hurst parameter H, which is commonly used to characterize the degree of long-range dependence [9] [37] [116] and, more specifically, $\upsilon = 2(1-H)$, where $0.5 \leq H < 1$. A traffic process with $H = 0.5$ corresponds to conventional traffic with a queue length distribution that decays exponentially. A larger H, or a smaller υ, corresponds to heavier-tailed LRD traffic.

The definition of WB shows that it is indeed a special case of a gSBB or v.b.c. stochastic arrival curve with bounding function $Ce^{-\eta w^{\upsilon}}$. Therefore, WB has all properties of a gSBB and v.b.c. stochastic arrival curve.

9.1.3 LRD Isolation of Flows

In Definition 9.2, the index parameter is what differentiates an LRD flow from a short-range dependent (SRD) flow. Although the decay rate and constant parameters also define the queuing process, these parameters form the exponential bound parameters commonly associated with an SRD flow. Hence their presence, by definition, is for the purpose of describing the SRD property of the flow.

The index parameter, found in the Weibull bound formula, was introduced to bound flows exhibiting LRD behavior that cannot be suitably bounded by just the constant parameter and the decay rate. Hence, the LRD property of a flow, by definition, is primarily due to its index parameter. Accordingly, we introduce the following notion of *LRD isolation*.

Definition 9.3. *A flow, when multiplexed with other flows in a queue system, is said to be LRD-isolated (from other flows) in that queue system if its resulting queue process has the same or larger index parameter (i.e., less heavy tailed) as the index parameter associated with its SSQ process with service rate equivalent to that guaranteed in the queue system.*

This notion of "LRD isolation" is different from the conventional understanding of *flow isolation*. In flow isolation, the major concern is the flow's service rate, and *a flow is said to be isolated from other flows if this flow is not adversely affected by these flows* [88]. Based on this, an LRD flow is flow isolated if and only if its queue process is not adversely affected after it is multiplexed and served with other flows in the GPS server.

It can be shown that flow isolation is guaranteed for a flow in a GPS server if the flow can be ordered first in a feasible ordering. The reason is that under this case, the flow is always guaranteed a service rate greater than its long term average rate based on (9.3), which is not affected by other flows. In addition, Lemma 9.11 will show that the flow's queue process in the GPS system is not adversely affected (with respect to its SSQ process) by other flows. However, if the flow cannot be ordered first in any feasible ordering, the guaranteed service rate to the flow may depend on the arrival rates of some other flows. In other words, it may vary over time and hence the queue process of the flow in the GPS system could be affected adversely.[1] As a result, if a flow cannot be ordered first in any feasible ordering, the flow may or may not be guaranteed to be flow isolated from other flows. However, a flow can still be guaranteed to be LRD-isolated (from heavier-tailed flows) even if some lighter-tailed flows have to be ordered before this flow in *all* feasible orderings, as will be discussed later in this section. Clearly, flow isolation implies LRD isolation but not vice versa. Since the index parameter is the most important measure of the LRD property (heaviness or lightness of the tail) of a flow, the notion of LRD isolation as defined above is useful when studying LRD flows.

9.2 Analysis of LRD Traffic

9.2.1 Single Arrival Process

In this subsection, we establish the relationship between a WB SSQ and a Weibull bounded burstiness (WBB) arrival process. We begin by defining the burstiness constraint qualifier that describes the arrival process of LRD traffic as follows.

Similar to the notation in Definition 9.2, let C denote the asymptotic constant, υ the index parameter, μ the decay rate[2], and ρ the long-term "upper rate" of the arrival process, which will be further elaborated in Lemma 9.5.

Definition 9.4. *A traffic arrival process $A(t)$ is WBB(ρ, C, μ, υ) with parameters ρ, C, μ, and υ if it satisfies*

$$P\{A(s,t) > \rho(t-s) + w\} < Ce^{-\mu w^{\upsilon}} \tag{9.5}$$

[1] Note that, when $\lambda_i = \phi_i \gamma$, flow i cannot be ordered first according to (9.3), although it is flow-isolated.

[2] Not to be confused with the symbol η, which denotes the decay rate of a WB SSQ process.

for all $w \geq 0$ *and all* $0 \leq s \leq t$.

Here again, $A(s,t)$ is the amount of arrival traffic accumulated in time interval $[s,t]$. In addition, the decay rate μ will increase with ρ, just as η will increase with γ in a WB SSQ process (see Definition 9.2).

It can be seen that WBB is a special case of a t.a.c. stochastic arrival curve with bounding function $Ce^{-\mu w^{\upsilon}}$. Therefore, WBB has all properties of a t.a.c. stochastic arrival curve. Additionally, WBB has interesting properties useful to the objectives of this chapter, which are presented below.

Lemma 9.5. *An arrival process* $A(t)$ *that is WBB(*ρ*,* C*,* μ*,* υ*) possesses the property that its parameter* ρ *is always larger than or equal to its long-term average rate*

$$\rho \geq \lim_{t-s\to\infty} \frac{E[\int_s^t A(u)du]}{t-s}. \tag{9.6}$$

Proof. First, we have

$$\begin{aligned} & E\left[\int_s^t A(u)\,du\right] \\ &= \int_0^\infty \Pr\left\{\int_s^t A(u)\,du > x\right\} dx \\ &= \int_0^{\rho(t-s)} \Pr\left\{\int_s^t A(u)\,du > x\right\} dx \\ &\quad + \int_0^\infty \Pr\left\{\int_s^t A(u)\,du > \rho(t-s) + x\right\} dx \\ &< \rho(t-s) + \int_0^\infty Ce^{-\mu x^{\upsilon} dx}. \end{aligned}$$

Second, as long as $\upsilon > 0$, we have

$$\lim_{t-s\to\infty} \frac{\int_0^\infty Ce^{-\mu x^{\upsilon}} dx}{t-s} = \lim_{t\to\infty} \frac{\int_0^t Ce^{-\mu x^{\upsilon}} dx}{t} = \lim_{t\to\infty} Ce^{-\mu t^{\upsilon}} = 0.$$

Therefore,

$$\rho \geq \lim_{t-s\to\infty} \frac{E[\int_s^t A(u)du]}{t-s}.$$

□

The long-term upper rate ρ is useful for the purpose of bounding the entire ensemble of sample time observations that constitute the stochastic arrival process $A(t)$. In particular, let $A_n(t)$ be the nth sample observation of $A(t)$ in $[s,t]$, and let $\lambda_n = \lim_{t-s\to\infty} \frac{\int_s^t A(n,u)du}{t-s}$ be the corresponding average arrival rate for this sample. If we were to repeat the observation of $A(t)$ infinitely many times using different start times, so that n approaches infinity, then we would have a corresponding list of average arrival rates $\lambda_1, \lambda_2, ..., \lambda_{n\to\infty}$. This long term upper rate ρ ranges between the lower limit $E[\lambda_n]$ and the higher limit $\rho_{max} = \max[\lambda_1, \lambda_2, ..., \lambda_{n\to\infty}]$. For a conservative (loose) WBB bound on $A(t)$, one may set ρ to the higher limit ρ_{max}. However, notice that the long-term upper rate, defined in (9.5), is applied continuously even if the arrival process is inactive. Therefore, a lower value of ρ, where $\rho_{max} \geq \rho \geq E[\lambda_n]$, may suffice to produce a tighter WBB bound on $A(t)$. To summarize, the use of the long-term upper rate ρ in (9.5) is essential for a general stochastic process that may not be stationary (i.e., $\lambda_1 \neq \lambda_2 \neq ... \neq \lambda_{n\to\infty}$). However, in practical arrival processes, stationarity is an implicit property for a flow that has some fixed arrival rate λ. This means that if this flow is presented to the queue at different start times, the same average rate λ applies. Hence, for the case of practical flows, $\lambda_1 = \lambda_2 = ... = \lambda_{n\to\infty} = \lambda$ and therefore $\rho = \lambda$. Although many of the later derivations following this definition are still based on ρ, readers should be aware that for practical considerations ρ ought to be replaced by λ since practical arrival processes are by default implicitly stationary in property. In fact, in the consideration of the GPS and PGPS discipline in Sections 9.3, 9.4, and 9.5, we consider λ instead of ρ. Finally, it is also noted that, besides ρ, the WBB expression in (9.5) also contains other parameters, such as the decay rate μ, the index υ, and the asymptotic constant C. These parameters can similarly be modified to obtain either loose or tight WBB bounds.

