

Packet Routing over Parallel Time-Varying Queues with Application to Satellite and Wireless Networks

Michael Neely Eytan Modiano Charles Rohrs
MIT -- LIDS {mjneely@mit.edu, modiano@mit.edu, crohrs@mit.edu}

Abstract: We consider the problem of routing packets from an arbitrary input stream X over a collection of heterogeneous queues in parallel. When the processing rates (μ_1, \dots, μ_n) of the queues are constant, a simple work conserving routing strategy π_{WC} is shown to hold total system backlog within a fixed upper bound from the resulting backlog of any other policy. Similar results apply to systems with time varying processing rates $(\mu_1(t), \dots, \mu_n(t))$ when routing decisions can be postponed by placing packets in a pre-queue. For the case when routing decisions must be made immediately upon arrival, we demonstrate that the non-predictive policy of routing packets to the shortest queue ensures system stability (and maintains low packet loss rates for finite buffer systems) whenever the system is stabilizable. Finally, we consider a joint problem of routing and power allocation where, for each time varying channel state $c_i(t)$, the rate of each queue i can be varied by adjusting a power parameter p_i (subject to a total power constraint $\sum p_i \leq P_{tot}$) according to a rate-power curve $\mu_i(c_i, p_i)$. A throughput maximizing algorithm is developed for this joint problem.

I. INTRODUCTION

This paper treats a problem of routing packets from an arbitrary input stream X over a set of n heterogeneous queues in parallel (Fig. 1). Routing decisions are made based upon the current state of the queues, and the goal is to maintain an acceptably low level of backlog in the queues for all time. Such a scenario occurs, for example, in a satellite network where data packets arrive to a satellite and can reach the ground using one of several downlink channels. Packet transmission rates along different downlinks may vary as different channel codes are used to adjust to fluctuations in channel conditions (e.g., due to weather). Timescales of variation may be very long--in which case a constant transmission rate model is appropriate--or very short, in which case a time-varying channel model is used. Here, both the cases of constant and time-varying rates are considered from a queueing perspective. The instantaneous transmission rates are assumed to be known at the transmitter, although in the time varying case future rates are unknown and change according to arbitrary ergodic processes.

For constant processing rates, we develop a simple work conserving routing policy π_{WC} that ensures no more than a fixed amount of excess bits in the system above the amount there would be under any other routing strategy. Similar results hold when processing rates are time-varying and routing decisions can be postponed by storing packets in a pre-queue.




Figure 1: Routing over a set of parallel queues with time varying processing speeds $\mu_i(t)$.

When no such pre-queue is available and packets must be routed immediately upon their arrival, it is impossible to bound the number of packets or the amount of unprocessed bits in the queues unless complete knowledge of future events is known. However, we show that the simple, non-predictive strategy of placing an arriving packet in the queue with the least backlogged bits yields a stable system whenever the system is stabilizable. A new notion of stability is developed for this analysis, which is useful for addressing stability issues in general ergodic systems.

We extend these results to address a joint problem of routing and power allocation. A constrained resource of power is allocated to the multiple transmitters to vary the transmission rates in response to packet backlog and/or varying channel conditions. A throughput maximizing algorithm is developed.

Previous work on queue control policies in satellite and wireless systems is found in [1-9]. In [1] a server scheduling problem is formulated in discrete time for the case when a single server can be scheduled to serve one of n parallel queues, each queue having *iid* packet arrivals and *iid* channel outage states (ON or OFF) on each timeslot. In symmetric situations when all queues have the same arrival and outage processes, stochastic coupling is used to show optimality of the “serve-the-longest-connected-queue” policy. In [2] an asymmetric situation of packet routing over two heterogeneous queues is considered when the arrival stream is Poisson, and an optimal threshold policy for minimizing average packet occupancy is developed using dynamic programming techniques. In [6] a Lyapunov function is used to establish a stable routing policy in a wireless network when arrivals are memoryless and packets have memoryless service times. In [8] convexity properties of queueing systems with general $*/*$ inputs are developed and used to establish optimal static policies for routing multiple data streams over n parallel queues.

A related NP-complete problem of batch packet arrivals is considered in [10, 11] where the goal is to route packets to parallel queues to minimize the total delay for transmitting one batch. In [12-14] an online job scheduling problem is considered where n identical servers work on k jobs which arrive randomly over an interval of time $[0, T]$. Simple algorithms which finish within twice the minimum possible completion time are developed.

