Universal Scheduling for Networks with Arbitrary Traffic, Channels, and Mobility

Michael J. Neely

***Abstract*— We extend stochastic network optimization theory to treat networks with arbitrary sample paths for arrivals, channels, and mobility. The network can experience unexpected link or node failures, traffic bursts, and topology changes, and there are no probabilistic assumptions describing these time varying events. Performance of our scheduling algorithm is compared against an ideal T-slot lookahead policy that can make optimal decisions based on knowledge up to T-slots into the future. We develop a simple non-anticipating algorithm that provides network throughput-utility that is arbitrarily close to (or better than) that of the T-slot lookahead policy, with a tradeoff in the worst case queue backlog kept at any queue. The same policy offers even stronger performance, closely matching that of an ideal *infinite lookahead* policy, when ergodic assumptions are imposed. Our analysis uses a sample path version of Lyapunov drift and provides a methodology for optimizing time averages in general time-varying optimization problems.**

***Index Terms*— Queueing analysis, opportunistic scheduling, internet, routing, flow control, wireless networks, optimization**

I. Introduction

Networks experience unexpected events. Consider the network of Fig. 1 and focus on the session that sends a stream of packets from node A to node D. Suppose that several paths are used, but due to congestion on other links, the primary path that can deliver the most data is the path A, B, C, D. However, suppose that there is a failure at node B in the middle of the session. An algorithm with perfect knowledge of the future would take advantage of the path A, B, C, D while it is available, and would switch to alternate paths before the failure occurs. The algorithm would also be able to predict the traffic load on different links at different times, and would optimally route in anticipation of these events.

The above example holds if the network of Fig. 1 is a wireline network, a wireless network, or a mixture of wired and wireless connections. As another example, suppose the network contains an additional mobile wireless node E, and that the following unexpected event occurs: Node E moves into close proximity to node A, allowing a large number of packets to be sent to it. It then moves into close proximity to node D, providing an opportunity to transmit packets to this destination node. If this event could be anticipated, we could take advantage of it and improve the short term throughput by routing many packets over the relay E.

Michael J. Neely is with the Electrical Engineering department at the University of Southern California, Los Angeles, CA. (web: http://www-rcf.usc.edu/∼mjneely).

This material is supported in part by one or more of the following: the DARPA IT-MANET program grant W911NF-07-0028, the NSF Career grant CCF-0747525, and continuing through participation in the Network Science Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory.

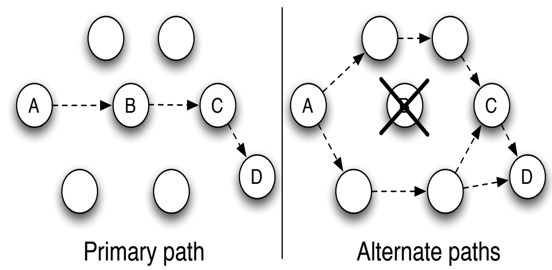

Fig. 1. A primary path from A to D, with alternative paths shown in the event of a failure at node B.

These examples illustrate different types of unexpected events that can be exploited to improve performance. There are of course even more complex sequences of arrival, channel, and mobility events that, if known in advance, could be exploited to yield improved performance. However, because realistic networks do not have knowledge of the future, it is not clear if these events can be practically used. Surprisingly, this paper shows that it is possible to reap the benefits of these time varying events without any knowledge of the future. We show that a simple non-anticipating policy can closely track the performance of an ideal *T-slot lookahead policy* that has perfect knowledge of the future up to T slots. Proximity to the performance of the T-slot lookahead policy comes with a tradeoff in the worst case backlog stored in any queue of the network, which also affects a tradeoff in network delay.