Following the relationship between the t.a.c. stochastic arrival curve and v.b.c. stochastic arrival curve, the following theorem establishes the relationship between a WBB arrival process and a WB SSQ process.

Theorem 9.6. *Consider a work-conserving SSQ that transmits at rate γ. Suppose the queue is fed with a single arrival process $A(t)$, and let $W^{SSQ,\gamma}(t)$ be the workload stored in the queue at time t. Then:*

(i) If $W^{SSQ,\gamma}(t)$ is WB, then $A(t)$ is WBB with long-term upper rate $\rho = \gamma$.
(ii) If $A(t)$ is WBB with long-term upper rate $\rho = \gamma - \varepsilon$ for some $\varepsilon > 0$, then $W^{SSQ,\gamma}(t)$ is WB.

Although LRD traffic is usually described in terms of some WB SSQ process, it is still insufficient to proceed on to GPS analysis since in GPS we are concerned with multiple arrival processes rather than a single arrival process. If there is no burstiness constraint on a single arrival process, there is not much that can be deduced about the stability of a GPS server that is serving a number of these arrival processes. With the introduction of Theorem 9.6, we can now proceed further since it is now known that any LRD arrival process resulting in a WB SSQ process must satisfy the WBB constraint with some long-term upper rate ρ. This means that for a GPS server serving a number of LRD sources, as long as the sum of the long-term upper rates of these LRD sources does not exceed the service capacity of the GPS server, the GPS queue will be stable and further analysis can proceed.

9.2.2 Aggregate Process

In this subsection, several bounds on the aggregate WB SSQ process are presented, based on the superposition property of the v.b.c. stochastic arrival curve. These bounds will later be used frequently for the analysis of GPS and PGPS.

Lemma 9.7. *Let $W_1(t)$ be WB(C_1, η_1, υ_1) and $W_2(t)$ be WB(C_2, η_2, υ_2). The two processes can either be dependent or independent. Then, $W_1(t) + W_2(t)$ is WB $((C_1 + C_2 + C^*), \eta, \upsilon)$, satisfying*

$$P\{W_1(t) + W_2(t) > w\} < (C_1 + C_2 + C^*)e^{-\eta w^{\upsilon}}, \tag{9.7}$$

where $\eta = \frac{1}{\frac{1}{\eta_1}+\frac{1}{\eta_2}}$, $\upsilon = \min(\upsilon_1, \upsilon_2)$, *and* $C^* = (C_1 + C_2)\left[e^{-\eta\left(w_0^{\upsilon_{\max}} - w_0^{\upsilon}\right)} - 1\right]$.

Proof. According to the superposition property of the v.b.c. stochastic arrival curve, we have

$$\begin{aligned} P\{W_1(t) + W_2(t) > w\} &< C_1 e^{-\eta_1 w^{\upsilon_1}} \otimes C_2 e^{-\eta_2 w^{\upsilon_2}} \\ &< C_1 e^{-\eta_1 p w^{\upsilon_1}} + C_2 e^{-\eta_2 (1-p) w^{\upsilon_2}}, \text{ for any } 0 \le p \le 1. \end{aligned}$$

We choose p such that $\eta_1 p = \eta_2(1-p)$, i.e., $p = \frac{\eta_2}{\eta_1+\eta_2}$. Defining $\eta = \frac{\eta_1\eta_2}{\eta_1+\eta_2}$ and $\upsilon = \min(\upsilon_1, \upsilon_2)$, we obtain

$$P\{W_1(t) + W_2(t) > w\} < C_1 e^{-\eta w^{\upsilon_1}} + C_2 e^{-\eta w^{\upsilon_2}}.$$

If $0 < w < 1$, then we have

$$P\{W_1(t) + W_2(t) > w\} < (C_1 + C_2)\, e^{-\eta w^{\upsilon_{\max}}}, \tag{9.8}$$

where $\upsilon_{\max} = \max[\upsilon_1, \upsilon_2]$.

If $w \ge 1$, then we have

$$P\{W_1(t) + W_2(t) > w\} < (C_1 + C_2)\, e^{-\eta w^{\upsilon}}, \tag{9.9}$$

where $\upsilon = \min[\upsilon_1, \upsilon_2]$.

It is noted that both (9.8) and (9.9) have the Weibull bound form except with different index parameters. Next, we try to combine (9.8) and (9.9) so the same index parameter, namely υ rather than $\upsilon_{\max}$, can also be used for the case where $0 < w < 1$. First, we notice that the bound using the index $\upsilon_{\max}$ in (9.8) is always larger than the bound using the index in (9.9) for the range $0 < w < 1$. Once $w > 1$, the bound in $0 < w < 1$ is always larger than the bound in (9.8). At $w = 0$ and $w = 1$, the bounds in (9.8) and (9.9) have exactly the same values. Hence, in order to extend (9.9), which uses the index υ, to provide a bound for the same case where $0 < w < 1$, we can always add an additional asymptotic constant factor C^* to raise the bound of (9.9). This additional asymptotic constant C can be easily obtained, as it is related to the maximum displacement between (9.8) and (9.9) when $0 < w < 1$. More specifically, let

$$f(w) = (C_1 + C_2)\, e^{-\eta w^{\upsilon_{\max}}} - (C_1 + C_2)\, e^{-\eta w^{\upsilon}}.$$

Notice that $f(w)$ is zero at $w = 0$ and $w = 1$, and $f(w) > 0$ only for $0 < w < 1$, where both $e^{-\eta w^{\upsilon_{\max}}}$ and $e^{-\eta w^{\upsilon}}$ monotonically decrease with w. Therefore, there exists a unique maximum point of $f(w)$ for $w \in (0,1)$. Let w_0 maximize $f(w)$ for $0 < w < 1$. Specifically, w_0 is the solution to the following non-algebraic equation:

$$\frac{e^{-\eta w^{\upsilon_{\max}}}}{w^{\upsilon_{\max}}}\upsilon_{\max} = \frac{e^{-\eta w^{\upsilon}}}{w^{\upsilon}}\upsilon.$$

Hence the additional asymptotic constant C^* is given by

$$\begin{aligned} C^* &= f(w_0)\, e^{-\eta w_0^{\upsilon}} \\ &= (C_1 + C_2)\left[e^{-\eta w_0^{\upsilon_{\max}}} - e^{-\eta w_0^{\upsilon}}\right] e^{-\eta w_0^{\upsilon}} \\ &= (C_1 + C_2)\left[e^{-\eta\left(w_0^{\upsilon_{\max}} - w_0^{\upsilon}\right)} - 1\right]. \end{aligned}$$

Therefore, we have

$$P\{W_1(t) + W_2(t) > w\} < (C_1 + C_2 + C^*)e^{-\eta w^{\upsilon}},$$

where $w \geq 0$. □

Lemma 9.7 can be applied step-by-step to obtain the following theorem for the case involving multiple WB processes.