The main contribution in this paper is to treat stochastic queueing systems and to provide tight, worst case bounds on system performance for arbitrary input processes. Using sample path techniques, we develop additive bounds on the amount of unfinished work in the queues for the entire timeline $t > 0$. Such additive bounds directly translate into statements about queue stability and buffer requirements in the system. This approach differs significantly from the dynamic programming and Lyapunov function techniques which address systems with particular types of input processes.

In the next section we introduce the routing problem by comparing two natural routing strategies: a greedy routing strategy and a work conserving routing strategy. In Section III a simple queueing inequality is developed for time-varying systems and is used as a tool for comparing routing schemes and proving performance bounds. In Sections IV and V we treat routing in systems with a pre-queue and without a pre-queue, respectively, and in Section VII we apply these results to a power allocation problem.

II. THE GREEDY AND WORK CONSERVING ROUTING ALGORITHMS

Consider the system of Fig. 1 with an arbitrary arrival stream X sending packets to be routed over the n queues. Assume all processing rates (μ_1, \dots, μ_n) are constant. The goal is to route packets in a manner that ensures an acceptably low level of unfinished work $U(t)$ and number of packets $N(t)$ in the system for all time. It can be shown that a policy π_{greedy} (described below) is optimal in minimizing $N(t)$ at every instant of time if all packets have fixed lengths and arrive in a single burst at time zero. However, for general arrival streams X , no policy exists that minimizes $N(t)$ or $U(t)$ at

every instant of time--even if the entire future is known in advance. We seek a robust strategy, one whose performance at each time t is sufficiently close to that of a system optimized to minimize backlog at that particular time instant.

Two natural routing strategies emerge, the greedy strategy π_{greedy} and the work conserving strategy π_{WC} :

1. The greedy strategy routes the current packet i to the queue that allows it to exit first:

$$\pi_{greedy}: \text{ Choose queue } k \text{ such that } k = \underset{j \in \{1, \dots, n\}}{\operatorname{argmin}} \left\{ \frac{L_i + U_j(t)}{\mu_j} \right\}.$$

2. The work conserving strategy routes the current packet to the queue that will empty first:

$$\pi_{WC}: \text{ Choose queue } k \text{ such that } k = \underset{j \in \{1, \dots, n\}}{\operatorname{argmin}} \left\{ \frac{U_j(t)}{\mu_j} \right\}.$$

(where L_i is the length of the current packet i , and $U_j(t)$ is the unfinished work in queue j at time t).

Notice that policies π_{greedy} and π_{WC} are identical if all server speeds μ_j are the same, although they may differ considerably under heterogeneous server rates. Because the greedy strategy uses the length of the current packet when making a routing decision, one would expect this policy to offer better performance. However, for suitable choices of the linespeeds (μ_1, \dots, μ_n) , a system under the greedy strategy can have arbitrarily more unfinished work within it than the same system operating under the work conserving strategy, as the following example illustrates.

Suppose a burst of B packets arrive to the system at time 0. After this initial burst, a single packet arrives periodically at times $\{1, 2, 3, \dots\}$. Assume all packets have fixed lengths $L=1$, and that $(\mu_1, \mu_2, \dots, \mu_n) = (1, \epsilon, \dots, \epsilon)$. Suppose that $\epsilon > 0$ is sufficiently small to ensure that all packets from the initial burst as well as all packets thereafter are routed to queue 1 under the greedy strategy π_{greedy} . Thus, under the greedy strategy π_{greedy} , there are always B packets in the system.

Now consider a different set of routing decisions $\tilde{\pi}$: route the B packets from the initial burst amongst the $n-1$ queues of rate ϵ , and route all packets from the periodic stream thereafter to queue 1. Under this policy, queue 1 always has exactly one packet within it. However, the B packets from the initial burst eventually depart from queues $\{2, 3, \dots, n\}$, leaving these queues empty after some finite time T . Hence, after time T the greedy strategy π_{Greedy} results in $B-1$ more packets in the system than policy $\tilde{\pi}$, where $B-1$ can be made arbitrarily large.

Alternatively, we shall see that the work conserving strategy π_{WC} is fundamental in that it never produces more than $n-1$ extra packets in the system compared to *any* other policy. Furthermore, because it utilizes the full bandwidth available, it can be seen to match the policy $\tilde{\pi}$ of this example in having no more than 1 packet within the system after time T .