Specifically, we treat networks with slotted time with normalized slots $t \in \{0, 1, 2, \ldots\}$. We measure network utility over an interval of timeslots according to a concave function of the time average throughput vector achieved over that interval. We show that for any positive integer frame size T, and any interval that consists of R frames of T slots, the utility achieved over the interval is greater than or equal to the utility achieved by using the T-slot lookahead policy over each of the R frames, minus a "fudge factor" that has the form:

$$\textit{fudge factor} = \frac{B_1 T}{V} + \frac{B_2 V}{RT}$$

where B_1 and B_2 are constants, and V is a positive parameter that can be chosen as desired to make B_1T/V arbitrarily small, with a tradeoff in worst case queue backlog that is $O(V)$. This shows that we reap almost the same benefits of knowing the future up to T slots if we choose V suitably large and if we wait for the completion of R frames of size T, where R

is sufficiently large to make $B_2V/(RT)$ small. Remarkably, the constants B_1 and B_2 can be explicitly computed in advance, without any assumptions on the underlying stochastic processes that describe the time varying events. Further, in the case when the utility function is linear, we have $B_2 = 0$.

This establishes a *universal scheduling paradigm* that shows a single network algorithm can provide strong mathematical guarantees for any network and for any time varying sample paths. The algorithm that we use is not new: It is a modified version of the backpressure based "drift-plus-penalty" algorithms that we previously developed and used in different contexts in our prior work [1][2][3][4]. These algorithms were originally developed for the case when new arrivals and new channel states are independent and identically distributed (i.i.d.) over slots, and were analyzed using a Lyapunov drift defined as an expectation over the underlying probability distribution. Extended non-i.i.d. models are treated for stability in [5][4][6], for joint stability and utility optimization in [7][8], and in a fluid limit sense in [9]. The works [10][4][11] address "instantaneous capacity regions" and "instantaneous traffic rates" for non-ergodic situations. However, the prior non-ergodic analysis [10][4][11] still assumes an underlying probability model, and makes assumptions about traffic rates and network capacity with respect to this model.

The analysis in this paper is new and uses a sample path version of Lyapunov drift, without any probabilistic assumptions. This allows treatment of realistic channels and traffic traces. Because arbitrary sample paths may not have well defined time averages, typical *equilibrium* notions of *network capacity* and *optimal time average utility* cannot be used. We thus use a new metric that measures performance with respect to ideal T-slot lookahead policies. This is a possible framework for treating the open questions identified in [12] concerning *non-equilibrium network theory*. Further, our results provide universal techniques for optimizing time averages that are useful for other types of time-varying systems.

A. Comparison to Related Work

We note that *universal algorithms* are important in other fields. For example, the universal Lempel-Ziv data compression algorithm operates on arbitrary files [13], and universal stock portfolio allocation algorithms hold for arbitrary price sample paths [14][15][16][17][18]. Prior work in the area of *competitive ratio analysis* considers network scheduling problems with arbitrary sample paths in a different context [19][20][21][22]. Work in [20] considers a large class of admission control problems for networks with random arrivals that earn revenue if accepted. An algorithm is developed that yields revenue that differs by a factor of $\Theta(\log(N))$ from that of an ideal algorithm with perfect knowledge of the future, where N is the number of network nodes. Further, this asymptotic ratio is shown to be optimal, in the sense that there is always a *worst case* sequence of packet arrivals that can reduce revenue by this amount. Related $\Theta(\log(N))$ competitive ratio results are developed for energy optimization in [21] and for wireless admission control in [22]. The works [19][20][21][22] do not consider networks with time varying channels or mobility, and do not treat (or exploit) network queueing. An adversarial queueing theory example in [23] shows that, if channels are time varying, the competitive ratio can be much worse than logarithmic, even for a simple packet-based network with a single link.

Our work treats the difficult case of multi-hop networks with arbitrary traffic, time varying channels, and mobility. However, rather than pursuing a competitive ratio analysis, we measure performance against a T-slot lookahead metric. We develop an algorithm that closely tracks the performance of an ideal T-slot lookahead policy, for any arbitrary (but finite) T. This does not imply that the algorithm has an optimal competitive ratio of 1, because the utility of a T-slot lookahead policy for finite T may not be as good as the performance of an *infinite* lookahead policy. However, it turns out that our policy indeed approaches an optimal competitive ratio of 1 (measured with respect to an infinite lookahead policy) under the special case when the time varying events are ergodic.

