

Power Allocation and Routing in Multibeam Satellites With Time-Varying Channels

Michael J. Neely, Eytan Modiano, and Charles E. Rohrs, *Senior Member, IEEE*

Abstract—We consider power and server allocation in a multibeam satellite downlink which transmits data to N different ground locations over N time-varying channels. Packets destined for each ground location are stored in separate queues and the server rate for each queue i depends on the power $p_i(t)$ allocated to that server and the channel state $c_i(t)$ according to a concave rate-power curve $\mu_i(p_i, c_i)$. We establish the capacity region of all arrival rate vectors $(\lambda_1, \dots, \lambda_N)$ which admit a stabilizable system. We then develop a power-allocation policy which stabilizes the system whenever the rate vector lies within the capacity region. Such stability is guaranteed even if the channel model and the specific arrival rates are unknown. Furthermore, the algorithm is shown to be robust to arbitrary variations in the input rates and a bound on average delay is established. As a special case, this analysis verifies stability and provides a performance bound for the *Choose-the-K-Largest-Connected-Queues* policy when channels can be in one of two states (ON or OFF) and K servers are allocated at every timestep ($K < N$). These results are extended to treat a joint problem of routing and power allocation in a system with multiple users and satellites and a throughput maximizing algorithm for this joint problem is constructed. Finally, we address the issue of interchannel interference and develop a modified policy when power vectors are constrained to feasible activation sets. Our analysis and problem formulation is also applicable to power control for wireless systems.

Index Terms—Delay, dynamic power allocation, power control, queueing analysis, satellite communication, stability, wireless downlink.

I. INTRODUCTION

IN THIS PAPER, we consider power allocation in a satellite which transmits data to N ground locations over N different downlink channels. Each channel is assumed to be time varying (e.g., due to changing weather conditions) and the overall channel state is described by the ergodic vector process $\vec{C}(t) = (c_1(t), \dots, c_N(t))$. Packets destined for ground location i arrive from an input stream X_i and are placed in an output queue to await processing (Fig. 1). The servers of each of the N output queues may be activated simultaneously at any time t by assigning to each a power level $p_i(t)$, subject to the total power constraint $\sum p_i(t) \leq P_{\text{tot}}$. The transmission rate of each server i depends on the allocated power $p_i(t)$ and on the current channel state $c_i(t)$ according to a general concave rate-power curve $\mu_i(p_i, c_i)$. A controller allocates power to each of the N queues at every instant of time in reaction to

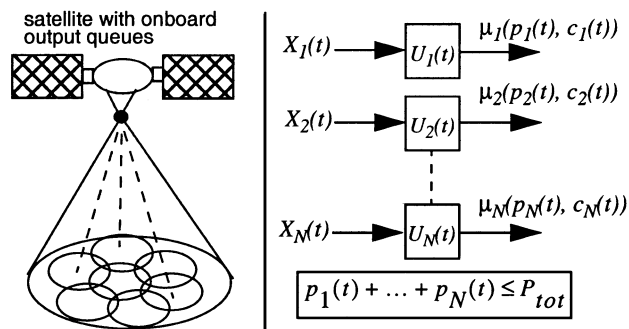


Fig. 1. Multibeam satellite with N time-varying downlink channels and N onboard output queues.

channel state and queue backlog information. The goal of the controller is to stabilize the system and thereby achieve maximum throughput and maintain acceptably low levels of unfinished work in all of the queues.

We establish the capacity region of the system by describing the multidimensional region of all arrival rate vectors $(\lambda_1, \dots, \lambda_N)$ which admit a stabilizable system under some power-allocation policy. Stability in this region holds for general ergodic channel and packet arrival processes. It is shown that if the channel model and arrival rates are known, any power-allocation policy which stabilizes the system—possibly by making use of special knowledge of future events—can be transformed into a stabilizing policy which considers only the current channel state.

We next consider the case of a slotted time system when arrivals and channel-state vectors are independent and identically distributed (i.i.d.) from one timeslot to the next, but the channel probabilities and the exact values of arrival rates $(\lambda_1, \dots, \lambda_N)$ are unknown. A particular power-allocation policy is developed which stabilizes the system whenever the rates $(\lambda_1, \dots, \lambda_N)$ are within the capacity region.¹ This policy is shown to maintain average queue occupancy within a fixed upper bound and is robust to arbitrary changes in the input rates. These results are extended to treat a joint routing and power-allocation problem with multiple users and multiple satellites and a simple policy is developed which maximizes throughput and ensures stability whenever the system is stabilizable. Finally, we address the issue of interchannel interference due to bandwidth limitations and develop a modified policy when power vectors are constrained to activation sets. This analysis makes use of a Lyapunov function defined over the state of the queues.

¹In [20], we show the same policy is stabilizing for Markov modulated input and channel dynamics (see simulation results in Section VII).

Manuscript received January 3, 2002; revised April 25, 2002; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Ajmone Marsan.

The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: mjneely@mit.edu; modiano@mit.edu; crohrs@mit.edu).

Digital Object Identifier 10.1109/TNET.2002.808401

Previous work on queue control problems for satellite and wireless applications is found in [1]–[8], [14], [15], [18], and [20]. In [2], a parallel queue system with a single server is examined, where every timeslot the transmit channels of the queues vary between ON and OFF states and the server selects a queue to service from those that are ON. The capacity region of the system is developed when packet arrivals and channel states are i.i.d. Bernoulli processes, and stochastic coupling is used to show optimality of the *Serve-the-Longest-Connected-Queue* policy in the symmetric situation that arrival and channel processes are identical for all queues (i.e., $\lambda_1 = \dots = \lambda_N$, $p_1^{\text{on}} = \dots = p_N^{\text{on}}$). Such a server-allocation problem can be viewed as a special case of our power-allocation formulation, and in Section IV, we verify stability of the *Serve-the-K-Longest-Queues* policy for symmetric and asymmetric systems with multiple servers.

In [3], a wireless network of queues is analyzed when input packets arrive according to memoryless processes and have exponentially distributed length. A Lyapunov function is used to establish a stabilizing routing and scheduling policy under network connectivity constraints. In [4], similar analysis is used to treat server allocation in a network with time-varying connectivity. Such a technique for proving stability has also been used in the switching literature [9]–[12]. In [10], an $N \times N$ packet switch with blocking is treated and input/output matching strategies are developed to ensure 100% throughput whenever arrival rates are within the capacity region. In [12] and [13], the method of Lyapunov stability analysis is used to prove that queues are not only stable but have finite backlog moments.

The main contribution in this paper is the formulation of a general power-control problem for multibeam satellites and the development of throughput maximizing power- and server-allocation algorithms for the system. The method extends to other wireless networking problems where power allocation and energy efficiency is a major issue. Recent work in [14] treats a problem of minimizing the total energy expended to transmit blocks of data arriving to a single queue, and it is shown that power control can be effectively used to extend longevity of network elements. In [15], power allocation for wireless networks is addressed. The authors consider ON/OFF-type power-allocation policies and observe that for random networks, capacity regions are not extended much by including more power quantization levels. Our capacity results in Sections III and VII illustrate that the capacity region is often considerably extended if multiple power levels are utilized for the satellite downlink problem.

In the next section, we introduce the power- and server-allocation problems. In Section III, we develop several stability results for single-queue systems with ergodic and nonergodic processing rates $\mu(t)$ and establish the capacity region of the satellite downlink with power control. In Section IV, a stabilizing power-allocation policy is developed for systems with i.i.d. inputs and channel states. In Section V, a joint routing and power-allocation policy is treated using similar analysis, and in Section VI, we extend the problem to treat channel interference issues. Numerical results are presented in Section VII.

II. POWER AND SERVER ALLOCATION

Consider the N queue system of Fig. 1. Each time-varying channel i can be in one of a finite set of states S_i . We




Fig. 2. Set of concave power curves $\mu_i(p_i, c_i)$ for channel states c_1, c_2, c_3 .

represent the channel process by the channel vector $\vec{C}(t) = (c_1(t), \dots, c_N(t))$, where $\vec{C}(t) \in S_1 \times \dots \times S_N$. Channels hold their state for timeslots of length T , with transitions occurring on slot boundaries $t = kT$. It is assumed that channel states are known at the beginning of each timeslot. Such information can be obtained either through direct channel measurement (where timeslots are assumed to be long in comparison to the required measurement time) or through a combination of measurement and channel prediction.² The channel process is assumed to be ergodic and yields time-average probabilities $\pi_{\vec{C}}$ for each state \vec{C} . At every timeslot, the server transmission rates can be controlled by adjusting the power-allocation vector $\vec{P}(t) = (p_1(t), \dots, p_N(t))$ subject to the total power constraint $\sum p_i(t) \leq P_{\text{tot}}$.

For any given state c_i of downlink channel i , there is a corresponding rate–power curve $\mu_i(p_i, c_i)$ which is increasing, concave, and continuous in the power parameter (Fig. 2). This power curve could represent the logarithmic Shannon capacity curve of a Gaussian channel, or could represent a rate curve for a specific set of coding schemes designed to achieve a sufficiently low probability of error in the given channel state. In general, any practical set of power curves will have the concavity property, reflecting diminishing returns in transmission rate with each incremental increase in signal power.

The continuity property is less practical. A real system will rely on a finite databank of coding schemes and, hence, actual rate–power curves restrict operation to a finite set of points. For such a system, we can create a new, *virtual power curve* $\tilde{\mu}_i(p_i, c_i)$ by a piecewise linear interpolation of the operating points [see Fig. 3(a)]. Such virtual curves have the desired continuity and concavity properties and are used as the true curves in our power-allocation algorithms. Clearly, a virtual system which allocates power according to the virtual curves has a capacity region which contains that of a system restricted to allocate power on the vertex points. However, when vertex points are equally spaced along the power axis and integrally divide the total power P_{tot} , the capacity regions are, in fact, the same, as any point on a virtual curve can effectively be achieved by time-averaging two or more feasible rate–power points over many timeslots. Indeed, in Section IV, we design a stabilizing policy for any set of concave power curves which naturally selects vertex points at every timeslot if power curves are piecewise linear.

This power-allocation formulation generalizes a simpler problem of server allocation. Assume that there are K servers

²In [23] and [24], it is shown that satellite channel states can be accurately predicted up to one second into the future, where attenuation levels in the Ka band are estimated to within ± 1 db of accuracy in both clear and rainy weather conditions.




Fig. 3. Virtual power curves for systems with a finite set of operating points.

and every timeslot the servers are scheduled to serve K of the N queues ($K < N$). A given queue i transmits data at a fixed rate μ_i whenever a server is allocated to it and transmits nothing when no server is allocated. This problem can be transformed into a power-allocation problem by defining the virtual power constraint $\sum p_i(t) \leq K$ and the virtual power curves

$$\tilde{\mu}_i(p) = \begin{cases} \mu_i p, & p \in [0, 1] \\ \mu_i, & p > 1. \end{cases} \quad (1)$$

Such a virtual curve contains the feasible points ($p=0, \tilde{\mu}_i=0$) and ($p=1, \tilde{\mu}_i=\mu_i$), corresponding to a server being either allocated or not allocated to queue i [see Fig. 3(b)]. However, it suffices to remove this feasible point restriction and treat the system as if it operates according to the continuous virtual power curve (1). This preserves the same capacity region, and later it is shown that any stabilizing algorithm which uses the virtual curves can be transformed into a stabilizing algorithm which conforms to the feasible point restriction.