Theorem 9.8. *Let $W_i(t)$, $1 \leq i \leq N$ be N WB processes with parameters (C_i, η_i, υ_i), respectively. These processes can either be dependent or independent. Then, $W_1(t) + W_2(t) + \cdots + W_N(t)$ is WB(($\sum_{i=1}^{N} C_i + C^*$), η, υ), satisfying*

$$Pr\left\{\sum_{i=1}^{N} W_i(t) > w\right\} < \left(\sum_{i=1}^{N} C_i + C^*\right) e^{-\eta w^{\upsilon}}, \tag{9.10}$$

where $\eta = \frac{1}{\sum_{i=1}^{N}\frac{1}{\eta_i}}$, $\upsilon = \min(\upsilon_1, \upsilon_2, ..., \upsilon_N)$, and C^* is a constant as given in Lemma 9.7.

Given Lemma 9.7, the proof of Theorem 9.8 is straightforward and hence omitted. For N WB queuing processes with the same LRD degree (i.e., the same υ), $\upsilon = \min(\upsilon_1, \upsilon_2, ..., \upsilon_N)$ is the tightest lower bound on the index parameter. But, for N WB queuing processes with different LRD degrees, it is a loose bound because the index parameter of multiplexed LRD flows is in general heavier tailed than the individual flows due to multiplexing gain.

Similar to the study on the independent case of two gSBB processes, in the case where $W_1(t)$ and $W_2(t)$ are two independent processes, Lemma 9.9 and Theorem 9.10 present alternate bounds to those obtained in Lemma 9.7 and Theorem 9.8, respectively. The alternate bounds are useful since in certain cases they are tighter.

Lemma 9.9. *Let $W_1(t)$ and $W_2(t)$ be two independent processes WB(C_1, η_1, υ_1) and WB(C_2, η_2, υ_2), respectively. If $\eta_2 \leq \eta_1$ and $\upsilon_2 \leq \upsilon_1$, then for $\forall w > 2$, $W_1(t) + W_2(t)$ has an upper bound of the form*

$$P\{W_1(t) + W_2(t) > w\} < C_2^{WB}(w)e^{-\eta w^{\upsilon}}, \tag{9.11}$$

where $\eta = \min\{\eta_1, \eta_2\} = \eta_2$, $\upsilon = \min\{\upsilon_1, \upsilon_2\} = \upsilon_2$, and

$$C_2^{WB}(w) = C_2 h(C_1) + C_1,$$

where $h(C_1) = 1 + C_1 \upsilon \eta (e^{\eta} - 1) + C_1 w^{\upsilon} \eta$.

Note that the η in Lemma 9.7 is $\frac{\eta_1 \eta_2}{\eta_1 + \eta_2}$, which is always less than or equal to $\min(\eta_1, \eta_2)$, which is the η in (9.11). Hence Lemma 9.9 yields a larger (thus better) decay rate. In fact, if $\eta_2 \approx \eta_1$, then η in Lemma 9.9 is almost twice the value of η in Lemma 9.7. However, the asymptotic constant in Lemma 9.9 increases with w, which is a trade-off. For a heavy-tailed arrival process where $\upsilon \to 0$ so that, for practical and finite values of w, $w^{\upsilon} \to 1$ and thus $C_2^{WB}(w)$ approaches a constant, the penalty for using Lemma 9.9 is insignificant. Conversely, if η_2 differs significantly from η_1 (i.e., $\eta_2 \ll \eta_1$), then $\eta \to \min(\eta_1, \eta_2)$, making Lemma 9.7 more attractive.

Table 9.1 summarizes the preferences (in terms of which lemma to use to obtain the bound) assuming that in all the scenarios the queue size of interest is larger than or equal to 2.

Table 9.1. Preference for Lemma 9.7 or Lemma 9.9 in different scenarios

Scenario	Preference
$\eta_2 \approx \eta_1$ and υ_2 is small	Lemma 9.9
$\eta_2 \ll \eta_1$ and υ_2 is large	Lemma 9.7
All other cases	Either is ok

Similar to the way in which Lemma 9.7 is extended to Theorem 9.8, we now extend Lemma 9.9 to Theorem 9.10, whose proof can be obtained by recursively applying Lemma 9.9.

Theorem 9.10. *Let $W_i(t)$, $1 \leq i \leq N$, be N independent WB processes with parameters (C_i, η_i, υ_i). If the queuing processes can be rearranged such that* *the* N*th queuing process has the property that* $\eta_N \leq \eta_j$ *and* $\upsilon_N \leq \upsilon_j$ *for* $1 \leq j \leq N-1$*, then, for* $\forall w > 2$*,* $W_1(t) + W_2(t) + \cdots + W_N(t)$ *has an upper bound of the form*

$$P\left\{\sum_{i=1}^{N} W_i(t) > w\right\} < C_N^{WB}(w) e^{-\eta w^{\upsilon}}, \tag{9.12}$$

where $\eta = \min\{\eta_1, \eta_2, ..., \eta_N\} = \eta_N$*,* $\upsilon = \min\{\upsilon_1, \upsilon_2, ..., \upsilon_N\} = \upsilon_N$ *and*

$$C_N^{WB}(w) = \sum_{j=1}^{N}\left[C_j \prod_{l=1}^{j-1} h(C_l)\right], \tag{9.13}$$

where $h(C_l) = 1 + C_l \upsilon \eta (e^{\eta} - 1) + C_l w^{\upsilon} \eta$*, and by convention* $\prod_{l=1}^{j-1} h(C_l) = 1$ *when* $j = 1$*.*

9.3 Sample Path Behavior of LRD Traffic in a GPS System

Recall from Theorem 9.6 that any LRD traffic input whose queue length distribution is characterized by a WB distribution has an arrival process that satisfies the WBB constraint with some long-term upper rate ρ. Hereafter, we consider N stationary flows that maintain the same long-term average rate λ_i, $i = 1, 2, ..., N$, irrespective of the start time of the flow. As mentioned earlier (in the discussion after Lemma 9.5), the long-term upper rate ρ reduces to the more familiar λ.

9.3.1 GPS Decomposition

Let A_i denote a sample path (or a single realization) of the random arrival process $A_i(t)$ and $Q_i^{GPS,\gamma}$ denote the corresponding sample path of the GPS queue backlog due to the sample arrival process A_i. To obtain relevant bounds on $Q_i^{GPS,\gamma}$, we use a method similar to that in [143] to decompose the GPS system into N fictitious WB single server queues (SSQs), denoted by $\delta_i^{SSQ,\gamma_i}(t)$, with individual rates γ_1, γ_2, ..., γ_N, where $\gamma_i > \lambda_i$, $\sum_{i=1}^{N} \gamma_i \leq \gamma$, and $\gamma_i \leq \varphi_i(\gamma - \sum_{j=1}^{i-1} \gamma_j)$. Now, the reason for considering the N fictitious WB SSQs is that their bounds are easier to obtain and would surely bound $Q_i^{GPS,\gamma}$ as well. This is because the N fictitious WB SSQs do not consider multiplexing gain, while the $Q_i^{GPS,\gamma}$ queue process does.

Without loss of generality, let 1, 2,..., N be a feasible ordering of the fictitious processes with respect to γ_i's. From Lemma 3 of [143], Lemma 9.11 can be derived.