III. A MULTIPLEXING INEQUALITY

Here we develop a queueing inequality useful for establishing performance bounds on routing policies. Let $X(t)$ represent an arbitrary packet arrival process on the interval $t \geq 0$. A particular $X(t)$ sample path is a non-decreasing staircase function representing the total number of bits that have arrived to the system during $[0, t]$. Jumps in the $X(t)$ function occur at packet arrival epochs and have magnitudes equal to the length of the arriving packet. We assume there is a maximum packet length L_{max} .

Consider the single server queueing system and the multi-server queueing system of Fig. 2. The

non-negative linespeed processes $\mu(t)$ and $\{\mu_i(t)\}$ represent instantaneous server processing rates (in units of bits per second). Here we assume that the rate of the single server queue of Fig. 2a is equal to the sum of the individual server rates in Fig. 2b, i.e., $\mu(t) = \mu_1(t) + \dots + \mu_n(t)$. Assume both systems of Fig. 2 are empty at time 0 and the same input process $X(t)$ is applied to each. Packets are immediately processed in the single server system according to a non-idling service policy. However, in the multi-server system packets are routed to the n queues using any conceivable routing mechanism. All buffers in the queues and in the router device are assumed to be infinite, so that no packets are lost. Let $U_{\text{single-server}}(t)$ represent the unfinished work (in bits) in the single server queue, and let $U_{\text{multi-server}}(t)$ represent the total amount of unfinished bits in the multi-queue system.

Lemma 1: $U_{\text{single-server}}(t) \leq U_{\text{multi-server}}(t)$ for all time $t \geq 0$. (1)

Proof: Let $D_{\text{single-server}}(t)$ and $D_{\text{multi-server}}(t)$ represent the amount of processed bits or “departures” from the single server system and multi-server system, respectively, during $[0, t]$. Clearly $D_{\text{single-server}}(t) = X(t) - U_{\text{single-server}}(t)$, $D_{\text{multi-server}}(t) = X(t) - U_{\text{multi-server}}(t)$, and it suffices to show that $D_{\text{single-server}}(t) \geq D_{\text{multi-server}}(t)$. Notice that the departure functions are both initially 0 at time $t=0$. Whenever the single-server system is empty, $D_{\text{single-server}}(t) = X(t) \geq D_{\text{multi-server}}(t)$, and hence the inequality is satisfied at such times. If the single-server system is not empty, then it is processing packets at the instantaneous rate $\mu(t)$, which is greater than or equal to the instantaneous departure rate of bits from the multi-server system. Hence, the departures from the multi-server system can never overtake the departures from the single-server system. \square

This multiplexing inequality demonstrates that it is always better to multiplex data streams from individual queues to a single queue whose rate is equal to the sum of the individual processing rates. It is useful to consider such a virtual single-server queue to provide a baseline for measuring the performance of routing policies. Strategies yielding unfinished work functions close to the $U_{\text{single-server}}(t)$ lower bound are desirable.

Consider now a particular work conserving routing strategy π_{WC} that operates on the multi-queue system of Fig. 2b. The policy π_{WC} places all incoming packets in a shared buffer device or “pre-queue.” Whenever any server becomes available, the pre-queue instantaneously routes the next packet to that server. If there is more than one server available, the choice is made arbitrarily. Thus, the policy π_{WC} is “work conserving” in that no servers are idle whenever there are packets waiting to be processed. Let $U_{\text{WC}}(t)$ represent the total unfinished work at time t in the multi-server system of Fig. 2b under this policy.

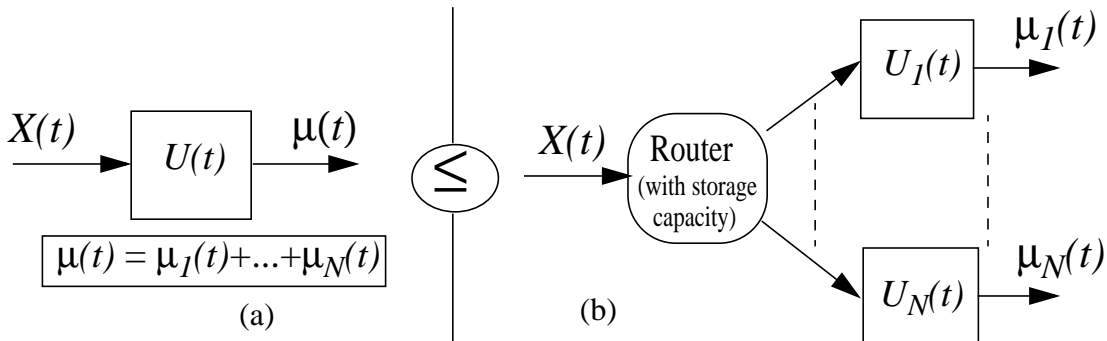


Figure 2: A single server queue and a set of multi-server queues with an arbitrary router device. The sum of the time varying server rates of the servers in (b) equals the rate $\mu(t)$ in (a).