Example Server-Allocation Algorithm: One might suspect the policy of serving the K fastest nonempty queues would maximize data output and achieve stability. However, we provide the following counterexample which illustrates that this is not the case. Consider a three-queue two-server system with constant processing rates $(\mu_1, \mu_2, \mu_3) = (1, 1, 1/2)$. All arriving packets have length $L = 1$ and arrive according to i.i.d. Bernoulli processes with packet arrival probabilities $(p_1, p_2, p_3) = (p, p, (1-p^2)/2 + \varepsilon)$, where $p < 1/2$ and $0 < \varepsilon < p^2/2$.

Note that the policy of serving the two fastest nonempty queues removes a server from queue 3 whenever there are simultaneous arrivals at queues 1 and 2. This happens with probability p^2 and, hence, the time-average processing rate at queue 3 is no more than $(1-p^2)/2$ (where the factor $1/2$ is due to the rate of server 3). This effective service rate cannot support the input rate and, hence, queue 3 is unstable under this server-allocation policy. However, the system is clearly stabilizable: The policy of always allocating a server to queue 3 and using the remaining server to process packets in queues 1 and 2 stabilizes all queues.

III. STABILITY AND THE DOWNLINK CAPACITY REGION

To understand the capacity region of the downlink system, we first develop a simple criterion for stability of a single queue with an input stream $X(t)$ and a time-varying processing rate $\mu(t)$. We assume the input stream is ergodic with rate λ . However, because an arbitrary power-control scheme could potentially yield a nonergodic processing rate, we must consider gen-

eral processes $\mu(t)$ which may or may not have well-defined time averages. We make the following definitions.

- $X(t)$ total amount of bits that arrived during $[0, t]$;
- $U(t)$ unprocessed bits in the queue at time t ;
- $\mu(t)$ instantaneous bit processing rate in the server.

$$\lambda = \lim_{t \rightarrow \infty} \frac{X(t)}{t}, \quad \underline{\mu} = \liminf_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mu(\tau) d\tau \quad (2)$$

where \liminf of a function $f(t)$ is defined as $\liminf_{t \rightarrow \infty} f(t) = \lim_{t \rightarrow \infty} [\inf_{\tau \geq t} f(\tau)]$. The \limsup is defined similarly.

The limits in (2) exist with probability 1. We assume the processing rate is always bounded above by some maximum value ($\mu(t) \leq \mu_{\max}$, for all t) and, hence, $0 \leq \underline{\mu} \leq \mu_{\max}$. As a measure of the fraction of time that the unfinished work in a queue is above a certain value M , we define the ‘‘overflow’’ function $g(M)$

$$g(M) = \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t 1_{[U(\tau) > M]} d\tau \quad (3)$$

where the indicator function 1_E used above takes the value 1 whenever event E is satisfied and 0 otherwise.

Definition: A single-server queueing system is *stable* if $g(M) \rightarrow 0$ as $M \rightarrow \infty$.

Notice that if sample paths of unfinished work in the queue are ergodic and a steady state exists, the overflow function $g(M)$ is simply the steady-state probability that the unfinished work in the queue exceeds the value M . Stability in this case is identical to the usual notion of stability defined in terms of a vanishing complementary occupancy distribution [3], [10], [13], [17].

Lemma 1 (Queue Stability): For a single-queue system with general input and server rate processes $X(t)$ and $\mu(t)$, a necessary condition for stability is $\lambda \leq \underline{\mu}$. If the arrival process $X(t)$ and the rate process $\mu(t)$ evolve according to an ergodic finite-state Markov chain, then a sufficient condition for stability is $\lambda \leq \underline{\mu}$.

Proof: The sufficient condition for Markovian arrivals and linespeeds is well known (see large deviations results in [16]). The necessary condition is proven in the Appendix by showing that if $\lambda > \underline{\mu}$, there exist arbitrarily large times t_i such that the average fraction of time that the unfinished work is above M during $[0, t_i]$ is greater than a fixed constant for any value of M . \square

We use this single-queue result to establish the capacity region of the power-constrained multichannel system of Fig. 1. We define the capacity region as the compact set of points $\Omega \subset [0, \infty)^N$ such that all queues of the system can be stabilized (with some power-allocation policy) whenever the vector of input bit rates $\vec{\lambda} = (\lambda_1, \dots, \lambda_N)$ is strictly in the interior of Ω and, conversely, no stabilizing policy exists whenever $\vec{\lambda} \notin \Omega$. (The system may or may not be stable if $\vec{\lambda}$ lies on the boundary of the capacity region.)

Assume arrivals and channel states are modulated by an ergodic finite-state Markov chain and transitions occur on timeslots of duration T . Let $\pi_{\vec{C}}$ represent the steady-state probability that the channel vector is in state $\vec{C} = (c_1, \dots, c_N)$.

Theorem 1 (Downlink Capacity): The capacity region of the downlink channel of Fig. 1 with power constraint P_{tot} and rate–power curves $\mu_i(p_i, c_i)$ is the set of all input rate vectors $\vec{\lambda}$ such that there exist power levels $p_i^{\vec{C}}$ satisfying $\sum_i p_i^{\vec{C}} \leq P_{\text{tot}}$ for all channel states \vec{C} and such that

$$\lambda_i \leq \sum_{\vec{C}} \pi_{\vec{C}} \mu_i(p_i^{\vec{C}}, c_i), \quad i \in \{1, \dots, N\}. \quad (4)$$

Proof: Using the stationary policy of allocating a power vector $\vec{P}^{\vec{C}} = (p_1^{\vec{C}}, \dots, p_N^{\vec{C}})$ whenever the system is in channel state \vec{C} creates a Markov modulated processing rate $\mu_i(t)$ for all queues i , with an average rate given by the right-hand side of inequality (4). Thus, Lemma 1 ensures stability whenever the vector $\vec{\lambda}$ satisfies (4) with strict inequality in all entries. We now show that restricting power control to such stationary policies (which use only the current channel state $\vec{C}(t)$ when making power-allocation decisions) does not restrict the capacity region and, hence, the region in (4) captures all input rates which yield stable systems.

Suppose all queues of the downlink channel can be stabilized with some power-control function $\vec{P}(t)$ which meets the power constraints—perhaps a function derived from a policy which knows future events. From the necessary condition of Lemma 1, we know that the \liminf of the resulting rate process satisfies $\lambda_i \leq \underline{\mu}_i$ for all queues $i \in \{1, \dots, N\}$.

We upper-bound $\underline{\mu}_i$ as follows. Let $T_{\vec{C}}(t)$ represent the subintervals of $[0, t]$ during which the channel is in state \vec{C} and let $\|T_{\vec{C}}(t)\|$ denote the total length of these subintervals. Fix $\varepsilon > 0$ and let $|\vec{C}|$ represent the total number of channel states of the system. Because the channel process is ergodic and because there are a finite number of queues and channel states, there exists a time \tilde{t} such that the time-average fraction of time in each channel state and the time-average processing rate of all queues are simultaneously within ε of their limiting values:

$$\frac{\|T_{\vec{C}}(\tilde{t})\|}{\tilde{t}} \leq \pi_{\vec{C}} + \varepsilon, \quad \text{for all channel states } \vec{C} \quad (5)$$

$$\underline{\mu}_i \leq \frac{1}{\tilde{t}} \int_0^{\tilde{t}} \mu_i(p_i(\tau), c_i(\tau)) d\tau + \varepsilon, \quad \text{for all } i \in \{1, \dots, N\}. \quad (6)$$

Thus, under power decisions $\vec{P}(t)$, we have, for all i

$$\lambda_i \leq \underline{\mu}_i \leq \sum_{\vec{C}} \frac{\|T_{\vec{C}}(\tilde{t})\|}{\tilde{t}} \frac{1}{\|T_{\vec{C}}(\tilde{t})\|} \int_{\tau \in T_{\vec{C}}(\tilde{t})} \mu_i(p_i(\tau), c_i) d\tau + \varepsilon \quad (7)$$

$$\leq \sum_{\vec{C}} \frac{\|T_{\vec{C}}(\tilde{t})\|}{\tilde{t}} \mu_i \left(\frac{1}{\|T_{\vec{C}}(\tilde{t})\|} \int_{\tau \in T_{\vec{C}}(\tilde{t})} p_i(\tau) d\tau, c_i \right) + \varepsilon \quad (8)$$

$$\leq \sum_{\vec{C}} (\pi_{\vec{C}} + \varepsilon) \mu_i \left(\frac{1}{\|T_{\vec{C}}(\tilde{t})\|} \int_{\tau \in T_{\vec{C}}(\tilde{t})} p_i(\tau) d\tau, c_i \right) + \varepsilon \quad (9)$$

where (8) follows from concavity of the $\mu_i(p, c_i)$ functions with respect to the power variable p and (9) follows from (5). We define for all states \vec{C} and queues i

$$\hat{p}_i^{\vec{C}} = \frac{1}{\|T_{\vec{C}}(\tilde{t})\|} \int_{\tau \in T_{\vec{C}}(\tilde{t})} p_i(\tau) d\tau + \varepsilon. \quad (10)$$

Hence, from (9) and (10)

$$\lambda_i \leq \sum_{\vec{C}} \pi_{\vec{C}} \mu_i(\hat{p}_i^{\vec{C}}, c_i) + \varepsilon \left(1 + |\vec{C}| \mu_{\text{max}}\right) \quad (11)$$

where μ_{max} is defined as the maximum processing rate of a queue (maximized over all queues and channel states) when it is allocated the full power P_{tot} .

Because the original power function satisfies the power constraint $\sum_i p_i(t) \leq P_{\text{tot}}$ for all times t , from (10) it is clear that the $\hat{p}_i^{\vec{C}}$ values satisfy the constraint $\sum_i \hat{p}_i^{\vec{C}} \leq P_{\text{tot}}$ for all channel states \vec{C} . Thus, (11) indicates that the arrival vector is arbitrarily close to a point in the region specified by (4). Because the region (4) is closed, it must contain $\vec{\lambda}$ and, hence, (4) represents the capacity region of the system. \square

In the case when the channel does not vary but stays fixed, we have one power curve $\mu_i(p)$ for each queue i and the expression for the downlink capacity region in Theorem 1 can be greatly simplified, as follows.