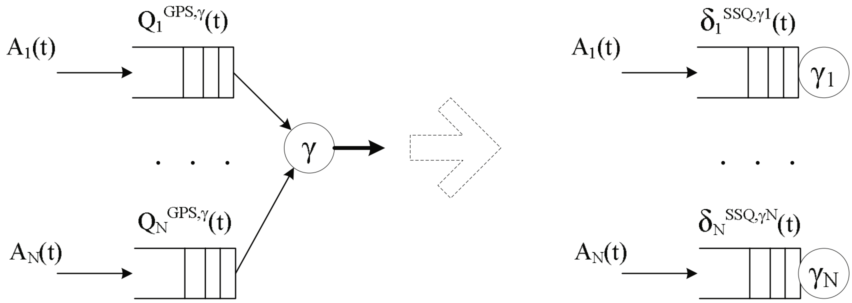

Fig. 9.1. Decomposing a GPS system into N fictitious SSQs

Lemma 9.11. *For any t,*

$$Q_i^{GPS,\gamma}(t) \le \varphi_i \sum_{j=1}^{i-1} \delta_j^{SSQ,\gamma_j}(t) + \delta_i^{SSQ,\gamma_i}(t), \tag{9.14}$$

where each δ_i^{SSQ,γ_i} SSQ process is independent.

Lemma 9.11 provides an upper bound on the queue length $Q_i^{GPS,\gamma}(t)$ of an individual session in the GPS system in terms of the queue length $\delta_i^{SSQ,\gamma_i}(t)$ in the fictitious system. It is clear from Lemma 9.11 that to bound the distribution of $Q_i^{GPS,\gamma}(t)$, it suffices to bound the following aggregate of fictitious queue length processes:

$$\varphi_i \delta_1^{SSQ,\gamma_1}(t) + \varphi_i \delta_2^{SSQ,\gamma_2}(t) + ... + \varphi_i \delta_{i-1}^{SSQ,\gamma_{i-1}}(t) + \delta_i^{SSQ,\gamma_i}(t) \tag{9.15}$$

In what follows, we will provide two bounds on (9.15), i.e., the right-hand side of (9.14).

9.3.2 A General Bound

For N individual LRD flows sharing a GPS server on the condition of queue stability (i.e., $\sum_{i=1}^{N} \lambda_i < \gamma$) and under the assumption that 1, 2,... N is a feasible ordering with respect to ϕ_i and λ_i, $\lambda_i < \gamma_i$ for $i = 1, 2, ..., N$, we present a GPS bound in Theorem 9.12 that is based on Theorem 9.8.

Theorem 9.12. *Each individual queue length distribution in the GPS system has an upper bound as follows:*

$$P\{Q_i^{GPS,\gamma}(t) > q\} < C_i^{GPS} e^{-\eta_i^{GPS} q^{v_i^{GPS}}}, \tag{9.16}$$

where

$$\upsilon_i^{GPS} = \min_{1 \le j \le i} \{\upsilon_j\}, \tag{9.17}$$

$$\eta_i^{GPS} = \frac{1}{\sum_{j=1}^{i} \frac{1}{\bar{\eta}_j}}, \tag{9.18}$$

$$C_i^{GPS} = \left(\sum_{j=1}^{i} C_j + C^* \right) e^{-\eta_i^{GPS}}, \tag{9.19}$$

and, in the above,

$$\bar{\eta}_j = \begin{cases} \frac{\eta_j}{\varphi_i^{\upsilon_j}} & 1 \le j < i \\ \eta_i & j = i \end{cases}$$

and C^* can be obtained similarly as in Theorem 9.8.

Proof. First, because all the input flows are LRD flows, we have

$$P\{\delta_j^{SSQ,\gamma_j}(t) > q\} < C_j e^{-\eta_j q^{\upsilon_j}} \quad , \quad j = 1, 2, ..., N. \tag{9.20}$$

Secondly letting

$$\delta_{j,eqv}^{SSQ,\gamma_j}(t) = \varphi_i \delta_j^{SSQ,\gamma_j}(t),\ j < i, \tag{9.21}$$

we have

$$P\{\delta_{j,eqv}^{SSQ,\gamma_j}(t) > q\} = Pr\{\delta_j^{SSQ,\gamma_j}(t) > \frac{q}{\varphi_i}\} < C_j e^{-\bar{\eta}_j q^{\upsilon_j}}, \tag{9.22}$$

for $1 \le j < i$, where $\bar{\eta}_j = \frac{\eta_j}{\varphi_i^{\upsilon_j}}$. Finally, since (9.15) can now be written as

$$\delta_{1,eqv}^{SSQ,\gamma_1}(t) + \delta_{2,eqv}^{SSQ,\gamma_2}(t) + ... + \delta_{i-1,eqv}^{SSQ,\gamma_{i-1}}(t) + \delta_i^{SSQ,\gamma_i}(t), \tag{9.23}$$

by combining (9.20), (9.22) and Lemma 9.11, one can easily verify the result in (9.16) based on Theorem 9.8. □

Theorem 9.12 gives a general upper bound on queue length distribution in a GPS system. It is important to note that the GPS upper bound on flow i is not affected by the flows that are ordered after flow i (because they do not factor in the upper bound expression for flow i). It is affected only by flows 1 to $i-1$, but the impact on the bound is *negligible* as long as flow i is heavier-tailed (i.e., has a smaller index parameter υ_i) than any of the flows 1 through $i-1$. In fact, the index parameter in the bound for flow i is not affected at all as long as flows 1 to $i-1$ are lighter tailed than flow i.

9.3.3 An Alternate Bound

In Theorem 9.13, an alternate upper bound on individual session queue length in GPS with LRD traffic is provided based on Theorem 9.10. Such a bound may be better than the bound previously given in Theorem 9.12 but can only be applied under the condition that for any given i there exists a $1 \le k \le i$ such that both $\bar{\eta}_k$ and υ_k are minimal in addition to the conditions stated for Theorem 9.12.

Theorem 9.13. *If there exists a* $1 \le k \le i$ *such that* $\bar{\eta}_k = \min_{1\le j\le i}\{\bar{\eta}_j\}$ *and* $\upsilon_k = \min_{1\le j\le i}\{\upsilon_j\}$*, then for* $\forall q > 2$*, the upper bound for individual session queue length is*

$$P\{Q_i^{GPS,\gamma}(t) > q\} < C_i^{GPS}(q) e^{\eta_i^{GPS} q^{\upsilon_i^{GPS}}}, \tag{9.24}$$

where

$$\upsilon_i^{GPS} = \min_{1\le j\le i}(\upsilon_j) = \upsilon_k, \tag{9.25}$$

$$\eta_i^{GPS} = \min_{1\le j\le i}(\bar{\eta}_j) = \bar{\eta}_k, \tag{9.26}$$

$$C_i^{GPS}(q) = C_k \prod_{l=1,l\neq k}^{i} h_i^{GPS}(C_l) + \sum_{j=1,j\neq k}^{i}\left[C_j \prod_{l=1,l\neq k}^{j-1} h_i^{GPS}(C_l)\right], \tag{9.27}$$

with $h_i^{GPS}(C_l) = 1 + C_l \upsilon_i^{GPS}\eta_i^{GPS}(e^{\eta_i^{GPS}} - 1) + C_l w^{\upsilon_i^{GPS}}\eta_i^{GPS}$*, and by convention* $\prod_{l=1,l\neq k}^{j-1} h_i^{GPS}(C_l) = 1$ *when* $j = 1$*.*