$$\textit{Lemma 2: } U_{\textit{single_server}}(t) \leq U_{WC}(t) \leq U_{\textit{single_server}}(t) + (n-1)L_{\textit{max}}. \quad (2)$$

Proof: The first inequality is just a particular case of Lemma 1. To prove the second inequality, we compare the departed bits $D_{WC}(t)$ and $D_{\textit{single_server}}(t)$ for the multi-server and single-server systems. It suffices to show that $D_{\textit{single_server}}(t) \leq D_{WC}(t) + (n-1)L_{\textit{max}}$. For simplicity, assume that $\mu_i < \infty$ for all i , so that departures drain continuously from the queues. For the above inequality to be violated, there must be some crossing time t^* such that $D_{\textit{single_server}}(t^*) = D_{WC}(t^*) + (n-1)L_{\textit{max}}$. If the single-server system is empty at this time, the departure function $D_{\textit{single_server}}(t)$ cannot increase. Otherwise, the multi-server system holds strictly more than $(n-1)L_{\textit{max}}$ bits, and hence contains at least n distinct packets. By the nature of the work conserving policy π_{WC} , all servers of the multi-server system must be actively processing packets. The departure functions $D_{\textit{single_server}}(t)$ and $D_{WC}(t)$ are thus increasing at the same rate at time t^* , and the departures from the single-server system cannot overtake the bound. \square

Notice that the bounds of Lemma 2 imply that the work conserving routing strategy π_{WC} is stable if and only if the single queue system with the same inputs is stable. Stability issues are addressed further in Section V.

IV. SYSTEMS WITH PRE-QUEUES

The results of the previous section allow routing policy π_{WC} --which is implemented with a pre-queue--to be compared to any other policy π which operates on the same multi-server system. Let $U_{WC}(t)$, $N_{WC}(t)$, $U_{\pi}(t)$, and $N_{\pi}(t)$ represent the unfinished work and number of packets in the multi-queue system of Fig. 2b under policies π_{WC} and some other policy π , respectively.

A. Variable Packet Lengths: Here we treat systems with variable length packets. All packets are bounded by a maximum packet length $L_{\textit{max}}$. We show that the π_{WC} routing strategy provides the best worst-case performance guarantee of all policies for routing packets non-preemptively over multiple time varying servers.

Consider the class of non-preemptive routing strategies: policies which do not interrupt a packet once it has initiated processing at a server. If a policy π has full knowledge of future events, it can be designed to minimize unfinished work at some particular time τ . Let $\min_{X, \mu}(\tau)$ represent this minimum value of unfinished work in a system with input stream X and linespeeds $\mu_i(t)$.

If policy π does not have knowledge of the future, we say it is ‘‘non-predictive.’’ Consider a game where a scheduler makes routing decisions (according to some non-predictive policy π) while an adversary dynamically creates an input stream X and varies the linespeeds $\mu_1(t), \dots, \mu_N(t)$ in order to maximize the difference between $U_{\pi}(t)$ and $\min_{X, \mu}(t)$. The goal of the scheduler is to minimize the worst case deviation from optimality, i.e., minimize the value of $\max_{t \geq 0} \{U_{\pi}(t) - \min_{X, \mu}(t)\}$.

Theorem 1: (a) For all policies π (possibly predictive and preemptive), we have:

$$U_{WC}(t) \leq U_{\pi}(t) + (n-1)L_{max}. \quad (3)$$

(b) The work conserving policy π_{WC} is the minimax optimal non-predictive, non-preemptive routing strategy over the set of all possible inputs and linespeed variations.