Corollary 1.1 (Static Channel Capacity): The capacity region for static channels is the set of all $\vec{\lambda}$ vectors such that

$$\sum_{i=1}^N \mu_i^{-1}(\lambda_i) \leq P_{\text{tot}}$$

where

$$\mu_i^{-1}(\lambda_i) = \begin{cases} \text{The smallest } p \text{ such that } \mu_i(p) = \lambda_i \\ \infty, & \text{if no such } p \text{ exists.} \end{cases}$$

\square

In Fig. 4(a), we illustrate a general capacity region for $N = 2$ channels with fixed channel states and concave power curves $\mu_1(p) + \mu_2(p)$. In this case of fixed channel states, one might suspect the optimal solution to be the one which maximizes the instantaneous output rate at every instant of time: Allocate full power to one queue whenever the other is empty and allocate power to maximize the sum output rate $\mu_1(p_1) + \mu_2(p_2)$ subject to $p_1 + p_2 \leq P_{\text{tot}}$ whenever both queues are full. Doing this restricts the capacity region to linear combinations of the three operating points, as illustrated in Fig. 4(a). The shaded regions in the figure represent the capacity gains obtained by power allocation using the full set of power levels. Note that the region is restricted further if only ON/OFF allocations are considered.

Corollary 1.2 (Server-Allocation Capacity): For the K -server-allocation problem where the channel rate of queue i is μ_i when it is allocated a server (and 0, otherwise), the capacity region is the polytope set of all $\vec{\lambda}$ vectors such that

$$\sum_i \frac{\lambda_i}{\mu_i} \leq K \quad (12)$$

$$\lambda_i \in [0, \mu_i], \quad i \in \{1, \dots, N\}. \quad (13)$$

Proof: Using the virtual power curves and constraints given in Section II, we find by Corollary 1.1 that the polytope region described by (12) and (13) contains the true capacity




Fig. 4. Capacity regions for static channels. (a) Two-queue system with power allocation. (b) K -server-allocation problem with $K = 2$, $N = 3$.

region. However, the K -server problem is constrained to allocate rates only on the vertex points of the polytope [see Fig. 4(b)]. Timesharing among vertex points, however, achieves any desired point within the polytope. \square

IV. STABILIZING POWER-ALLOCATION ALGORITHM

Theorem 1 implies that stability of the downlink channel can be achieved by a stationary power-allocation policy which allocates power levels $p_i^{\vec{C}}$ whenever the channel is in state \vec{C} . Such power levels can, in principle, be calculated with full knowledge of arrival rates λ_i and channel-state probabilities $\pi_{\vec{C}}$. However, such computation is impractical if the number of channel states is large and cannot be done if the arrival and channel-state probabilities are unknown. Here, we develop a power-allocation policy which stabilizes the system at every point of the capacity region (4) without using the arrival and channel-state probabilities. In essence, the policy learns the system parameters indirectly by basing power-allocation decisions both on channel state and queue backlog information. Furthermore, because the policy is not bound to a particular set of system parameters, it is shown to be robust to arbitrary changes in the input rates λ_i .

We assume that channel-state vectors \vec{C} vary i.i.d. from timeslot to timeslot with probability distribution $\pi_{\vec{C}}$. Likewise, assume that packets bring a new batch of unfinished work i.i.d. from timeslot to timeslot in the form of an arrival vector $\vec{A} = (a_1, \dots, a_N)$, with distribution $f(a_1, \dots, a_N)$ and expectation $E[\vec{A}] = \vec{\lambda}T$. Note that entries of the channel-state vector and the arrival vector may be correlated within the same timeslot. New arrivals are assumed to have bounded second moments: $E[a_i^2] < \infty$.

The i.i.d. assumption on channel-state variation facilitates system analysis and enables a simple bound on packet occupancy and delay to be calculated. However, note that this assumption does not necessarily accurately model satellite downlinks. Channel modeling experiments show that channel states could be modeled as i.i.d. during clear weather conditions (due to the observed rapid fluctuation of signal attenuation from scintillations in the Ka band [23]–[27]). However, in rainy weather, future channel states are highly dependent on the current state. In [23] and [24], it is shown that the channel-state variations can be modeled as a Markov process. In [20], we show that the control schemes developed here (derived using i.i.d. assumptions) offer the same stability properties under general Markov modulated arrival and channel dynamics, although the analysis and corresponding performance bounds are slightly more complex and are omitted from this paper for brevity.

Let $\vec{U}(t) = (U_1(t), \dots, U_N(t))$ represent the vector of unfinished work in each queue at time t (where $t = kT$). We assume channel- and queue-state vectors $\vec{C}(t)$ and $\vec{U}(t)$ are known at the beginning of each timeslot and seek a control policy which allocates power based on this information. Assuming this power allocation $\vec{P}(t)$ is held constant during the full timeslot $[t, t+T]$, the unfinished work dynamics proceed according to the one-step equation

$$U_i(t+T) = \max[U(t) - \mu_i(p_i(t), c_i(t))T, 0] + a_i(t). \quad (14)$$

Notice that for a given stationary power-allocation policy, the unfinished work vector at timestep $t+T$ is independent of the past, given the current value of unfinished work. Hence, the system can be viewed as evolving according to a Markov chain on an N -dimensional uncountably infinite state space \vec{U} .

A. Lyapunov Stability

For stability analysis, we define the Lyapunov function $L(\vec{U}) = \sum \theta_i U_i^2$ (for arbitrary positive weights $\{\theta_i\}$) and make use of a well-developed theory of stability in Markov chains using negative Lyapunov drift [3], [10], [12], [13], [17]. Below, we state a sufficient condition for the system to be stable and have a well-defined steady-state distribution of unfinished work \vec{U} . The proof is a straightforward extension to those given for countably infinite state spaces in [13] and [17] and is omitted for brevity.

Theorem 2 (Lyapunov Drift): For the given Lyapunov function $L(\vec{U})$, if there exists a compact region Λ of \mathfrak{R}^N and a number $\alpha > 0$ such that

- 1) $E[L(\vec{U}(t+T)) | \vec{U}(t)] < \infty$, for all $\vec{U} \in \mathfrak{R}^N$;
- 2) $E[L(\vec{U}(t+T)) - L(\vec{U}(t)) | \vec{U}(t)] \leq -\alpha$, whenever $\vec{U}(t) \notin \Lambda$;
- 3) whenever $\vec{U}(t) \in \Lambda$, there is a nonzero probability p that $\vec{U}(t+mT) = 0$ for some finite integer m ;

then a steady-state distribution on the vector \vec{U} exists (clearly, with the property that $\Pr[U_i > u] \rightarrow 0$ as $u \rightarrow \infty$, for all i) and, hence, the system is stable. \square

The first two conditions of Theorem 2 are similar to the conditions for Lyapunov stability in Markov chains with countably infinite state spaces. They ensure that, because of the negative drift, the mean recurrence time to the Λ region is finite. The third condition is a necessary modification to address systems with uncountably infinite state spaces. It implies that the zero state is reached infinitely often with finite mean recurrence times and ensures that the Markov chain reduces to a single ergodic class. Using renewal theory [21], it can be shown that the steady-state distribution is equal to the time-average integral of an indicator function over a sample path

$$\Pr[\vec{U} \leq \vec{u}] = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t 1_{[\vec{U}(\tau) \leq \vec{u}]} d\tau$$

where the limit exists with probability 1.

The first and third conditions of Theorem 2 are rather mild. The first is satisfied in the Markov chain for the downlink system because of the assumption $E[a_i^2] < \infty$, for all i . The third is satisfied as long as there is a probability of drifting negatively in one timestep for any $\vec{U} \in \Lambda$, such as when there is

a nonzero probability that no arrivals occur to any queue during a timeslot. We assume throughout that these properties hold. The second condition for negative drift is, thus, the most important for proving stability. In [12] and [13], it is shown that if a stronger drift condition is satisfied (such that the negative drift gets larger in magnitude as $|\vec{U}|$ increases), the moments of unfinished work are finite and can be bounded. Using a proof technique similar to that of [12], we have the following corollary to Theorem 2 for Lyapunov drift.

Corollary 2.1 (Performance Bound): Assuming condition 3 of Theorem 2 holds, if there exist positive values B and $\{\alpha_i\}$ such that

$$E[L(U(t+T)) - L(U(t))|U(t)] \leq B - \sum_i \alpha_i U_i(t) \quad (15)$$

then there exists a steady-state distribution with bounded first moments \bar{U}_i , such that $\sum \alpha_i \bar{U}_i \leq B$.

Proof: The proof relies on a telescoping series argument similar to the proof given in [12]. A full proof is provided in Theorem 4 (Section IV-D), where a more general time-varying Markov chain is considered. \square

B. Power-Allocation Policy

Consider now the following power-allocation policy for the downlink system. At the beginning of each timeslot, observe $\vec{U}(t)$ and $\vec{C}(t)$ and allocate a power vector $\vec{P}(t) = (p_1(t), \dots, p_N(t))$ (satisfying the power constraint) to maximize the quantity $\sum \theta_i U_i \mu_i(p_i, c_i)$, where $\{\theta_i\}$ is any set of positive weights. (If the weights are chosen to be different, the more heavily weighted queues can be given better delay guarantees, as described subsequently). Notice that the policy acts only through the current value of $\vec{U}(t)$ and $\vec{C}(t)$ without specific knowledge of the arrival rate vector $\vec{\lambda}$ or the channel-state probabilities. Intuitively, we desire a policy that gives more power to queues with currently high processing rates (to achieve maximum throughput) as well as giving more power to queues with large backlog (to ensure that these queues are stabilized). The above policy does both by considering as a metric the product of backlog and data rate for each queue.

Theorem 3 (Dynamic Power Allocation): The power-allocation policy of choosing the power vector

$$\vec{P}(t) = \arg \max_{\sum p_i \leq P_{\text{tot}}} \sum \theta_i U_i(t) \mu_i(p_i, c_i(t)) \quad (16)$$

stabilizes the system whenever the arrival rate vector $\vec{\lambda}$ is interior to the capacity region given by Theorem 1.