Proof. Without loss of generality, assume that $k < i$. The aggregate process in (9.23) can be rewritten such that the *kth* process with the minimum decay rate as well as with the minimum index parameter appears last in the sequence as follows:

$$\delta_{1,eqv}^{SSQ,\gamma_1}(t) + \delta_{2,eqv}^{SSQ,\gamma_2}(t) + ... + \delta_{k-1,eqv}^{SSQ,\gamma_{k-1}}(t) + \delta_{k+1,eqv}^{SSQ,\gamma_{k+1}}(t) + ... + \delta_i^{SSQ,\gamma_i}(t) + \delta_{k,eqv}^{SSQ,\gamma_k}(t). \tag{9.28}$$

Hence, by applying Theorem 9.10, Theorem 9.13 can be easily verified. □

Theorem 9.13 provides an upper bound on an actual session i's backlog $Q_i^{GPS,\gamma}(t)$ in the GPS system when there exists a very heavy-tailed LRD flow with the smallest index parameter (as well as the smallest decay rate). One implication of Theorem 9.13, similar to Theorem 9.12, is that it is desirable to order the flows that are heavier-tailed as close to the end of a feasible ordering as possible, again since the index parameter in the upper bound for the individual queue length of flow i will not be affected if and only if the flows 1 through $i-1$ are all lighter-tailed than flow i.

9.4 Technique to Check and Ensure LRD Isolation

Recall from the discussion immediately following Definition 9.1 that in a stable GPS system where $\sum_{j=1}^{N}\lambda_j < \gamma$, there exists at least one feasible ordering for a given weight assignment. Before we discuss LRD isolation, it is useful to revisit the concept of flow isolation with Lemma 9.14.

Lemma 9.14. *In a stable GPS system, if flow i satisfies the condition:*

$$\lambda_i < \gamma \frac{\phi_i}{\sum_{j=1}^{N} \phi_j}, \tag{9.29}$$

then the flow is flow-isolated.

Proof. The proof is straightforward since the right-hand side of (9.29) is the minimum guaranteed rate. Yet, we still provide the required proof for this lemma since several intermediate results of this proof will be used later to prove newer results pertaining to LRD isolation. Relabel flow i as flow 1 and all other $N-1$ flows to be flows 2 to N. Note that flow 1 now satisfies (9.3), and hence all we need to show is that the remaining $N-1$ flows can be feasibly ordered after flow 1. To this end, consider a new GPS system with service rate $\gamma' = \gamma - \lambda_1$. Since $\gamma' > \sum_{j=2}^{N} \lambda_j$, the new GPS system is also stable, and hence there always exists a feasible ordering such that (after relabeling the flows 2 to N) we have for any flow $2 \leq i \leq N$

$$\begin{aligned} \lambda_i &< \frac{\phi_i}{\sum_{j=i}^{N} \phi_j} \left(\gamma' - \sum_{j=2}^{i-1} \lambda_j \right) \\ &= \frac{\phi_i}{\sum_{j=i}^{N} \phi_j} \left(\gamma - \sum_{j=1}^{i-1} \lambda_j \right). \end{aligned} \tag{9.30}$$

Note that the equations above becomes the same as (9.3), which means that if flow 1 is ordered first, the remaining $N-1$ flows can also be ordered to yield a feasible ordering. Therefore, flow 1 is flow-isolated. □

It should be noted that (9.29) is only a sufficient condition for flow isolation, not a necessary condition. In fact, it is a sufficient condition to guarantee a flow to be flow-isolated. However, as mentioned earlier, a flow can still be isolated even if it cannot be "guaranteed" to be flow-isolated, or even if it does not satisfy (9.29).

Based on Lemma 9.14, an obvious method to guarantee the flow isolation of every flow is to assign the weight of every flow according to (9.29) such that every flow i can be ordered in the first place in a feasible ordering. As mentioned earlier, being able to order a flow first in any feasible ordering is the most applicable condition to guarantee a flow to be flow-isolated for, say, admission control purposes. That, however, is not necessary to guarantee just LRD isolation of a flow, which is less strict than flow isolation, as will be discussed later.

9.4.1 Limitations of Existing Methods

In this subsection, we discuss the shortcomings of the existing methods for assigning weights to achieve flow isolation and testing whether a flow can be flow-isolated for a given weight assignment.

A GPS system may support the following three types (or service classes) of flows. A Type 1 flow requires a higher QoS than that provided by flow isolation, so it requires a contract weight that is much larger than $\frac{\lambda_i}{\gamma}\sum_{j=1}^{N}\phi_j$; A Type 2 flow requires flow isolation and thus needs a contract weight that is a little larger than $\frac{\lambda_i}{\gamma}\sum_{j=1}^{N}\phi_j$. A Type 3 flow only requires LRD isolation (but not flow isolation) and thus can have a contract weight less than $\frac{\lambda_i}{\gamma}\sum_{j=1}^{N}\phi_j$. Note that the contract weight cannot be changed as long as the service level agreement (SLA) is in effect. On the other hand, a (lightly loaded) GPS system may assign a flow an *extra* weight (if available) to provide the flow with better service, and such extra weights can be adjusted (e.g., transferred to other flows) by the GPS system.

The method of assigning weights based on (9.29) has a limited applicability in supporting Types 1 and 2 but is not applicable to Type 3 flows. From users or applications' viewpoint, having Type 3 flows is useful because certain applications may require a less strict performance guarantee than that given by flow isolation, and such flows can be admitted into a GPS system and with less costs to the users or applications. In addition, from the GPS system's viewpoint, supporting Type 3 flows allows it to admit more flows than otherwise possible, thus increasing its utilization and potential revenues.

For example (hereafter referred to as Example 1), consider a GPS system with $\gamma = 16$ and five flows numbered 1 through 5 in descending order of their index parameters whose $\lambda_i = i$ where, $1 \le i \le 5$. Assume that the total weight is $\sum_{j=1}^{5}\phi_j = 16$, and in addition flows 1 and 2 have been assigned contract weights of $\phi_1 = 1.1$ and $\phi_2 = 4$, respectively. Since the remaining weight for flows 3, 4 and 5 is 10.9 but the sum of their arrival rates is 12, it is clear that (9.29) cannot be used to assign the weights to all three remaining flows to guarantee their flow isolation.

In general, due to the existence of Type 1 flows (e.g., flow 2 in Example 1), flow isolation may not always be achievable by every flow, and accordingly the existing approach based on (9.29) may not be useful. Note that, even if flow i does not satisfy (9.29), it may still be LRD isolated. In Example 1, one can assign 2.6 to flow 3 to ensure its LRD isolation (which can be verified using the technique to be proposed later), even though such a weight violates (9.29).

As another example (hereafter called Example 2) showing the deficiency of the existing approaches, assume that the weight assignment for the same five flows as in Example 1 is now $\{1.1, 2.1, 1, 4, 7.8\}$. It is clear that (9.29) cannot be used to test if flows 3 and 4 (both of which violate (9.29)) are LRD-isolated or not. In addition, (9.3) in Definition 9.1 is not effective either. More specifically, in order to use it to test whether flow 4 can be guaranteed to be LRD isolated or not, a naive approach will test if the ordering of 1, 2, 3, 4, 5 is feasible, and because it is not, it will have to examine the ordering of 1, 3, 2, 4, 5 and then the ordering of 2, 3, 1, 4, 5 and so on. In the worst case, to test if flow i can be guaranteed to be LRD-isolated or not, all possible orderings involving j lighter-tailed flows, where $0 \leq j \leq (i-1)$, have to be tested. Thus, the (worst-case) time complexity of the testing process is $O(i!)$. When the number of flows is large, such an approach is clearly infeasible.