Proof: Part (a) follows immediately from Lemmas 1 and 2. To establish that the policy π_{WC} is minimax optimal, it suffices to show that any non-predictive, non-preemptive policy π can be forced to meet the $(n-1)L_{max}$ bound. The idea is to force policy π to route maximum length packets to $(n-1)$ distinct servers, and then to trap these packets by setting the server rates to 0. Specifically, $(n-1)$ maximum length packets are sent at time 0. The adversary maintains a constant output rate of $\mu_i(t) = \mu$ for all servers i that have no packets within them. Whenever a packet is routed to a server under policy π , that server rate is set to 0. After $(n-1)L_{max}/\mu$ seconds have elapsed, there must be some server j that has remained empty for the full time interval, with an unused processing rate $\mu_j(t) = \mu$. Hence, an alternative routing scheme that sends all packets to this server would have allowed the system to be empty at this time, and so the $(n-1)L_{max}$ bound is met with equality. \square

B. Fixed Length Packets: In the case when all packets have fixed lengths, a version of inequality (3) can be established in terms of the number of packets $N(t)$ in the system.

Theorem 2: If all packets have fixed lengths L , then:

$$N_{WC}(t) \leq N_{\pi}(t) + n - 1. \quad (4)$$

Proof: Define a *completely busy period* as an interval of time when all servers of the parallel queue system are busy. The inequality holds whenever t is not within a completely busy period. Suppose now that t lies within a completely busy period, and let τ_B be the beginning time of this period. We have:

$$N_{WC}(t) = N_{WC}(\tau_B^-) + A(\tau_B^-, t) - D_{WC}(\tau_B^-, t) \quad (5)$$

$$N_{\pi}(t) = N_{\pi}(\tau_B^-) + A(\tau_B^-, t) - D_{\pi}(\tau_B^-, t) \quad (6)$$

where τ_B^- represents the time just before the arrival initiating the completely busy period,

$A(\tau_B^-, t)$ is the number of arrivals during the interval (τ_B^-, t) , and $D_{WC}(\tau_B^-, t)$, $D_{\pi}(\tau_B^-, t)$ respectively represent the total number of packet departures from the π_{WC} system and the π system.

The above departure functions are composed of individual terms D_{WC}^i and D_{π}^i representing departures from queue i in the π_{WC} and π systems, respectively (we suppress the (τ_B^-, t) arguments for ease of notation). During the time interval $[\tau_B^-, t]$ each queue of the π_{WC} system continuously processes fixed length packets. Thus, $D_{WC}^i \geq D_{\pi}^i - 1$ for all queues i , and equality holds only if there was a packet being processed by queue i under the π routing policy at time τ_B^- . Suppose there are k such cases (where $k \leq n$). Then, from (5) and (6):

$$\begin{aligned} N_{WC}(t) &\leq N_{WC}(\tau_B^-) + A(\tau_B^-, t) - D_{\pi}(\tau_B^-, t) + k \\ &= N_{WC}(\tau_B^-) + N_{\pi}(t) - N_{\pi}(\tau_B^-) + k \\ &\leq (n-1) + N_{\pi}(t) - k + k. \quad \square \end{aligned}$$

A technique similar to the one used in the proof of Theorem 1 shows that this $(n-1)$ bound is tight and is minimax optimal over all non-predictive, non-preemptive strategies.

Similar routing problems have been treated in [2,3,16] for systems with two heterogeneous servers ($n=2$) when packet inputs are Poisson. With this stochastic formulation, it is possible to prove that the optimal routing strategy for minimizing expected occupancy in the system has a threshold structure. A complex dynamic program can be developed to calculate the exact threshold function. However, here we find that the simple work conserving strategy π_{WC} ensures no more than 1 extra packet in the system compared to any other strategy at any time. Furthermore, this bound applies to arbitrary input processes $X(t)$.

V. SYSTEMS WITHOUT A PRE-QUEUE

A. Constant Rate Servers: The implementation of the work conserving policy π_{WC} for time varying servers uses a pre-queue to store packets until the next server becomes available. In many systems it is undesirable or even impossible to implement a pre-queue. Here we show that the same results can be obtained without a pre-queue if the server rates (μ_1, \dots, μ_n) are known constants.

Observation: For constant server rates (μ_1, \dots, μ_n) , the strategy of routing an incoming packet to the queue k with the smallest value of $U_k(t)/\mu_k$ (the π_{WC} strategy as described in Section II) is the same as the π_{WC} strategy described for time varying servers in Section III. Thus, a pre-queue is not needed.

Proof: The strategy always routes a new packet to the queue that will empty first. Thus, if there is ever an empty server i , there can be no more than 1 packet in each of the $(n-1)$ other queues. \square

Thus, the bounds in Theorems 1 and 2 apply to heterogeneous, constant rate servers when this routing method is used.