Proof: Consider the one-step drift in the Lyapunov function $L(\vec{U}) = \sum \theta_i U_i^2$ from Theorem 2. For ease of notation, let $U_i = U_i(t)$, $a_i = a_i(t)$ and let $\mu_i = \mu_i(p_i(t), c_i(t))$. From (14), we have

$$\begin{aligned} U_i^2(t+T) &\leq (U_i - \mu_i T)^2 + a_i^2 + 2a_i \max(U_i - \mu_i, 0) \\ &\leq (U_i - \mu_i T)^2 + a_i^2 + 2a_i U_i \\ &= U_i^2 - 2TU_i \left(\mu_i - \frac{a_i}{T} \right) + \mu_i^2 T^2 + a_i^2. \end{aligned} \quad (17)$$

From (17), it is clear that Property 1 of Theorem 2 holds. Now define the following constants:

$$\beta = \max_{\vec{C}, \sum_i p_i = P_{\text{tot}}} \left[\sum \theta_i \mu_i^2(p_i, c_i) \right] \quad (18)$$

$$B = \beta + \sum \theta_i E \left[\left(\frac{a_i}{T} \right)^2 \right]. \quad (19)$$

Taking conditional expectations of (17), scaling by weights θ_i and summing over all i , we have

$$\begin{aligned} E \left[L(\vec{U}(t+T)) - L(\vec{U}(t)) | \vec{U}(t) \right] \\ \leq T^2 B - 2T \sum \theta_i U_i \left(E \left[\mu_i | \vec{U}(t) \right] - \lambda_i \right) \end{aligned} \quad (20)$$

where the i.i.d. nature of the packet arrivals has been used in the identity $E[a_i/T | \vec{U}(t)] = \lambda_i$. Now, notice that the term $\sum \theta_i U_i E[\mu_i | \vec{U}(t)]$ maximizes the value of $\sum \theta_i U_i \gamma_i$ over all vectors $\vec{\gamma} = (\gamma_1, \dots, \gamma_N)$ in the capacity region (4). To see this, note that for any $\vec{\gamma}$ in the capacity region, there is a set of $\{p_i^{\vec{C}}\}$ values satisfying the power constraint such that

$$\sum_i \theta_i U_i \gamma_i \leq \sum_i \theta_i U_i \sum_{\vec{C}} \pi_{\vec{C}} \mu_i(p_i^{\vec{C}}, c_i) \quad (21)$$

$$= \sum_{\vec{C}} \pi_{\vec{C}} \sum_i \theta_i U_i \mu_i(p_i^{\vec{C}}, c_i) \quad (22)$$

$$\leq \sum_{\vec{C}} \pi_{\vec{C}} \max_{\sum p_i \leq P_{\text{tot}}} \left[\sum_i \theta_i U_i \mu_i(p_i, c_i) \right] \quad (23)$$

$$= \sum_{\vec{C}} \pi_{\vec{C}} \sum_i \theta_i U_i \mu_i^*(\vec{U}, \vec{C}) \quad (24)$$

where we define $\mu_i^*(\vec{U}, \vec{C})$ as the rate of queue i resulting from the given power-allocation policy (16) when the queue- and channel-state vectors are \vec{U} and \vec{C} , respectively. We thus have for all $\vec{\gamma}$ in the capacity region

$$\sum_i \theta_i U_i \gamma_i \leq \sum_i \theta_i U_i E \left[\mu_i | \vec{U}(t) \right]. \quad (25)$$

Now, because the arrival rate vector $\vec{\lambda}$ is assumed to be strictly in the interior of the capacity region, we can add a positive vector $\vec{\varepsilon} = (\varepsilon, \dots, \varepsilon)$ to produce another vector $(\vec{\lambda} + \vec{\varepsilon})$, which is in the capacity region. Hence, from (25) we have $\sum \theta_i U_i E[\mu_i | \vec{U}(t)] \geq \sum \theta_i U_i (\lambda_i + \varepsilon)$ and, thus

$$\sum \theta_i U_i \left(E \left[\mu_i | \vec{U} \right] - \lambda_i \right) \geq \varepsilon \sum \theta_i U_i. \quad (26)$$

Using (26) in (20), we find that

$$E \left[L(\vec{U}(t+T)) - L(\vec{U}(t)) | \vec{U}(t) \right] \leq T^2 B - 2T\varepsilon \sum \theta_i U_i. \quad (27)$$

Choose any number $\alpha > 0$ and define the compact region

$$\Lambda = \left\{ \vec{U} \in \mathbb{R}^N | U_i \geq 0, \sum \theta_i U_i \leq \left(\frac{T^2 B + \alpha}{2T\varepsilon} \right) \right\}.$$

We find from (27) that the Lyapunov drift is less than $-\alpha$ whenever $\vec{U} \notin \Lambda$. \square

Using the bounded first-moment result from Corollary 2.1, we find that the strong negative drift condition in (27) implies that the steady-state queue occupancies have bounded first moments.

Corollary 3.1 (Downlink Delay Bound): For the downlink system under the dynamic power-allocation algorithm of Theorem 3, the steady-state unfinished work has a finite mean \bar{U}_i for all queues i and satisfies

$$\sum \theta_i \bar{U}_i \leq \frac{TB}{2\varepsilon}. \quad (28)$$

Thus, by Little's Theorem, the average bit delay \bar{D}_i satisfies $\sum \theta_i \lambda_i \bar{D}_i \leq TB/(2\varepsilon)$. \square

The $1/\varepsilon$ behavior that this bound exhibits is worth noting. Recall that ε can be viewed as the minimum distance from the input rate vector $\vec{\lambda}$ to the boundary of the capacity region (in the sense that is chosen as the largest value such that $\vec{\lambda} + \vec{\varepsilon}$ remains in the capacity region). Thus, the bound grows asymptotically like $1/\varepsilon$ as the rate vector is pushed toward the boundary. Such behavior is characteristic of queueing systems, as illustrated by the standard $P-K$ formula for average occupancy in an $M/G/1$ queue [21].

In [20], we show the same policy guarantees stability and provides a bound on average delay when the input and channel processes are Markov modulated, provided that the steady-state input-rate vector $\vec{\lambda} = (\lambda_1, \dots, \lambda_N)$ is within a distance ε of the boundary of the capacity region (so that $\vec{\lambda} + \vec{\varepsilon} \in \Omega$). This holds because we can analyze the Lyapunov drift of the system every K timeslots, rather than every timeslot, where K is chosen such that time-average channel probabilities and data rates over K slots are within δ of their steady-state values (for some small value $\delta > 0$), regardless of the initial state of the Markov chain. Hence, the effects of the initial conditions and the dependencies introduced by the Markovian dynamics are negligible if the system is analyzed every K timesteps. Intuitively, such a K slot interval can be viewed as a "super-timeslot," and the resulting occupancy bound has a form similar to (28), exhibiting $1/\varepsilon$ behavior as well as a linear dependence on the super-timeslot length KT . The exact bound is computed in [20]. Similar techniques for stability analysis of Markovian systems are demonstrated in [4] and for fluid limits in [13].

Note that the positive weights $\{\theta_i\}$ in the dynamic power-allocation algorithm (16) can be chosen arbitrarily. Larger weights can be given to specific queues to improve their relative performance according to the downlink performance bound³ in Corollary 3.1. Choosing weights $\theta_i = 1$, for all i , yields a policy which chooses a power vector that maximizes $\sum U_i \mu_i$ at every timestep. The following corollary makes use of a different set of weights.

Consider again the K -server-allocation problem where each queue has only two channel states, ON or OFF, and these states vary i.i.d. over each timeslot as an N -dimensional vector. When a server is allocated to queue i while it is in the ON state, the server transmits data from the queue at a rate μ_i (the transmission rate is zero when in the OFF state or when no server is allo-

cated). Defining the virtual rate-power curves $\tilde{\mu}_i(\cdot)$ as in Section II, we have the following corollary:

Corollary 3.2 (Dynamic Server Allocation): For the K -server-allocation problem with ON/OFF channel states, the policy of allocating the K servers to the K longest ON queues stabilizes the system whenever the system is stabilizable.

Proof: Assume the system operates according to virtual power curves as in Section II (1) and define the Lyapunov function $L(\vec{U}) = \sum (U_i^2)/\mu_i$ (using weights $\{\theta_i\} = \{1/\mu_i\}$). With this Lyapunov function, we know that allocating power to maximize $\sum (U_i(t)/\mu_i) \tilde{\mu}_i(p_i, c_i)$ (where $c_i \in \{\text{ON}, \text{OFF}\}$) stabilizes the system. Clearly, the optimization does not need to place any power on queues in the OFF state, so the summation can be restricted to queues that are ON

$$\text{Maximize } \sum_{i|c_i=ON} U_i(t) \frac{\tilde{\mu}_i(p_i, c_i)}{\mu_i}, \text{ subject to } \sum p_i \leq K. \quad (29)$$

However, notice that the above maximization effectively chooses a rate vector $\vec{\mu}$ within the polytope capacity region specified in (12) and (13). The optimal solution for maximizing a linear function over a polytope will always be a vertex point. Fortunately, such a vertex point corresponds to the feasible allocation of K servers (with full power $p_i = 1$) to K queues. Considering (29), the optimal way to do this is to choose the K queues with the largest value of $U_i(t)$. \square

Using the same reasoning as in the proof above, it follows that the power-allocation policy (16) naturally chooses a vertex point when power curves are piecewise linear, such as the virtual curves described in Section II. It follows that optimization can be restricted to searches over the vertex points without loss of optimality.

C. Real-Time Implementation

The dynamic power-allocation policy of the previous section requires solving a nonlinear optimization problem every timeslot (16). However, because the rate curves $\mu_i(\cdot)$ are concave in the power parameter for every fixed channel state, the solution can be computed efficiently. Indeed, for positive weights $\{\theta_i\}$ and known unfinished work and channel-state vectors $\vec{U}(t)$ and $\vec{C}(t)$, the problem (16) becomes a standard concave maximization problem: Maximize $\sum \theta_i U_i(t) \mu_i(p_i, c_i(t))$, subject to the simplex constraint $\sum p_i \leq P_{\text{tot}}$. Using standard Lagrange multiplier techniques [22], it can be shown that a solution is optimal if and only if power is allocated according to the constraints so that scaled derivatives $\theta_i U_i(t) (d/dp_i) \mu(p_i, c_i(t))$ are equalized to some value γ^* for all queues i which receive nonzero power, while all queues which receive zero power have scaled derivatives less than γ . A fast bisection-type algorithm can be constructed to find such a solution, where a bracketing interval $[\gamma_1, \gamma_2]$ is found which contains γ^* and the interval size is decreased iteratively by testing the midpoint γ value to see if the corresponding powers sum to more or less than the power constraint P_{tot} . Such an algorithm yields power allocations whose proximity to the optimal solution converges geometrically with each iteration.

³Note that in Corollary 3.1, the constant B is proportional to $\sum \theta_i$, as expressed in (19). Hence, scaling all weights equally does not change the bound.

An important set of rate–power curves to consider are the standard curves for Shannon capacity

$$\mu_i(p_i, \alpha_i) = \log(1 + \alpha_i p_i)$$

where α_i represents the attenuation-to-noise level for downlink channel i during a particular timeslot. With these curves, the solution to (16) is found by the following computation:

$$\begin{aligned} \Lambda &= \text{Set of downlinks } i \in \{1, \dots, N\}, \text{ such that } U_i(t) > 0 \\ p_i &= \frac{\theta_i U_i(t) \left(P_{\text{tot}} + \sum_{j \in \Lambda} \frac{1}{\alpha_j} \right)}{\sum_{j \in \Lambda} \theta_j U_j(t)} - \frac{1}{\alpha_i}, \quad \text{if } i \in \Lambda \\ p_i &= 0, \quad \text{if } i \notin \Lambda. \end{aligned} \quad (30)$$

The above equations produce the optimal power allocations whenever the resulting p_i values are nonnegative. If any p_i values are negative, these are set to zero, the corresponding i indices are removed from the set Λ , and the calculation is repeated—a process ending in at most $N - 1$ iterations.