9.4.2 Necessary and Sufficient Condition

We now determine, for a given flow i, not only the set of lighter-tailed flows, denoted by f_i, that can be ordered before flow i in a feasible ordering, but also the minimum contract weight to ensure the LRD isolation of flow i. To this end, we first initialize f_i to be empty. Then, if there exists a flow k where $1 \leq k < i$ that satisfies

$$\frac{\lambda_k}{\phi_k} < \frac{\gamma - \sum_{j \in f_i} \lambda_j}{\sum_{j=1}^{N} \phi_j - \sum_{j \in f_i} \phi_j}, \tag{9.31}$$

we add flow k to f_i and update the right-hand side of (9.31), which will be denoted by $R(f_i)$. We repeat the process above until no such flow k exists and denote the resulting set by F_i and accordingly the final value of $R(f_i)$ by $R(F_i)$. Note that this process of obtaining F_i has the worst-case time complexity of $O(i^2)$.

One can easily verify that when a flow k that satisfies (9.31) is added to f_i, the resulting $R(f_i)$ increases; i.e., $R(f_i) < R(f_i \cup k) \leq R(F_i)$ if $f_i \subseteq F_i$. Conversely, if we were to add a flow k' that does not satisfy (9.31) to f_i, then $R(f_i \cup k') \leq R(f_i)$. In other words, $R(F_i)$ is the maximum value that flow i can obtain from all flows that are lighter-tailed than flow i. This observation is important for proving the following theorem, which provides a both necessary and sufficient condition for the LRD isolation guarantee of flow i.

Theorem 9.15. *Suppose there are N flows in a GPS system that are numbered in the descending order of their index parameters as $1, 2, ..., N$, and their contract weights are ϕ_1, ϕ_2, ..., ϕ_N, respectively. Then flow i is guaranteed to be LRD-isolated from other flows if and only if*

$$\frac{\lambda_i}{\phi_i} < \frac{\gamma - \sum_{j \in F_i} \lambda_j}{\sum_{j=1}^{N} \phi_j - \sum_{j \in F_i} \phi_j} = R(F_i). \tag{9.32}$$

Proof. (i) To show that (9.32) is a sufficient condition, we note that flow i also satisfies (9.31), just as any flow $k < i$ in F_i does. Accordingly, if we let $F_i' = F_i \cup \{i\}$ and note that when f_i is empty $R(f_i) = \frac{\gamma}{\sum_{j=1}^{N} \phi_j}$, we have the following (based on the observation drawn preceding the theorem):

$$\frac{\gamma - \sum_{j \in F_i'} \lambda_j}{\sum_{j=1}^{N} \phi_j - \sum_{j \in F_i'} \phi_j} > \frac{\gamma - \sum_{j \in F_i} \lambda_j}{\sum_{j=1}^{N} \phi_j - \sum_{j \in F_i} \phi_j} > \frac{\gamma}{\sum_{j=1}^{N} \phi_j}.$$

Accordingly, we can easily conclude that

$$\frac{\sum_{j \in F_i'} \lambda_j}{\sum_{j \in F_i'} \phi_j} < \frac{\gamma}{\sum_{j=1}^{N} \phi_j}.$$

The above means that if we treat the flows in F_i' as one *big* flow with arrival rate $\sum_{j \in F_i'} \lambda_j$ and weight $\sum_{j \in F_i'} \phi$, it satisfies (9.29). Hence, according to Lemma 5.1, there exists a feasible ordering with this *big* flow ordered first. In other words, flow i can be feasibly ordered before any heavier-tailed flow. Note that the exact ordering of the flows within F_i will not affect the LRD isolation of flow i. In fact, the flows in F_i can be feasibly ordered according to the order in which they are added to F_i in (9.31), with flow i being ordered right after them.

(ii) We now prove that (9.32) is necessary by contradiction. Suppose (9.32) does not hold for flow i but there still exists a feasible ordering with flow i ordered before any heavier-tailed flows. Denote the set of all the (lighter-tailed) flows that are feasibly ordered before flow i by F_i^* (which may be empty). According to (9.3), we should have:

$$\frac{\lambda_i}{\phi_i} < \frac{\gamma - \sum_{j \in F_i^*} \lambda_j}{\sum_{j=1}^{N} \phi_j - \sum_{j \in F_i^*} \phi_j} = R(F_i^*).$$

However, since F_i^* contains zero or more flows in F_i and zero or more flows not in F_i, we have $R(F_i^*) \leq R(F_i)$ based on the discussion preceding the theorem; or in other words,

$$\frac{\lambda_i}{\phi_i} < R(F_i^*) \leq R(F_i),$$

which contradicts the assumption that (9.32) does not hold for flow i. □

Note that if a flow satisfies (9.29), it will satisfy (9.31) but not vice versa. With (9.32), whether a flow is guaranteed to be LRD isolated or not depends only on the weights assigned to the flows in F_i, and flow i itself. In Example 1, one can easily verify that $F_3 = \{1, 2\}$ and $R(F_3) = (16 - 3)/(16 - 5.1) = 1.19$. Hence, if $\phi_3 = 2.6$, flow 3 satisfies (9.32) and thus is guaranteed to be LRD-isolated. On the other hand, in previous Example 2 (where the weight assignment for five flows is $\{1.1, 2.1, 1, 4, 7.8\}$), one can easily verify that since $F_3 = F_4 = \{1, 2\}$ and $R(F_3) = R(F_4) = 13/12.8$, (9.32) cannot be satisfied by flow 3, and thus flow 3 is not guaranteed to be LRD-isolated. On the other hand, flow 4 satisfies (9.32) and thus is guaranteed to be LRD-isolated.

9.4.3 Weight Adjustment and Assignment

Theorem 9.15 is also useful for weight assignment and adjustment in order to guarantee a flow's LRD isolation. More specifically, the observation drawn preceding the theorem (i.e., $R(F_i)$ is maximum with respect to flow i) serves as the base for determining a minimal ϕ_i to guarantee the LRD isolation of flow i.

For instance, consider again Example 2, but now assume that only the weights assigned to flows 1, 2, and 4 are contract weights (i.e., non-adjustable). If we want to ensure LRD isolation of flow 3, we must increase ϕ_3 to above 13/12.8. Such an increase can be accomplished if ϕ_5 has an *extra* weight of 2 that can be transferred to flow 3 (and, as a result, ϕ_5 is reduced to 5.8 from 7.8).

The technique above to adjust the weight of a single flow to ensure its LRD isolation can certainly be extended to ensure LRD isolation of more than one flow provided that there are extra weights in the GPS system that can be adjusted or transferred. As a slightly different example from those above (which we call Example 3), consider five flows numbered in descending order of their index parameters whose arrival rates are more or less randomly distributed as $\{2, 4, 5, 1, 3\}$. Suppose that $\gamma = 17$ (which is sufficient to make the system stable) and the total weight is a constant 17. In addition, suppose that flow 2 (which is a Type 1 flow) has been assigned a contract weight of 7 (and thus the method based on (9.29) cannot be used for weight assignment to guarantee flow isolation of all the other flows, as discussed earlier). If all four other flows are Type 3 flows that only require LRD isolation, we can use Theorem 9.15 to assign contract weights to them to guarantee their LRD isolation as follows (note that one can easily verify that flow 2 can be ordered first in any feasible ordering, so it is already flow-isolated).