B. Time Varying Servers: Consider routing over a multi-queue system with time-varying processing rates when no pre-queue is available and all routing decisions are made immediately upon packet arrival. We examine the Join-the-Shortest-Queue policy π_{JSQ} : route the current packet to the queue i with the smallest value of $U_i(t)$. Intuitively, this strategy is the closest match to the work conserving strategy given that we cannot predict future values of server rates.

We seek to prove that this strategy stabilizes the multi-queue system whenever it is stabilizable. This is done for general ergodic input sources and linespeed processes by introducing a new notion of stability defined in terms of finite buffer systems. Consider the single queue system of Fig. 3a with ergodic input process $X(t)$ of bit rate λ , linespeed process $\mu(t)$ with ergodic rate μ_{av} , and assume this system has a finite buffer capacity of M bits. Let $D(M)$ represent the drop rate (measured in bits) of the system as a function of buffer space M .

Definition: A system is *loss rate stable* if the drop rate can be made arbitrarily small by increasing buffer capacity, i.e., if $D(M) \rightarrow 0$ as $M \rightarrow \infty$.

It can be shown that a necessary condition for loss rate stability is $\lambda \leq \mu_{av}$. Furthermore, if the input stream and server rate process evolve according to an underlying finite state Markov chain, a sufficient condition for loss rate stability is $\lambda < \mu_{av}$. This notion of stability is closely tied to the standard notion defined in terms of a vanishing complementary occupancy distribution for infinite buffer capacity queues.




Figure 3: A single stage system with an aggregated server processing rate $\mu(t)$ compared to a multi-queue system operating under the JSQ routing strategy.

Consider the two systems of Fig. 3. The server rates $(\mu_1(t), \dots, \mu_n(t))$ of the multi-queue system sum to the single-server linespeed $\mu(t)$. Let $D_{JSQ}(M)$ represent the drop rate of the multi-queue system operating under the JSQ policy if all queues have finite buffer storage M .

Theorem 3: For all M , $D_{JSQ}(M + nL_{max}) \leq D_{single-queue}(M)$, and hence the multi-queue system under JSQ is stable whenever the single queue system is stable.

This theorem is a consequence of the following lemma:

Lemma 3: Suppose the input stream $X(t)$ and linespeed process $\mu(t)$ are such that the unfinished work $\tilde{U}(t)$ in the single-server system of Fig. 3a is always less than or equal to M . Then:

$$U^{JSQ}_i(t) \leq M + (n - 1)L_{max} \quad (7)$$

for all queues i and for all time $t \geq 0$.

Proof: Let t_1 represent any time when at least one queue of the multi-queue system under JSQ is empty, and suppose (7) holds at this time. Let t_2 represent the next time one of the queues in the system empties after all queues have been busy (if no such time exists, let $t_2 = \infty$). If the n queues are never simultaneously busy, then (7) holds for all time and we are done. Otherwise, there is some initial activation time t_a ($t_1 < t_a < t_2$) that all queues become simultaneously active. The result holds for $t \in [t_1, t_a]$. For all times $t \in [t_a, t_2]$, the multi-queue system is processing bits at the full rate $\mu(t)$, and hence the amount of bits which depart from the multi-queue system is greater than or equal to the departures from the single queue system during this interval. We thus have:

$$\begin{aligned} \sum U_i(t) &\leq \sum U_i(t_a^-) + \tilde{U}(t) - \tilde{U}(t_a^-) \\ &\leq (n - 1)(M + (n - 1)L_{max}) + M - 0 \\ &= nM + (n - 1)^2 L_{max} \end{aligned} \quad (8)$$

Now, if there is a queue i such that $U_i(t) > M + (n - 1)L_{max}$ at some time $t \in [t_a, t_2]$, there must be a time $\tau < t$ when the last packet (of size L) arrived to queue i such that:

$$U_i(\tau^-) > M + (n - 1)L_{max} - L \quad (9)$$

$$U_j(\tau^-) > M + (n - 1)L_{max} - L \quad (\text{for all } j \neq i) \quad (10)$$

where (10) follows from (9) and the fact that a packet is added to queue i under the JSQ strategy at

time τ^+ . Hence at time τ^+ :

$$\sum U_i(\tau^+) > nM + (n-1)^2 L_{max} \quad (11)$$

which contradicts (8). Iteration over all such intervals $[t_1, t_2]$ proves the lemma. \square

Using the above lemma, it can be shown that the multi-queue system (under JSQ) with a buffer size of $M + nL_{max}$ in each queue drops fewer bits than the single queue system with buffer size M .