II. Network Model

Consider a network with N nodes that operates in slotted time. There are M sessions, and we let $\boldsymbol{A}(t) = (A_1(t), \ldots, A_M(t))$ be the vector of data that exogenously arrives to the transport layer for each session on slot t (measured either in integer units of *packets* or real units of *bits*). We assume that arrivals are bounded by constants A_m^{max}:

$$0 \leq A_m(t) \leq A_m^{max} \ \ \forall t$$

A network where all sources always have data to send can be modeled as one with $A_m(t) = A_m^{max}$ for all t.

Each session $m \in \{1, \ldots, M\}$ has a particular source node and destination node. Data delivery takes place by transmissions over possibly multi-hop paths. We assume that a *transport layer flow controller* observes $A_m(t)$ every slot and decides how much of this data to add to the network layer at its source node, and how much to drop. Let $\boldsymbol{x}(t) = (x_1(t), \ldots, x_M(t))$ be a vector of *flow control decision variables* on slot t. These decisions are made subject to the constraints:

$$0 \leq x_m(t) \leq A_m(t) \ \ \forall m \in \{1, \ldots, M\}, \forall t \qquad (1)$$

All data that is intended for destination node $c \in \{1, \ldots, N\}$ is called *commodity c data*, regardless of its particular session. For each $n \in \{1, \ldots, N\}$ and $c \in \{1, \ldots, N\}$, let $\mathcal{M}_n^{(c)}$ denote the set of all sessions $m \in \{1, \ldots, M\}$ that have source node n and commodity c. All data is queued according to its commodity, and we define $Q_n^{(c)}(t)$ as the amount of commodity c data in node n on slot t. We assume that $Q_n^{(n)}(t) = 0$ for all t, as data that reaches its destination is removed from the network. Let $\boldsymbol{Q}(t)$ denote the matrix of current queue backlogs for all nodes and commodities.

The queue backlogs change from slot to slot as follows:

$$Q_n^{(c)}(t+1) = Q_n^{(c)}(t) - \sum_{j=1}^{N} \tilde{\mu}_{nj}^{(c)}(t) + \sum_{i=1}^{N} \tilde{\mu}_{in}^{(c)}(t) + \sum_{m \in \mathcal{M}_n^{(c)}} x_m(t)$$

where $\tilde{\mu}_{ij}^{(c)}(t)$ denotes the actual amount of commodity c data transmitted from node i to node j (i.e., over link (i,j)) on slot

t. It is useful to define *transmission decision variables* $\mu_{ij}^{(c)}(t)$ as the transmission rate *offered* by link (i, j) to commodity c data, where this full amount is used only if there is that much commodity c data available at node i, so that:

$$\tilde{\mu}_{ij}^{(c)}(t) \leq \mu_{ij}^{(c)}(t) \;\; \forall i, j, c \in \{1, \ldots, N\}, \forall t$$

Thus, for all $n \neq c$:

$$Q_n^{(c)}(t+1) \leq \max\left[Q_n^{(c)}(t) - \sum_{j=1}^N \mu_{nj}^{(c)}(t), 0\right] + \sum_{i=1}^N \mu_{in}^{(c)}(t) + \sum_{m \in \mathcal{M}_n^{(c)}} x_m(t) \quad (2)$$

The inequality is because the actual endogenous arrivals $\tilde{\mu}_{in}^{(c)}(t)$ may not be as large as $\mu_{in}^{(c)}(t)$ if there is little or no data available for transmission over that link. Our results hold for any scheduling choices that satisfy (2). For example, if there are 5 units of commodity c data in node n, and there are two outgoing links offering non-zero rates $\mu_{n1}^{(c)}(t) = 4$ and $\mu_{n2}^{(c)}(t) = 3$, then the 5 units will be cleared from node n, but can be distributed over the two links in different ways, such as $\tilde{\mu}_{n1}^{(c)}(t) = 4$, $\tilde{\mu}_{n2}^{(c)}(t) = 1$, or $\tilde{\mu}_{n1}^{(c)}(t) = 3$, $\tilde{\mu}_{n2}^{(c)}(t) = 2$, etc.