For the first flow, from Theorem 9.15, we need to have $\phi_1 > \lambda_1 = 2$, so we set $\phi_1 = 2.1$ (theoretically speaking, we can set $\phi_1 = 2 + \epsilon$, where $\epsilon > 0$ can have a very small value). For flow 3, we first obtain $F_3 = \{1, 2\}$, and then from (9.32) we have

$$\begin{aligned}\phi_3 &> \lambda_3 \frac{\sum_1^5 \phi - \phi_1 - \phi_2}{\gamma - \lambda_1 - \lambda_2} \\ &= 5 \cdot \frac{17 - 2.1 - 7}{17 - 2 - 4} = 3.59.\end{aligned}$$

Accordingly, we set $\phi_3 = 3.6$ to flow 3. Similarly, we set $\phi_4 = 0.72$ and $\phi_5 = 0$[3]. The extra weight available in the system is $17 - 2.1 - 7 - 3.6 - 0.72 = 1.58$, which may be distributed among the five flows in an arbitrary manner.

To further illustrate the usefulness of the proposed technique, let us consider the following corollary of Theorem 9.15 that may be used in the case of online admission control.

Corollary 9.16. *If a flow i is provided a contract weight ϕ_i that guarantees it to be either flow-isolated or just LRD-isolated, it will be guaranteed to be flow-isolated or just LRD-isolated after a new flow j is admitted as long as the system remains stable.*

Proof. Note that, in the corollary, flow j cannot take away any existing contract weights already assigned to the other flows, so its contract weight

[3] Note that, with $\phi_5 = 0$, flow 5 gets best-effort service.

can only come from the extra weight available in the system before it was admitted.

If flow i is guaranteed to be flow-isolated before flow j is admitted, flow i must satisfy (9.29). Hence, flow i is guaranteed to be flow-isolated after flow j is admitted.

Now assume flow i was only guaranteed to be LRD-isolated but not guaranteed to be flow-isolated before flow j was admitted. Under this assumption, there existed a feasible ordering in which the set of flows ordered in front of flow i, F_i, are all lighter-tailed than flow i. Now treat $F_i \cup \{i\}$ as one big flow F_i'. Just as shown in Part (i) of the proof for Theorem 9.15, this big flow F_i' satisfies (9.29) and thus can be ordered in the first place in a feasible ordering. Thus, flow i is still guaranteed to be LRD-isolated. □

Let us continue Example 3 by assuming that the online admission control receives a request for a new flow (flow 6). Suppose that its arrival rate is $\lambda_6 = 1$, and its index parameter is in between those of flows 2 and 3. To ensure its LRD isolation, we first obtain $F_6 = \{1, 2\}$ and then conclude that we need a contract $\phi_6 > 0.718$. Since we have an extra weight of 1.58, we can assign $\phi_6 = 0.72$ and redistribute the remaining extra weight $1.58 - 0.72 = 0.86$ among all six flows.

Note that, from Corollary 9.16, admitting flow 6 as done in the case above will not affect either the flow isolation or LRD isolation of any flows, or in other words, their guaranteed (contracted) performance. There are, however, cases where a heavy-tailed flow has been assigned a weight more than its arrival rate, and hence the remaining weight is not enough to ensure the LRD isolation of the newly arrived flow. An example is that for the same set of five flows described in Example 3, but this time flow 4 instead of flow 2 is a Type 1 flow that requires a contract weight of 5. To ensure each of the four other flows are LRD-isolated, we need the assignment $\{2.1, 4.1, 5.1, 5, 0\}$, which leaves an extra weight of only 0.7. Hence, when flow 6 arrives, it needs $\phi_6 > 1$ to ensure its LRD isolation. In such a case, the system may decide not to admit flow 6 or admit it without ensuring its LRD isolation.

9.5 Sample Path Behavior of LRD Traffic in a PGPS System

The results obtained for the GPS system are now extended to the PGPS system. While the GPS discipline assumes that the input traffic behaves like a fluid such that multiple sessions can be served bit by bit, the packet-based GPS (PGPS) is a more practical discipline in that only one packet at a time may be served. In other words, a PGPS server considers the arrival of a packet only after its last bit has been received. To manage this difference, the PGPS server is often taken to consist of two parts, a regulator and a PGPS core that is a GPS scheduler (see Chapter 4 in [112]), as illustrated in Figure 9.2.

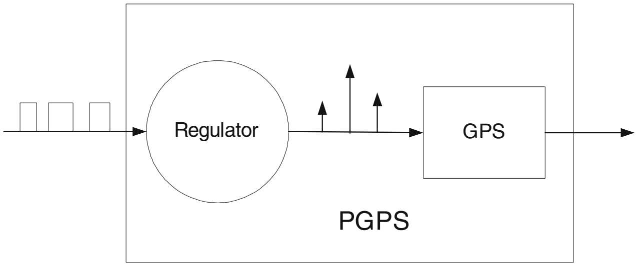

Fig. 9.2. PGPS server

Partially complete (or partially arrived) packets are queued in the regulator, which passes only complete (or arrived) packets to the PGPS core. The output of this regulator, which is the input to the PGPS core, is a series of impulses whose heights represent the sizes of the packets.

Let A_i be the session i input traffic to the PGPS server, which is also the input to the regulator, $A_{i,reg}$ be the output traffic from the regulator, which is the input traffic to the PGPS core, and finally $A(s,t)$ be the total amount of traffic that arrived in time interval $[s,t]$.

It is not difficult to verify that the queuing process of $A_{i,reg}(s,t)$ is also bounded by the queuing process of $A_i(s,t)$ with an extra length L; i.e., $Q_i^{PGPS}(s,t) \leq Q_i^{GPS}(s,t) + L$, where L is the maximum length of all arrived packets (e.g. see Corollary 1 in [113]). From the queuing process Q_i of A_i, which is WB(C,η,υ), we obtain the queuing process Q_i^{PGPS} of $A_{i,reg}$,

$$\begin{aligned} P\{Q_i^{PGPS}(s,t) > q\} &\leq P\{Q_i^{GPS}(s,t) + L > q\} \\ &= P\{Q_i^{GPS}(s,t) > q - L\} \\ &< C_i e^{-\eta_i (q-L)^{\upsilon_i}} \\ &\leq C_i e^{\eta_i L^{\upsilon_i}} e^{-\eta_i q^{\upsilon_i}}, \end{aligned} \tag{9.33}$$

which is WB$(Ce^{\eta L^{\upsilon_i}}, \eta, \upsilon)$. In other words, the two GPS upper bounds derived in Theorems 9.12 and 9.13 in the previous section can be extended to the PGPS domain via a simple transformation of the asymptotic constant $C_i \rightarrow C_i e^{\eta_i L^{\upsilon_i}}$ provided that the queue length or backlog is large enough to exceed the maximum packet length L; i.e., $q > L$. Note that this assumption $(q > L)$ is reasonable because in practice the buffer size B is much larger than L (i.e., $B \gg L$) and, in addition, since the main concern is whether the backlog is about to exceed B, the values of q that are of interest should be close to B and thus larger than L.

For completeness, we now present Theorems 9.17 and 9.18 which are derived from Theorems 9.12 and 9.13 respectively via the use of the simple transformation $C_i \rightarrow C_i e^{\eta_i L^{\upsilon_i}}$ as follows.

Theorem 9.17. *Let $Q_i^{PGPS,\gamma}$, $1 \leq i \leq N$ represent the ith queuing process of the PGPS system with N LRD arrival processes. Then, at any time t, for* *any queue length $q > L$, where L is the maximum packet length of all the N sessions, we have*

$$P\left\{Q_i^{PGPS,\gamma} > q\right\} < C_i^{PGPS} e^{-\eta_i^{GPS} q^{\upsilon_i^{GPS}}}, \tag{9.34}$$

where υ_i^{GPS}, and η_i^{GPS} have already been defined in (9.17) and (9.18), and

$$C_i^{PGPS} = \left(\sum_{j=1}^{i} C_j e^{\bar{\eta}_j L^{\upsilon_j}} + C^*\right),$$

where C^* can be obtained similarly as in Theorem 9.8.