D. Robustness to Input Rate Changes

The dynamic power-allocation algorithm of Theorem 3 uses i.i.d. assumptions on packet arrivals and channel states to establish the negative Lyapunov drift condition. In [20], we show that the same dynamic policy stabilizes the system when inputs and channel states are Markov modulated (see simulations in Section VII). Here, we demonstrate that the policy is robust to arbitrary changes in the input rate $\vec{\lambda}$ as long as $\vec{\lambda}$ remains within the capacity region at each timestep. Specifically, suppose that the input rate to the downlink system is $\vec{\lambda}_1$ for a certain duration time, then changes to $\vec{\lambda}_2$ —perhaps due to changing user demands. This change will be reflected in the backlog that builds up in the queues of the system. Because the power-allocation algorithm bases decisions on the size of the queues, it reacts smoothly to such changes in the input statistics. Formally, this situation is modeled by defining an input distribution $f_k(\vec{a})$ on the arrival vector \vec{A}_k at each timestep k . The $f_k(\vec{a})$ distributions are arbitrary and unknown to the controller, although we assume they yield input rates $\vec{\lambda}_k = E[\vec{A}_k/T]$, all of which are within the capacity region. The dynamics of the system, thus, proceed according to a time-varying Markov chain $\vec{U}(0), \vec{U}(T), \dots, \vec{U}(kT)$. Although there is no notion of a steady-state distribution for a time-varying chain,⁴ we show that time averages are still well behaved.

Assume that all arrival distribution functions have uniformly bounded second moments, so that a value C exists where $\sum_i \theta_i E[a_i^2] < C$ for all distributions $f_k(\vec{a})$. Further suppose that there is some distance ε such that all instantaneous arrival rate vectors $\vec{\lambda}_k$ are at least a distance ε from the capacity region boundary, i.e., $\vec{\lambda}_k + \vec{\varepsilon}$ is in the capacity region for all k . For this system, we can again define a Lyapunov function $L(\vec{U}) = \sum_i \theta_i U_i^2$. Note that the one-step drift equation for

⁴Recall that the $f_k(\vec{a})$ distributions can vary arbitrarily. Thus, notions of steady state do not exist unless the dynamics are further described by some probabilistic model.

timestep k can be expressed as in (20), where the first and second moment of arrivals is taken over the k th arrival distribution $f_k(\vec{a})$. Because second moments are uniformly bounded for all k and because each rate vector $\vec{\lambda}_k$ is bounded by an ε distance away from the boundary of the capacity region, we have for each timestep k

$$\begin{aligned} E \left[L \left(\vec{U}(k+1)T \right) - L \left(\vec{U}(kT) \right) \mid \vec{U}(kT) \right] \\ \leq T^2(C + \beta) - 2T\varepsilon \sum_i \theta_i U_i(kT) \end{aligned}$$

where the constant β is defined in (18). In Theorem 4, given below, a simple telescoping series argument is used to show that the above inequality implies

$$\limsup_{M \rightarrow \infty} \frac{1}{M} \sum_{k=1}^M \left[\sum_i \theta_i E[U_i(kT)] \right] \leq \frac{(C + \beta)T}{2\varepsilon}.$$

The \limsup is needed for the above limit because the arrival distributions $f_k(\vec{a})$ may not be ergodic in k . Thus, the expression above indicates that, for arbitrary input rate changes, the power-allocation algorithm (16) stabilizes the system in that it maintains bounded time averages of unfinished work. This result follows as an immediate consequence of the following Theorem 4, which uses a telescoping series argument similar to a technique in [12].

Theorem 4 (Time-Varying Drift): If a time-varying Markov chain $\{\vec{U}(kT)\}$ has a Lyapunov function $L(\vec{U})$ that satisfies

$$\begin{aligned} E \left[L \left(\vec{U}((k+1)T) \right) - L \left(\vec{U}(kT) \right) \mid \vec{U}(kT) \right] \\ \leq B - \sum_i \alpha_i U_i(kT) \end{aligned} \quad (31)$$

for positive constants $B, \{\alpha_i\}$, then

$$\limsup_{M \rightarrow \infty} \frac{1}{M} \sum_{k=0}^{M-1} \left[\sum_i \alpha_i E[U_i(kT)] \right] \leq B. \quad (32)$$

Furthermore, if the chain is time invariant and condition 3 of Theorem 2 holds, then a stationary distribution for \vec{U} exists with the property that $\sum \alpha_i \bar{U}_i \leq B$.

Proof: Taking expectations of (31) over the distribution of $\vec{U}(kT)$ and summing over k from 0 to $M - 1$ yields

$$\begin{aligned} E \left[L \left(\vec{U}(MT) \right) \right] - E \left[L \left(\vec{U}(0) \right) \right] \\ \leq BM - \sum_{k=0}^{M-1} \left[\sum_i \alpha_i E[U_i(kT)] \right]. \end{aligned}$$

Hence

$$\frac{1}{M} \sum_{k=0}^{M-1} \left[\sum_i \alpha_i E[U_i(kT)] \right] \leq B + \frac{E \left[L \left(\vec{U}(0) \right) \right]}{M}.$$

Taking the \limsup of the above inequality yields (32). In the case when the chain is time invariant, inequality (32) implies that the first two conditions of the Lyapunov drift theorem (Theorem 2) hold. If the third condition also holds, a steady-




Fig. 5. Multiuser multisatellite system with joint routing and power control. User X_i can route to queues within set Q_i . Satellite s allocates power subject to $\sum_{j \in \text{Sat}(s)} p_j(t) \leq P_{tot}^{(s)}$.

state distribution exists. For this steady state, by (32), we have $\sum \alpha_i \bar{U}_i \leq B$. \square

V. JOINT ROUTING AND POWER ALLOCATION

We consider now a collection of S multibeam satellites and develop a method for jointly routing packets and allocating power over the downlinks. Each satellite has multiple output queues (corresponding to multiple downlink channels) and operates according to individual power constraints $\{P_{tot}^{(s)}\}_{s=1}^S$ (Fig. 5). Every timeslot, packets enter the system from M input streams according to input processes $(X_i)_{i=1}^M$ with arrival rates $(\lambda_1, \dots, \lambda_M)$. Each input stream i can route incoming packets to a subset of the output queues, where the subsets may overlap with each other and may contain queues from different satellites. The problem is to jointly route packets and allocate power to each of the downlinks in order to stabilize the system and ensure maximum throughput.

Such a scenario arises, for example, when several satellites have a connection to the same ground unit, and, hence, packets destined for this unit have several routing options. Alternatively, the routing options may represent a cluster of ground locations connected together by a reliable ground network. In this case, packets arrive to the cluster from the downlinks and are routed to their final destinations using the wire lines on the ground. We note that the formulation of this joint routing and power-allocation problem also applies to wireless systems, where base stations communicate with users over a wireless network.

As before, we consider slotted time and assume the M input streams produce an arrival vector \vec{A} i.i.d. every timeslot, where $E[\vec{A}/T] = (\lambda_1, \dots, \lambda_M)$. Let J represent the total number of output queues (summed over all satellites), and let each output queue be indexed with a single integer $j \in \{1, \dots, J\}$. For each satellite s , let $\text{Sat}(s)$ represent the set of output queues which it contains (hence, $\text{Sat}(s) \subseteq \{1, \dots, J\}$ for all $s \in \{1, \dots, S\}$). Likewise, for each input stream i , let Q_i represent the set of all output queues that input i can route packets to (where $Q_i \subseteq \{1, \dots, J\}$). Note that the Q_i subsets are arbitrary and need not be disjoint.

Channel states vary according to an i.i.d. state vector $\vec{C}(t) = (c_j(t))_{j=1}^J$. Let $\vec{U}(t) = (U_j(t))_{j=1}^J$ represent the vector state of unfinished work in all queues at time t . Every

timeslot, routing decisions are made and a power vector $\vec{P}(t)$ is allocated according to the system constraints. In general, the full queue- and channel-state vectors $\vec{U}(t)$ and $\vec{C}(t)$ are important in both the routing and power-allocation decisions. For example, more power should be allocated to queues which are expected to grow large—which is dependent on the state of unfinished work in other satellites as well as on future routing decisions. Likewise, a router should place packets in faster queues (especially if these rates are likely to be high for one or more timeslots) and should avoid queues likely to be congested because of high contention with other input sessions.

However, here we show that the routing and power-allocation decisions can be *decoupled* into two policies: a routing policy which considers only $\vec{U}(t)$ and a power-allocation policy which considers both $\vec{U}(t)$ and $\vec{C}(t)$. Furthermore, a router for stream i needs only to consider the entries of the unfinished work vector $\vec{U}(t)$ within the set Q_i of queues to which it can route. Likewise, the power-allocation decisions use information local to each satellite: Power is allocated in satellite s based only on the unfinished work and channel-state information for queues in $\text{Sat}(s)$. The resulting strategy stabilizes the system whenever the system is stabilizable.

Theorem 5 (Joint Routing and Power Allocation): The capacity region Ω for the multisatellite system with joint routing and power allocation is the set of all arrival vectors $\vec{\lambda} = (\lambda_1, \dots, \lambda_M)$ such that there exist splitting rates r_{ij} and power levels $p_j^{\vec{C}}$ such that

- 1) $\sum_{j \in Q_i} r_{ij} = \lambda_i$, for all $i \in \{1, \dots, M\}$;
- 2) $\sum_{j \in \text{Sat}(s)} p_j^{\vec{C}} \leq P_{tot}^{(s)}$, for all s and all channel states \vec{C} ;
- 3) $\sum_i r_{ij} \leq \sum_{\vec{C}} \pi_{\vec{C}} \mu_j(p_j^{\vec{C}}, c_j)$.

Intuitively, the above theorem says that the system is stabilizable if the input rates can be split amongst the various queues (in accordance with the routing restrictions) so that the aggregate input rates allow each satellite to be stabilized individually.

Proof that $\vec{\lambda} \in \Omega$ is necessary for stability: Suppose a stabilizing algorithm exists for some set of routing decisions and power controls $\vec{P}(t)$. Define $X_{ij}(t)$ to be the total amount of data the algorithm routes from input i to queue j during the time interval $[0, t]$. For simplicity, we assume the routing process is ergodic so that $\lim_{t \rightarrow \infty} X_{ij}(t)/t$ is well defined for all i and j . (The general nonergodic case can be handled similarly, according to our treatment in Theorem 1). Let $\{r_{ij}\}$ represent these limiting values. The i th input stream $X_i(t)$ can be written $X_i(t) = \sum_{j \in Q_i} X_{ij}(t)$. Dividing both sides by t and taking limits, it follows that $\sum_{j \in Q_i} r_{ij} = \lambda_i$ for all i , and, hence, condition 1) holds. Note that the aggregate data rate entering any queue $j \in \{1, \dots, J\}$ is $\sum_i r_{ij}$. Because the system is stable, the stability conditions of Theorem 1 must be satisfied for each satellite, and, hence, the remaining conditions 2) and 3) must also hold. \square

We prove the sufficiency part of Theorem 5 by considering the following decoupled power-allocation and routing policies. The policies use only local information about queue and channel states, and do not require knowledge of the input rates $\vec{\lambda}$ or the channel-state probabilities $\pi_{\vec{C}}$. We assume that enough is known about the channel to identify and remove from the set of routing options any queues which produce zero output rate

for all channel states and power allocations. Hence, in the algorithms and analysis below, we assume that all queues j have some nonzero probability of being in a functional channel state.

Joint Routing and Power-Allocation Algorithm:

Power allocation: At each timestep, each satellite s allocates power as before, using the $\vec{U}(t)$ and $\vec{C}(t)$ vectors to maximize

$$\sum_{j \in \text{Sat}(s)} U_j \mu_j(p_j(t), c_j(t)), \text{ subject to } \sum_{j \in \text{Sat}(s)} p_j(t) \leq P_{\text{tot}}^{(s)}.$$

Routing: All packets from stream i are routed to the queue $j \in Q_i$ with the smallest amount of unfinished work.