VI. JOINT ROUTING AND POWER ALLOCATION

Suppose that the transmission rates (μ_1, \dots, μ_n) of the system can be controlled by adjusting the power levels p_i allocated to each server. Specifically, suppose that each channel $i \in \{1, 2, \dots, n\}$ has an associated channel state c_i which takes values on a finite set of states

($c_i \in \{S^i_1, S^i_2, \dots, S^i_{M_i}\}$). Let $\mu_i(p_i, c_i)$ represent a concave rate-power curve for each channel state (see Fig. 4).

Let $\vec{C}(t) = (c_1(t), c_2(t), \dots, c_n(t))$ represent the vector of channel states at time t , and assume that $\vec{C}(t)$ varies according to a discrete time Markov chain with timeslots of length T . As before, packets enter the system according to an input process $X(t)$ and routing decisions are made immediately upon arrival. In addition to making routing decisions, a controller must choose a power allocation $\vec{P}(t) = (p_1(t), \dots, p_n(t))$ for each instant of time, subject to a total power constraint $\sum p_i \leq P_{tot}$. The problem is to design a joint routing and power allocation strategy that maximizes throughput and stabilizes the system whenever the system is stabilizable. Such a policy makes decisions using the observable system state vectors $\vec{U}(t)$ and $\vec{C}(t)$.

In general, both state vectors $\vec{U}(t)$ and $\vec{C}(t)$ are important in both the routing and power allocation decisions. For example, clearly any power allocated to an empty queue is wasted and should be re-allocated to improve processing rates amongst the non-empty queues. Likewise, a router is inclined to place packets in faster queues, especially if the rates of those queues are guaranteed to operate at high levels for one or more timeslots. However, below we show that the routing and power allocation problem can be *decoupled* into two policies: a routing policy which considers only $\vec{U}(t)$, and a power allocation policy which considers only $\vec{C}(t)$. The resulting strategy maximizes total system throughput even when the underlying Markov chain describing $\vec{C}(t)$ is unknown.

Let λ represent the bit rate of input stream X , and let $\xi_{\vec{C}}$ represent the steady-state probability




Figure 4: A set of concave power curves $\mu_i(p_i, c_i)$ for channel states c_1, c_2, c_3 .

that the channel vector is in state \vec{c} .

Observation: The maximum throughput possible in the system is $\bar{\mu}$, where:

$$\bar{\mu} = \sum_{\vec{c}} \xi_{\vec{c}} [\max_{p_1+\dots+p_n=P_{tot}} \sum \mu_i(p_i, c_i)], \quad (10)$$

The value of $\bar{\mu}$ is the average total rate offered by the system when power is allocated to maximize total rate at every instant of time.

Assume that the input process X generates packets according to a finite state Markov chain, and let λ represent the total bit rate of stream X .

Theorem 4: The capacity of the multi-queue system with joint routing and power allocation is $\bar{\mu}$, i.e., a necessary condition for stability is $\lambda \leq \bar{\mu}$, and a sufficient condition is $\lambda < \bar{\mu}$.

The fact that $\lambda \leq \bar{\mu}$ is necessary is clear. Sufficiency is established by design of the following decoupled policy π^* which stabilizes the system whenever $\lambda < \bar{\mu}$:

Power Allocation: At every new timeslot, observe the channel state $\vec{c}(t)$ and allocate power $\{p_i(t)\}$ to maximize $\sum \mu_i(p_i, c_i(t))$ subject to $\sum p_i = P_{tot}$. Keep this power allocation constant until the next timeslot.

Routing: Whenever a new packet enters the system, we observe the value of $\vec{U}(t) = (U_1(t), \dots, U_n(t))$ and route to the shortest queue.

Proving that this strategy π^* achieves maximum throughput is now a straightforward application of the results of the previous sections. The method of power allocation maximizes the sum rate $\mu(t) = \mu_1(p_1(t), c_1(t)) + \dots + \mu_n(p_n(t), c_n(t))$ at every instant of time. From Lemma 1 it follows that a single server system with a server of rate $\mu(t)$ (and the same input stream X) outperforms the multi-queue system under any routing and power allocation strategy. Hence, whenever it is possible to stabilize the multi-queue system (by using *some* power allocation and routing strategy) the single queue system will be stable.