Theorem 9.18. *Under the same assumptions used for Theorem 9.13 except that the server is now a PGPS server, at any time t, for any $q > L > 2$, where L is the maximum packet length of all the N sessions*

$$Pr\{Q_i^{PGPS,\gamma} > q\} < C_i^{PGPS}(q) e^{-\eta_i^{GPS} q^{\upsilon_i^{GPS}}}, \tag{9.35}$$

where υ_i^{GPS}, and η_i^{GPS} have already been defined in (9.25) and (9.18), and

$$\begin{aligned} C_i^{PGPS}(q) = {} & C_k e^{\bar{\eta}_k L^{\upsilon_k}} \prod_{l=1, l\neq k}^{i} h_i^{GPS}(C_l e^{\eta_l L^{\upsilon_l}}) \\ & + \sum_{j=1, j\neq k}^{i} \left[C_j e^{\bar{\eta}_j L^{\upsilon_j}} \prod_{l=1, l\neq k}^{j-1} h_i^{GPS}(C_l e^{\eta_l L^{\upsilon_l}}) \right] \end{aligned}$$

with

$$\begin{aligned} h_i^{GPS}(C_l e^{\eta_l L^{\upsilon_l}}) = {} & (1 + C_l e^{\bar{\eta}_l L^{\upsilon_l}} \upsilon_i^{GPS} \eta_i^{GPS} (e^{\eta_i^{GPS}} - 1) \\ & + C_l e^{\bar{\eta}_l L^{\upsilon_l}} q^{\upsilon_i^{GPS}} \eta_i^{GPS}) \end{aligned}$$

and, by convention, $\prod_{l=1, l\neq k}^{j-1} h_i^{GPS}(C_l e^{\eta_l L^{\upsilon_l}}) = 1$ when $j = 1$.

Note that the bounds above shed light on the LRD isolation among LRD sources sharing a PGPS server. To illustrate this, consider a simple case of two independent LRD sources with a feasible ordering of 1, 2. From Theorem 9.12, the source that appears first in the feasible ordering is always guaranteed to be LRD-isolated. Therefore, the queuing process that is of interest is the last queuing process in the feasible ordering, i.e., $Q_2^{PGPS,\gamma}$. By applying Theorems 9.17 and 9.18, three possible sets of bounds can be obtained as follows:
(i) If $\eta_1 \leq \eta_2$ and $\upsilon_1 \leq \upsilon_2$ then from Theorem 9.18 we have

$$P\{Q_2^{PGPS,\gamma} > q\} < C_{2'}^{PGPS}(q) e^{-\eta_1 q^{\upsilon_1}}, \tag{9.36}$$

where

$$C_{2'}^{PGPS}(q) = C_1 e^{\eta_1 L^{\upsilon_1}} h(C_2 e^{\eta_2 L^{\upsilon_2}}) + C_2 e^{\eta_2 L^{\upsilon_2}}.$$

(ii) Otherwise, $\eta_2 \le \eta_1$ and $\upsilon_2 \le \upsilon_1$, then from Theorem 9.18

$$P\{Q_2^{PGPS,\gamma} > q\} < C_2^{PGPS}(q) e^{-\eta_2 q^{\upsilon_2}}, \tag{9.37}$$

where

$$C_2^{PGPS}(q) = C_2 e^{\eta_2 L^{\upsilon_2}} h(C_1 e^{\eta_1 L^{\upsilon_1}}) + C_1 e^{\eta_1 L^{\upsilon_1}}.$$

(iii) In general, regardless of the relationship between η_1 and η_2 and that between υ_1 and υ_2, from Theorem 9.17, we have

$$P\{Q_2^{PGPS,\gamma} > q\} < (C_1 e^{\eta_1 L^{\upsilon_1}} + C_2 e^{\eta_2 L^{\upsilon_2}}) \times e^{-\eta(q_0^{\upsilon max} - q_0^{\upsilon})} e^{-\eta q^{\upsilon}}, \tag{9.38}$$

where

$$\eta = \frac{\eta_1 \eta_2}{\eta_1 + \eta_2} \quad \textit{and} \quad \upsilon = \min\{\upsilon_1, \upsilon_2\}.$$

The index parameter (as well as the decay rate parameter) of the three bounds shown in (9.36)–(9.38) indicates the influence of source 1 on source 2. In the first case, the bound on $Q_2^{PGPS,\gamma}$ decays slower, and in fact it adopts the same index parameter as that in the bound on the heavier-tailed queuing process δ_1^{SSQ,γ_1}. This means that source 2 is not guaranteed to be LRD-isolated from source 1. In the second case, source 2 is not much affected by source 1 since the bound on $Q_2^{PGPS,\gamma}$ adopts the same index parameter as the bound on δ_2^{SSQ,γ_2}. Finally, in the third case, which is useful when neither of the first two cases is applicable, the bound on $Q_2^{PGPS,\gamma}$ decays slower than both the bound on δ_1^{SSQ,γ_1} and the bound on δ_2^{SSQ,γ_2}.

9.6 Summary and Bibliographic Comments

In this chapter, by applying the relationship between the t.a.c. stochastic arrival curve and v.b.c. stochastic arrival curve, we have established the relationship between a Weibull bounded burstiness (WBB) arrival process and a Weibull bounded (WB) queuing process, which brings more validity to the analysis of the upper bounds on the queuing process with long-range dependent (LRD) traffic inputs.

In addition, this chapter develops several upper bounds on the queue length distribution of the generalized processor sharing (GPS) scheduling discipline with LRD traffic inputs. The GPS bounds have also been extended to a packet-based GPS (PGPS) system. These explicit bounds contribute additional results to stochastic network calculus. In addition, they show that the long range dependency and queue length distribution of an LRD source in a GPS system will in general not be adversely affected despite the presence of other admitted sources as long as it can be feasibly ordered before other heavier tailed flows.

The content of this chapter is mainly based on [141] by Yu, Thng, Jiang, and Qiao. Also in [141], some numerical results on a PGPS system with LRD input flows are given to demonstrate the usefulness of the bounds. There is a vast body of literature on GPS and PGPS. Some closely related works include [131] [139] [143]. While in [131] the focus is on deterministic constraint inputs, an upper bound is developed for the individual session queue length when the input traffic is short-range dependent and particularly has exponentially bounded burstiness (EBB) [138]. The notion of flow isolation can be found in [88] and the notion of LRD isolation was initially introduced by Yu, Thng, Jiang and Qiao in [141].

Problems

9.1. It is said that (9.29) is only a sufficient condition for flow isolation but not a necessary condition. Give an example scenario where flow is isolated but the condition (9.29) is not satisfied.

9.2. Prove Lemma 9.9.

9.3. Prove Lemma 9.11.

9.4. Prove Theorem 9.17.

9.5. Prove Theorem 9.18.

9.6. For Example 2 where the weight assignment for five flows is $\{1.1, 2.1, 1, 4, 7.8\}$, assume that only weights assigned to flows 1, 2, and 3 are contract weights. Find how to adjust the weight to ensure LRD isolation for flow 4.

9.7. In Example 2, let the weight assignment for five flows be $\{1.3, 0.9, 1.2, 3.8, 7.6\}$. Find which flows are guaranteed to be flow-isolated and which flows are guaranteed to be LRD-isolated.