Proof of stability whenever $\vec{\lambda}$ is strictly interior to Ω : Suppose the $\vec{\lambda}$ vector is strictly interior to Ω so that conditions 1)–3) of Theorem 5 are satisfied even with an additional input stream of rate ε applied to each queue $j \in \{1, \dots, J\}$, i.e., there exist r_{ij} and $p_j^{\vec{C}}$ values such that conditions 1) and 2) hold, and such that

$$\sum_i r_{ij} + \varepsilon \leq \sum_{\vec{C}} \pi_{\vec{C}} \mu(p_j^{\vec{C}}, c_j), \text{ for all } j. \quad (33)$$

Again, define the quadratic Lyapunov function $L(\vec{U}) = \sum U_j^2$. Let $A_i(t)$ represent the total amount of bits from packets arriving from stream i in timeslot $[t, t+T]$, and let $(a_{i1}(t), \dots, a_{iJ}(t))$ represent the bit length of packets from stream i routed to queues $j \in \{1, \dots, J\}$ (where $A_i(t) = \sum_j a_{ij}(t)$, and $E[A_i] = \lambda_i T$). Let μ_j represent the transmission rate $\mu_j(p_j(t), c_j(t))$ of queue j during timeslot $[t, t+T]$ under the specified power-allocation policy. Likewise, let $\sum_i a_{ij}$ represent the total arrivals to queue j in a timeslot. As in the stability proof for the dynamic power-allocation policy of Theorem 3, we have for all queues j [from (17)]

$$U_j^2(t+T) \leq U_j^2 - 2TU_j \left(\mu_j - \frac{1}{T} \sum_i a_{ij} \right) + \mu_j^2 T^2 + \left(\sum_i a_{ij} \right)^2 \quad (34)$$

Define the constant

$$D = E \left[\left(\sum_i \frac{A_i}{T} \right)^2 \right] + \left[\sum_{\vec{C}} \pi_{\vec{C}} \sum_j \mu_j^2(p_j^{\vec{C}}, c_j) \right] \quad (35)$$

where the $\bar{p}_j^{\vec{C}}$ values are chosen to maximize the second term in (35) subject to the power constraints $\sum_{j \in \text{Sat}(s)} \bar{p}_j^{\vec{C}} \leq P_{\text{tot}}^{(s)}$ for all satellites s and channel states \vec{C} . Summing (34) over all $j \in \{1, \dots, J\}$ and taking conditional expectations, it can be shown that

$$E \left[L(\vec{U}(t+T)) - L(\vec{U}(t)) \mid \vec{U}(t) \right] \leq T^2 D - 2T \cdot \sum_j U_j \left(E[\mu_j | \vec{U}(t)] - \sum_i E \left[\frac{a_{ij}}{T} \mid \vec{U}(t) \right] \right). \quad (36)$$

The $E[\mu_j | \vec{U}(t)]$ and $E[a_{ij}/T | \vec{U}(t)]$ values in the above inequality are influenced by the power control and routing algorithm, respectively, and will determine the performance of the system. To examine the impact of routing, we switch the sum above to express the routing term as

$$2T \sum_i \sum_{j \in Q_i} U_j E \left[\frac{a_{ij}}{T} \mid \vec{U}(t) \right].$$

Notice that the given routing strategy of placing all bits from stream i in the queue $j \in Q_i$ with the smallest value of unfinished work minimizes the above term over all possible routing strategies, including the strategy of routing according to flow rates r_{ij} of condition (1) in Theorem 5, and, hence

$$\sum_i \sum_{j \in Q_i} U_j E \left[\frac{a_{ij}}{T} \mid \vec{U}(t) \right] \leq \sum_i \sum_{j \in Q_i} U_j r_{ij}. \quad (37)$$

To examine the power-allocation term in (36), we rewrite the single summation as a double summation over all satellites: $s \in \{1, \dots, S\}$

$$2T \sum_j U_j E[\mu_j | \vec{U}(t)] = \sum_s \sum_{j \in \text{Sat}(s)} U_j E[\mu_j | \vec{U}(t)]. \quad (38)$$

By an argument similar to (21)–(25), it can be shown that the given power-allocation policy maximizes (38) over all allocation policies, including the policy of allocating a power vector $\vec{P}^{\vec{C}}(t) = (p_1^{\vec{C}}, \dots, p_J^{\vec{C}})$ whenever the channel is in state \vec{C} . Hence

$$\sum_j U_j E[\mu_j | \vec{U}(t)] \geq \sum_j U_j \sum_{\vec{C}} \pi_{\vec{C}} \mu(p_j^{\vec{C}}, c_j). \quad (39)$$

Using (37), (39), and (33) in (36), we find

$$E \left[L(\vec{U}(t+T)) - L(\vec{U}(t)) \mid \vec{U}(t) \right] \leq T^2 D - 2T\varepsilon \sum_j U_j.$$

Defining any $\alpha > 0$ and choosing the compact set Λ to be

$$\Lambda = \left\{ \vec{U} \in \mathbb{R}^N \mid U_j \geq 0, \sum U_j \leq \left(\frac{T^2 D + \alpha}{2T\varepsilon} \right) \right\}$$

ensures the negative drift condition of Theorem 2 whenever $\vec{U} \notin \Lambda$. \square

Corollary 5.1 (Performance Bound of the Joint Routing and Power-Allocation Strategy): Steady-state values of unfinished work are bounded and satisfy $\sum \bar{U}_j \leq TD/(2\varepsilon)$. \square

An important special case of the above theorem is when inputs can route to the full set of available queues, i.e., $Q_i = \{1, \dots, J\}$ for all inputs i . The goal is to simply transmit all the data to the ground as soon as possible. Such a situation arises when the ground units are connected together via a reliable ground network, and the wireless paths from satellite to ground form the rate bottleneck (see Fig. 6). In this case, it is shown in [18] that the capacity region of the Joint Routing and Power-Allocation Theorem 5 simplifies to the simplex set of all input rates $\vec{\lambda} = (\lambda_1, \dots, \lambda_M)$ such that

$$\lambda_1 + \dots + \lambda_M \leq \bar{\mu}_{\text{out}} \quad (40)$$




Fig. 6. Joint routing/power-allocation problem where the goal is to transmit the data to any node of the reliable ground network ($Q_i = \{1, \dots, J\}$ for all input streams i).




Fig. 7. Capacity region for a two-queue system with routing and power control. The region Ω corresponds to the routing constraints shown in the figure, and is dominated by the simplex region for unconstrained routing.

where

$$\bar{\mu}_{\text{out}} = \sum_{\vec{c}} \pi_{\vec{c}} \sum_{s=1}^S \max_{\left\{ \sum_{j \in \text{Sat}(s)} p_j \leq P_{\text{tot}}^{(s)} \right\}} \left[\sum_{j \in \text{Sat}(s)} \mu_j(p_j, c_j) \right]$$

that is, $\bar{\mu}_{\text{out}}$ is the average output rate of the system when power is allocated to maximize the instantaneous processing rate at every instant of time.

In Fig. 7, we illustrate the capacity region for a two-queue system with and without routing constraints. As expected, exploiting the full set of routing options considerably expands the capacity region of the system. Indeed, the simplex region (40) always contains the capacity region specified in Theorem 5 for joint routing and power allocation. This capacity gain is achieved by utilizing the extra resources offered by the ground network.

We note that this joint routing and power-allocation problem has been formulated for the case when data already contained within a single satellite or within a constellation of satellites is to be routed through a choice of downlinks. Hence, it is reasonable to assume the unfinished work values $U_i(t)$ are known to the controllers when making routing decisions. However, it can be shown that one can apply the same strategy when only *estimates* of the true unfinished work values are known, and the system is still stable for all arrival rates within the stability region.

VI. CONNECTIVITY CONSTRAINTS

It has been assumed throughout that all transmit channels can be activated simultaneously, subject only to the total power constraint $\sum_{j \in \text{Sat}(s)} p_j(t) \leq P_{\text{tot}}^{(s)}$ for all time t . Hence, it is

implicitly assumed that there is no interchannel interference. Such an assumption is valid when there is sufficient bandwidth to ensure potentially interfering channels can transmit using different frequency bands. However, in bandwidth-limited scenarios, power-allocation vectors $\vec{P}(t)$ may be additionally restricted to channel *activation sets*: finite sets $\mathcal{P}_1, \dots, \mathcal{P}_R$, where each set \mathcal{P}_r is a convex set of points (p_1, \dots, p_N) representing power vectors which, when allocated, ensure that interchannel interference is at an acceptable level. This use of activation sets is similar to the treatment in [3], where activation link sets for scheduling ON/OFF links in a wireless network are considered. Here, the definition has been extended from sets of links to sets of power vectors to treat power control.

As an example of an activation set, consider the single-satellite system of Fig. 1 with N output queues, and suppose that downlink channels 1, 2, and 3 can be activated simultaneously if all other transmitters are silent. Such an activation set can be represented by

$$\mathcal{P}_r = \left\{ (p_1, p_2, p_3, 0, \dots, 0) \in \mathbb{R}^N \mid p_j \geq 0, \sum_{j=1}^3 p_j \leq P_{\text{tot}} \right\}.$$

Another type of system constraint is when power allocation is further restricted so that no more than K transmitters are active at any given time. Such a constraint corresponds to $\binom{N}{K}$ convex activation sets. Multisatellite systems can also be treated using this activation set model. Indeed, the N output queues of Fig. 1 may be physically located in several different satellites. In the following, we assume that each activation set incorporates the power constraints $\sum_{j \in \text{Sat}(s)} p_j \leq P_{\text{tot}}^{(s)}$.

Consider the downlink system of Fig. 1. Packets arrive according to a random arrival vector (i.i.d. on each timeslot) with rates $(\lambda_1, \dots, \lambda_N)$, and channel states $\vec{C}(t)$ vary i.i.d. every timeslot with steady-state probabilities $\pi_{\vec{c}}$. Each timeslot, a power-allocation vector $\vec{P}(t)$ is chosen such that it lies within one of the activation sets $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_R\}$.

Theorem 6 (Power Allocation With Connectivity Constraints): For the multiqueue system of Fig. 1 with power constraints $\vec{P} \in \mathcal{P}$:

- 1) The capacity region of the system is the set Ω of all arrival rate vectors $\vec{\lambda}$ such that

$$\vec{\lambda} \in \Omega \triangleq \sum_{\vec{c}} \pi_{\vec{c}} \text{Convex Hull} \left(\left\{ \bar{\mu}(\vec{P}, \vec{c}) \mid (\vec{P} \in \mathcal{P}_r) \right\}_{r=1}^R \right) \quad (41)$$

where addition and scalar multiplication of sets has been used above.⁵

- 2) The policy of allocating a power vector $\vec{P} = (p_1, \dots, p_N)$ at each timestep to maximize the quantity $\sum U_j(t) \mu_j(p_j, c_j(t))$ (subject to $\vec{P} \in \mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_R\}$) stabilizes the system whenever the $\vec{\lambda}$ vector is in the interior of the capacity region.