Furthermore, the routing strategy is simply an application of JSQ routing in reaction to the rates determined by the power allocation decisions. Thus, from Theorem 3 we know that the multi-queue system is stable whenever the single queue system is stable. This establishes Theorem 4.

VII. CONCLUSIONS

The problem of routing packets from an arbitrary stream over a set of parallel queues has been considered in the context of constant and time-varying processing rates. Using sample path analysis, a simple work conserving strategy was developed to provide fundamental performance bounds on the unfinished work in the system at any instant of time. In time varying systems, this strategy can be implemented with a pre-queue and guarantees that performance closely follows the performance of a superior single-queue system with an aggregated data rate.

The pre-queue was shown to be unnecessary when service rates are constant. In the case of time-varying rates, removing the pre-queue precludes the design of a routing strategy which meets the tight performance bounds guaranteed by a work conserving policy. However, a general stability result was established for the Join-the-Shortest-Queue policy, and the result was extended to treat a joint problem of routing and power allocation. This analysis was performed using a new and useful notion of stability defined in terms of finite buffer systems.

Our approach differs significantly from other approaches in that we provide tight, worst case bounds on system performance with arbitrary input and linespeed processes, rather than analyzing systems with particular stochastic inputs and linespeeds. We believe this approach can be applied to more complex queueing structures to provide performance bounds and stability guarantees for systems with very general input processes and control laws.

REFERENCES:

- [1] L. Tassiulas and A. Ephremides, "Dynamic Server Allocation to Parallel Queues with Randomly Varying Connectivity," *IEEE Trans. on Information Theory*, vol.39, no.2, March 1993.
- [2] A. Ephremides, P. Varaiya, and J. Walrand, "A simple Dynamic Routing Problem," *IEEE Trans. on Automatic Control*, vol.25, no.4, Aug. 1980.
- [3] I. Viniotis and A. Ephremides, "Extension of the Optimality of the Threshold Policy in Heterogeneous Multiserver Queueing Systems," *IEEE Trans. on Automatic Control*, vol.33, no.1, Jan. 1988.
- [4] Bruce Hajek, "Optimal Control of Two Interacting Service Stations," *IEEE Trans. on Automatic Control*, vol.29,no.6, June 1984.
- [5] P. Wu and M.J.M. Posner, "A level-crossing approach to the solution of the shortest-queue problem," *Operations Research Letters* 21, 1997.
- [6] L. Tassiulas and A. Ephremides, "Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, Dec. 1992.
- [7] M. Carr and Bruce Hajek, "Scheduling with Asynchronous Service Opportunities with Applications to Multiple Satellite Systems," *IEEE Trans. on Automatic Control*, vol.38, no.12, Dec. 1993.
- [8] M.J.Neely and E. Modiano, "Convexity and Optimal Load Distributions in Work Conserving $\infty/1$ Queues," IEEE Infocom Proceedings, 2001.
- [9] M.J.Neely, E. Modiano, and C. Rohrs, "Power and Server Allocation in a Multi-Beam Satellite with Time Varying Channels," MIT LIDS Report: P-2515 Sept. 2001.
- [10] R.L.Graham, "Bounds on Multiprocessing Timing Anomalies," *SIAM J. Appl. Math.*, vol.17,no.2, March 1969.
- [11] E.G. Coffman, M.R. Garey, and D.S.Johnson, "An Application of Bin-Packing to Multi-processor scheduling," *SIAM J. of Comput.*, vol.7, no.1, Feb. 1978.
- [12] D. B. Shmoys, J. Wein, and D.P. Williamson, "Scheduling Parallel Machines On-line," IEEE 32nd Ann. Symp. on Foundations of Comp. Sci, 1991.
- [13] Y.Bartal, A.Fiat, H.Karloff, and R. Vohra, "New Algorithms for an Ancient Scheduling Problem," Proc. of the 24th Annual ACM Symp. on the Theory of Computing, 1992.
- [14] D.R. Karger, S.J.Phillips, E.Torng, "A Better Algorithm for an Ancient Scheduling Problem," Proc. of the 5th Annual ACM-SIAM Symp. on Discrete Algorithms, 1994.
- [15] Tom Leighton, "Methods for Message Routing in Parallel Machines," Proc. of the 24th Annual ACM Symp. on the Theory of Computing, 1992.
- [16] Jean Walrand, *An Introduction to Queueing Networks*. Prentice Hall: NJ.
- [17] Sheldon Ross, *Stochastic Processes*. New York: John Wiley & Sons, Inc. 1996.