⁵For sets A, B and scalars α, β , the set $\alpha A + \beta B$ is defined as $\{\gamma \mid \gamma = \alpha a + \beta b \text{ for some } a \in A, b \in B\}$.

We note that the allocation policy specified in part 2) of the theorem involves the nonconvex constraint $\vec{P} \in \mathcal{P}$. Optimizing the given metric over individual activation sets \mathcal{P}_r is a convex optimization problem, although a complete implementation of the given policy is nontrivial if the number of activation sets is large.

However, the proof of parts 1) and 2) are simple extensions of the analysis presented in Sections III and IV.

Proof of 1): To establish that $\vec{\lambda} \in \Omega$ is a necessary condition for stability, suppose the system is stable using some power-allocation function $\vec{P}(t)$ which satisfies $\vec{P}(t) \in \mathcal{P}$ for all time. Thus, we know that $\lambda_i \leq \underline{\mu}_i$ for all i (Lemma 1 for queue stability), and the proof proceeds as the proof of the downlink capacity theorem (Theorem 1), where for any fixed $\varepsilon > 0$ we can find a large time \tilde{t} such that the following entrywise vector inequality is satisfied (where $\vec{\varepsilon} = (\varepsilon, \dots, \varepsilon)$)

$$\vec{\lambda} \leq \underline{\mu} \leq \frac{1}{\tilde{t}} \int_0^{\tilde{t}} \underline{\mu}(\vec{P}(\tau), \vec{C}(\tau)) d\tau + \vec{\varepsilon}.$$

The main difference from Theorem 1 is that the above integral is broken into a double summation over intervals when the channel is in state \vec{C} and when the power vector is in set \mathcal{P}_r . Let $T_{\vec{C}}(\tilde{t})$ represent the intervals of time during $[0, \tilde{t}]$ when the channel is in state \vec{C} , and let $T_{\vec{C}, \mathcal{P}_r}(\tilde{t})$ represent the subintervals of $T_{\vec{C}}(\tilde{t})$ when the power function $\vec{P}(t)$ is in activation set \mathcal{P}_r . Similar to our treatment in (7)–(9), the concavity of the rate–power curve allows the integration to be pushed inside the $\underline{\mu}(\cdot)$ function, and we have

$$\vec{\lambda} \leq \sum_{\vec{C}} \pi_{\vec{C}} \left[\sum_{\mathcal{P}_r} \frac{\|T_{\vec{C}, \mathcal{P}_r}(\tilde{t})\|}{\|T_{\vec{C}}(\tilde{t})\|} \underline{\mu}(\vec{P}_{\vec{C}, \mathcal{P}_r}, \vec{C}) \right] + O(\varepsilon)$$

where

$$\vec{P}_{\vec{C}, \mathcal{P}_r} \triangleq \frac{1}{\|T_{\vec{C}, \mathcal{P}_r}(\tilde{t})\|} \int_{[\tau \in T_{\vec{C}, \mathcal{P}_r}(\tilde{t})]} \vec{P}(\tau) d\tau.$$

Note that any point $\vec{\gamma}$ in the convex hull of a collection of convex sets can be written as a linear combination of points $\vec{\gamma}_1, \dots, \vec{\gamma}_R$ in the sets: $\vec{\gamma} = \alpha_1 \vec{\gamma}_1 + \dots + \alpha_R \vec{\gamma}_R$ where $\alpha_r \geq 0$ and $\sum \alpha_r = 1$. Letting $\alpha_r = \|T_{\vec{C}, \mathcal{P}_r}(\tilde{t})\| / \|T_{\vec{C}}(\tilde{t})\|$, we see the inequality above indicates that $\vec{\lambda}$ is arbitrarily close to a point in the closed region Ω defined in (41) and, hence, $\vec{\lambda} \in \Omega$.

The sufficiency condition is implied by part 2). \square

Proof of 2): Define the Lyapunov function $L\vec{U} = \sum U_i^2$. The proof of Theorem 3 can literally be repeated up to (27)

$$\begin{aligned} E \left[L(\vec{U}(t+T)) - L(\vec{U}(t)) \mid \vec{U}(t) \right] \\ \leq T^2 B - 2T \sum_j U_j(t) \left(E[\mu_j \mid \vec{U}(t)] - \lambda_j \right). \end{aligned}$$

From this point, negative drift of the Lyapunov function can be established by again noting that the value of $E[\mu_j \mid \vec{U}(t)]$ maxi-

Log-normal distribution of α_i for each of the three channel conditions:

	mean	variance
Good	15 db	.264 db-squared
Medium	10 db	.868 db-squared
Bad	0 db	.145 db-squared




Fig. 8. Three-state Markov Chain representing Good, Medium, and Bad conditions for a single downlink from satellite to ground. In each state, an attenuation level α_i is chosen according to a log-normal distribution with means and variances as shown.

mizes $\sum U_j \gamma_j$ over all vectors $\vec{\gamma}$ within the region Ω specified in (41). To see this, note that any $\vec{\gamma}$ in Ω can be written as

$$\vec{\gamma} = \sum_{\vec{C}} \pi_{\vec{C}} \sum_{\mathcal{P}_r} \alpha_{\vec{C}, \mathcal{P}_r} \underline{\mu}(\vec{P}_{\vec{C}, \mathcal{P}_r}, \vec{C})$$

for some vectors $\vec{P}_{\vec{C}, \mathcal{P}_r} \in \mathcal{P}_r$ and some scalar values $\alpha_{\vec{C}, \mathcal{P}_r} \geq 0$ such that $\sum_{\mathcal{P}_r} \alpha_{\vec{C}, \mathcal{P}_r} = 1$ for all channel states \vec{C} . The result follows by an argument similar to (21)–(25). \square

VII. NUMERICAL RESULTS

Here, we present numerical and simulation results illustrating the capacity and delay performance provided by the dynamic power-allocation policy of Section IV (16) for a simple satellite downlink consisting of two channels and two queues. We assume the corresponding input streams consist of unit length packets arriving as Poisson processes with rates (λ_1, λ_2) . However, rather than simulating i.i.d. channel-state vectors every timeslot, we consider a Markov modulated channel state that is typical of a satellite downlink [24] and demonstrate the ability of the dynamic power-allocation policy (16) to perform well under general time-varying channel conditions.

Specifically, consider a Markov chain with three states corresponding to Good, Medium, and Bad channel conditions, with transition probabilities shown in Fig. 8. Such a three-state system has been considered in [24] and extends the well-known two-state Gilbert–Elliott model [28], [29] for satellite and wireless channels. In each state, we assume signal attenuation is log-normally distributed with a given mean and variance. Such a distribution is consistent with the Karasawa model [27] based on short-term fading measurements in the Ka band.

Total transmit power at the satellite is assumed to be 100 W. Factoring together the antenna gains, signal attenuation, and receiver noise, the average signal-to-noise-ratio when full power is allocated to a single channel is assumed to be 15, 10, and 0 db, for Good, Medium, and Bad conditions, respectively. The corresponding variances are .264, .868, and .145 db², respectively. These values are based on measurement data for Ka-band satellite channels given in [24]. We consider the Shannon capacity curves for data rate as a function of a normalized signal power

$$\mu_i(p_i, c_i) = \log(1 + \alpha_i p_i)$$




Fig. 9. Stability regions for the three power-allocation algorithms in the two-queue downlink system. The isolated points represent (a) rate points (λ_1, λ_2) used in the simulations, and (b) the corresponding average unfinished work $E[U_1 + U_2]$ obtained from simulations.

where $\sum p_i \leq 1$ and α_i represents the fading coefficients, chosen according to the specified log-normal distributions with mean and variance determined by the channel state c_i . For the simulation, we discretize the log-normal distribution with eleven quantization levels. The two channels from satellite to ground are assumed to vary independently, each according to the described Markov modulated process.

In Fig. 9(a), we plot the downlink capacity region given by Theorem 1 (4). Notice the nonlinear “bulge” curvature, representing capacity gains due to dynamic power allocation. This full region is achievable using the dynamic power-allocation algorithm of Theorem 3 (16), (30). We compare the capacity region offered by this algorithm to the corresponding stability regions when power is allocated according to the following alternative strategies.

- 1) ON/OFF power allocation: Only one transmitter can be activated at any time.
- 2) Static power allocation: Constant power $P_{\text{tot}}/2$ is allocated to each channel for all time.

The ON/OFF strategy allocates full power to the channel with the largest rate-backlog index $U_i(t)\mu_i(P_{\text{tot}}, c_i(t))$, which, by Theorem 6, achieves full capacity among all policies restricted to using a single transmitter. Notice that the stability region is slightly nonlinear, because of the capacity boost due to the diversity offered by the independently time-varying channels. The stability region for the static power-allocation algorithm has a rectangular shape, as shown in Fig. 9(a). The capacity for this static algorithm is expanded beyond the stability region for the single-transmitter algorithm when the input rates λ_1 and λ_2 are roughly within a factor of two of each other, although the single-transmitter algorithm is better for highly asymmetric data rates. Both policies are stable on a significantly reduced subset of the capacity region offered by the dynamic power-allocation policy. Note that even in the completely symmetric case $\lambda_1 = \lambda_2$, the stability point of the static power-allocation policy is slightly below the stability point of the dynamic power-allocation policy, because the static policy cannot take advantage of the time-varying channel conditions.

In addition, we simulate system dynamics for two million iterations using the three power-allocation policies and a variety of data rates which linearly approach a boundary rate point $(\lambda_1, \lambda_2) = (2.05, 3.79)$ of the capacity region. The rates tested

are shown in Fig. 9(a). In Fig. 9(b), we plot the empirical average occupancy $E[U_1 + U_2]$ for the two-queue system when the multibeam dynamic power-allocation algorithm is used, where power is allocated according to (30) (with weights $\theta_i = 1$, for all i). The plot illustrates that the dynamic power-allocation policy achieves stability throughout the entire capacity region, with an average delay growing asymptotically as the input data rates approach the boundary point $(2.05, 3.79)$.

We compare the dynamic power-allocation algorithm to the two other strategies, whose simulated performance is also shown in Fig. 9(b). From the figure, it is clear that the average occupancy (and, hence, average delay) of the multibeam dynamic power-allocation algorithm is significantly lower than the corresponding averages for the other algorithms at all data rates (note that the asymptotes for instability occur earlier for the other two algorithms). For the rate regime tested, the stability region for the constant power-allocation algorithm is slightly larger than the single-transmitter dynamic algorithm, and, hence, the corresponding average occupancies are lower. However, the static policy cannot adjust to asymmetries in data rate and, thus, the single-transmitter algorithm will perform better in the regime where one input rate is much larger than the other [see the capacity plot in Fig. 9(a)]. The figures illustrate that to enable high data rates and low delay in a satellite downlink, it is essential to dynamically allocate power to the multiple beams.

VIII. CONCLUSION

We have treated data transmission over multiple time-varying channels in a satellite downlink using power control. Processing rates for each channel i were assumed to be determined by concave rate-power curves $\mu_i(p_i, c_i)$, and the capacity region of all stabilizable arrival rate vectors $\vec{\lambda}$ was established. This capacity region is valid for general Markovian input streams, and inputs with arrival rates $\vec{\lambda}$ in the interior of the capacity region can be stabilized with a power-allocation policy which only considers the current channel state $\vec{C}(t)$. In the case when arrival rates and channel probabilities $\vec{\lambda}$ and $\pi_{\vec{C}}$ are unknown, but packet arrivals and channel-state transitions are i.i.d. every timeslot, a stabilizing policy which considers both current channel state and current queue backlog was developed. Intuitively, the policy favors queues with large backlogs and better channels by allocating power to maximize $\sum U_i \mu_i$ at every timeslot. The policy reacts smoothly to channel-state changes and arbitrary variations in the input rates. A real-time implementation of the algorithm was described, and an analytical bound on average bit delay was established. This power-control formulation was shown to contain the special case of a server-allocation problem, and analysis verified stability and provided a performance bound for the *Serve-the-K-Largest-Connected-Queue* policy.

A joint routing and power-allocation scenario was also considered for a system with multiple users and multiple satellites, and a throughput maximizing algorithm and a corresponding performance bound was developed. The structure of this algorithm allows for decoupled routing and power-allocation decisions to be made by each user and each satellite based on local channel state and queue backlog information.

In the case of interchannel interference, modified power-allocation policies were developed when power vectors are constrained to a finite collection of activation sets. The policies offer 100% throughput, although they are difficult to implement if the number of activation sets is large.

Stability properties of these systems were established by demonstrating negative drift of a Lyapunov function defined over the current state of unfinished work in the queues. Robustness to arbitrary input rate changes was demonstrated by establishing an upper bound on time-average queue occupancy in the case when the arrival rate vector $\vec{\lambda}_k$ is inside the capacity region for all timesteps k . We show in [20] that the given control strategies provide similar performance guarantees for general Markovian arrival and channel processes whenever the steady-state arrival rate vector $\vec{\lambda}$ is in the capacity region. Thus, they offer desirable performance under a variety of input processes and time-varying channel conditions.

Our focus was power control for a satellite downlink, although the results extend to other wireless communication scenarios where power allocation and energy efficiency is a major issue. The use of dynamic power allocation can considerably extend the throughput and performance properties of such systems.

APPENDIX

Lemma 1b: If an input stream $X(t)$ to a single-queue system is rate-ergodic of input rate λ , a necessary condition for queue stability is $\lambda \leq \underline{\mu}$.

Proof: Suppose $\lambda > \underline{\mu}$ and choose $\varepsilon > 0$ such that $\lambda - \underline{\mu} - 2\varepsilon > 0$. The limits in (2) ensure that, with probability 1, we can find a set of times $\{t_i\} (i \in \{1, 2, \dots\})$ where $t_i \rightarrow \infty$ with increasing i and such that, for all t_i

$$\frac{X(t_i)}{t_i} \geq \lambda - \varepsilon, \quad \frac{1}{t_i} \int_0^{t_i} \mu(\tau) d\tau \leq \underline{\mu} + \varepsilon. \quad (42)$$

However, it is clear that

$$U(t_i) \geq X(t_i) - \int_0^{t_i} \mu(\tau) d\tau. \quad (43)$$

From (42) and (43), it follows that $U(t_i) \geq (\lambda - \underline{\mu} - 2\varepsilon)t_i$, for all t_i . Define $\alpha = \lambda - \underline{\mu} - 2\varepsilon$, and let T_i represent the extra time it takes the unfinished work in the queue to empty below a threshold value M , starting at value $U(t_i)$ at time t_i . Clearly, $T_i \geq (\alpha t_i - M)/\mu_{\max}$ and, hence, at any time $t_i + T_i$ the empirical fraction of time that the unfinished work in the queue exceeded the value M is greater than or equal to $T_i/(t_i + T_i)$, which is greater than or equal to $(\alpha t_i - M)/(\mu_{\max} t_i + \alpha t_i - M)$. Taking limits as $t_i \rightarrow \infty$ reveals that $g(M) \geq \alpha/(\alpha + \mu_{\max})$ for all M and, hence, the system is unstable. \square

REFERENCES

[1] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power and server allocation in a multibeam satellite with time-varying channels," in *Proc. IEEE INFOCOM*, vol. 3, 2002, pp. 1451–1460.

[2] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inform. Theory*, vol. 39, pp. 466–478, Mar. 1993.

[3] —, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Automat. Contr.*, vol. 37, pp. 1936–1948, Dec. 1992.

[4] L. Tassiulas, "Scheduling and performance limits of networks with constantly changing topology," *IEEE Trans. Inform. Theory*, vol. 43, pp. 1067–1073, May 1997.

[5] B. Hajek, "Optimal control of two interacting service stations," *IEEE Trans. Automat. Contr.*, vol. AC-29, pp. 491–499, June 1984.

[6] I. Viniotis and A. Ephremides, "Extension of the optimality of the threshold policy in heterogeneous multiserver queueing systems," *IEEE Trans. Automat. Contr.*, vol. AC-33, pp. 104–109, Jan. 1988.

[7] M. Carr and B. Hajek, "Scheduling with asynchronous service opportunities with applications to multiple satellite systems," *IEEE Trans. Automat. Contr.*, vol. 38, pp. 1820–1833, Dec. 1993.

[8] G. Bongiovanni, D. T. Tang, and C. K. Wong, "A general multibeam satellite switching algorithm," *IEEE Trans. Commun.*, vol. COM-29, pp. 1025–1036, July 1981.

[9] A. Mekikittikul and N. McKeown, "A practical scheduling algorithm to achieve 100% throughput in input-queued switches," in *Proc. IEEE INFOCOM*, 1998, pp. 792–799.

[10] N. McKeown, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," in *Proc. IEEE INFOCOM*, Mar. 1996, pp. 296–302.

[11] M. Karol, M. Hluchyj, and S. Morgan, "Input versus output queueing in a space division switch," *IEEE Trans. Commun.*, vol. COM-35, pp. 1347–1356, Dec. 1987.

[12] E. Leonardi, M. Mellia, F. Neri, and M. A. Marson, "Bounds on average delays and queue size averages and variances in input-queued cell-based switches," in *Proc. IEEE INFOCOM*, vol. 2, 2001, pp. 1095–1103.

[13] P. R. Kumar and S. P. Meyn, "Duality and linear programs for stability and performance analysis of queueing networks and scheduling policies," *IEEE Trans. Automat. Contr.*, vol. 41, pp. 4–17, Jan. 1996.

[14] E. Uysal-Biyikoglu, B. Prabhakar, and A. El Gamal, "Energy-efficient packet transmission over a wireless link," *IEEE/ACM Trans. Networking*, vol. 10, pp. 487–499, Aug. 2002.

[15] S. Toumpis and A. Goldsmith, "Some capacity results for Ad Hoc networks," in *Proc. 38th Annu. Allerton Conf.*, 2000, pp. 74–75.

[16] I. C. Paschalidis, "Large deviations in high speed communication networks," Ph.D. dissertation, Lab. Information and Decision Syst., Massachusetts Inst. Technol., Cambridge, 1996.

[17] S. Asmussen, *Applied Probability and Queues*. New York: Wiley, 1987.

[18] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic routing to parallel time-varying queues with applications to satellite and wireless networks," in *Proc. Conf. Information Sciences and Systems*, Mar. 2002, pp. 36–39.

[19] M. J. Neely and E. Modiano, "Convexity and optimal load distributions in work conserving $*/ * / 1$ queues," in *Proc. IEEE INFOCOM*, vol. 2, 2001, pp. 1055–1064.

[20] M. J. Neely, "Dynamic power allocation and routing in satellite and wireless networks with time-varying channels," Ph.D. dissertation, Lab. Information and Decision Syst., Massachusetts Inst. Technol., Cambridge, 2003.

[21] R. G. Gallager, *Discrete Stochastic Processes*. Boston, MA: Kluwer, 1996.

[22] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena, 1995.

[23] J. P. Choi and V. W. S. Chan, "Predicting and adapting satellite channels with weather-induced impairments," *IEEE Trans. Aerospace Electron. Syst.*, vol. 38, pp. 779–790, July 2002.

[24] J. P. Choi, "Channel prediction and adaptation over satellite channels with weather induced impairments," Master's thesis, Lab. Information and Decision Syst., Massachusetts Inst. Technol., Cambridge, 2000.

[25] C. E. Mayer, B. E. Jaeger, R. K. Crane, and X. Wang, "Ka-Band scintillations: Measurements and model predictions," *Proc. IEEE*, vol. 85, pp. 936–945, June 1997.

[26] M. S. Alouini, S. A. Borgsmiller, and P. G. Steffes, "Channel characterization and modeling for Ka-Band very small aperture terminals," *Proc. IEEE*, vol. 85, pp. 981–997, June 1997.

[27] Y. Karasawa, M. Yamada, and J. E. Allnett, "A new prediction method for tropospheric scintillation on earth-space paths," *IEEE Trans. Antennas Propagat.*, vol. AP-36, pp. 1608–1614, Nov. 1988.

[28] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J.*, vol. 42, pp. 1977–1997, Sept. 1963.

[29] E. N. Gilbert, "Capacity of a burst-noise channel," *Bell Syst. Tech. J.*, vol. 39, pp. 1253–1265, Sept. 1960.

Michael J. Neely received B.S. degrees in electrical engineering and mathematics from the University of Maryland, College Park, in 1997. He received a NDSEG Fellowship from the Department of Defense for graduate study at the Massachusetts Institute of Technology (MIT), Cambridge, where he received the M.S. degree in electrical engineering and computer science in 1999. He is currently working toward the Ph.D. degree in the Laboratory for Information and Decision Systems (LIDS), MIT.

His research interests are in the areas of satellite and wireless networks and queueing theory.

Mr. Neely is a member of Tau Beta Pi and Phi Beta Kappa.

Eytan Modiano received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, in 1986 and the M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, in 1989 and 1992, respectively.

He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Postdoctoral Fellow during 1992–1993, while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999, he was with the Communications Division at the Lincoln Laboratory, Massachusetts Institute of Technology (MIT), Cambridge, where he designed communication protocols for satellite, wireless, and optical networks, and was the Project Leader for Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999, he has been a Member of the Faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high-speed networks.

Charles E. Rohrs (S'78–M'82–SM'88) received the B.S. degree from the University of Notre Dame, Notre Dame, IN, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 1976, 1978, and 1982, respectively.

He is a Principle Research Scientist at MIT, and is also affiliated with the Tellabs Research Center, Mishawaka, IN, which is the research arm of Tellabs Operations, Inc., a manufacturer of telecommunications equipment for public service network providers. Before becoming a Tellabs Fellow in 1995, he was Director of the Tellabs Research Center for ten years. He also served on the faculty at the University of Notre Dame from 1982 to 1997, and was a Visiting Professor at MIT from 1997 to 2000, at which time he began his current position. While best known for his early contributions to the study of robustness in adaptive control, he has also contributed work in adaptive signal processing, communication theory, and communication networks.