A. Transmission Variables

Let $S(t)$ represent the *topology state* of the network on slot t, observed on each slot as in [1]. The value of $S(t)$ is an abstract and possibly multi-dimensional quantity that describes the current link conditions between all nodes under the current slot. The collection of all transmission rates that can be offered over each link (i, j) of the network is given by a general transmission rate function $\boldsymbol{b}(I(t), S(t))$:[1]

$$\boldsymbol{b}(I(t), S(t)) = (b_{ij}(I(t), S(t)))_{i,j \in \{1,\ldots,N\}, i \neq j}$$

where $I(t)$ is a general network-wide resource allocation decision (such as link scheduling, bandwidth selection, modulation, etc.) and takes values in some abstract set $\mathcal{I}_{S(t)}$ that possibly depends on the current $S(t)$. We assume the transmission rate function $b_{ij}(I(t), S(t))$ is non-negative and bounded by a finite constant b_{ij}^{max} for all (i, j), $I(t)$, and $S(t)$.

Every slot the network controller observes the current $S(t)$ and makes a resource allocation decision $I(t) \in \mathcal{I}_{S(t)}$. The controller then chooses $\mu_{ij}^{(c)}(t)$ variables subject to the following constraints:

$$\mu_{ij}^{(c)}(t) \geq 0 \quad \forall i, j, c \in \{1, \ldots, N\} \quad (3)$$

$$\mu_{ii}^{(c)}(t) = \mu_{ij}^{(i)}(t) = 0 \quad \forall i, j, c \in \{1, \ldots, N\} \quad (4)$$

$$\mu_{ij}^{(c)}(t) = 0 \text{ if } (i, j) \notin \mathcal{L}^{(c)} \quad \forall i, j, c \in \{1, \ldots, N\} \quad (5)$$

$$\sum_{c=1}^N \mu_{ij}^{(c)} \leq b_{ij}(I(t), S(t)) \quad \forall i, j \in \{1, \ldots, N\} \quad (6)$$

Constraint (4) ensures data is not transmitted from a node to itself, or transmitted again once it reaches its destination. Constraint (6) ensures the total transmission rate offered over link (i, j) on slot t is at most $b_{ij}(I(t), S(t))$. Constraint (5) restricts transmission of commodity c data to a pre-specified set of links $\mathcal{L}^{(c)}$, which is sometimes useful. The case of unrestricted routing is covered by defining $\mathcal{L}^{(c)}$ as the set of all network links.

[1] It is worth noting now that for networks with orthogonal channels, our "max-weight" transmission algorithm (to be defined in the next subsection) decouples to allow nodes to make transmission decisions based only on those components of the current topology state $S(t)$ that relate to their own local channels. Of course, for wireless interference networks, all channels are coupled, although distributed approximations of max-weight transmission exist in this case, see Section 4.7 and Corollary 5.2 in [1].

B. The Utility Optimization Problem

For simplicity, let $\omega(t) \triangleq [\boldsymbol{A}(t), S(t)]$ represent the random *network events* that are observed on slot t. Let $\alpha(t)$ represent the collection of *control actions* taken on slot t, constrained to a set $\mathcal{A}_{\omega(t)}$. Specifically:

$$\alpha(t) \triangleq [I(t), (\mu_{ij}^{(c)}(t)), \boldsymbol{x}(t)]$$

The constraint $\alpha(t) \in \mathcal{A}_{\omega(t)}$ is a simple way of representing all constraints (1), (3)-(6), and $I(t) \in \mathcal{I}_{S(t)}$.

Fix a finite implementation time t_{end}, and define $\overline{x}_m$, $\overline{\mu}_{ab}^{(c)}$ as time averages over $\tau \in \{0, 1, \ldots, t_{end} - 1\}$:

$$\overline{\mu}_{ab}^{(c)} \triangleq \frac{1}{t_{end}} \sum_{\tau=0}^{t_{end}-1} \mu_{ab}^{(c)}(\tau)$$

$$\overline{\boldsymbol{x}} \triangleq \frac{1}{t_{end}} \sum_{\tau=0}^{t_{end}-1} \boldsymbol{x}(\tau) = (\overline{x}_1, \ldots, \overline{x}_M) \quad (7)$$

Our objective is to solve the following problem:[2]

$$\text{Maximize:} \quad \phi(\overline{\boldsymbol{x}}) \quad (8)$$

Subject to:

$$\sum_{m \in \mathcal{M}_n^{(c)}} \overline{x}_m + \sum_{i=1}^N \overline{\mu}_{in}^{(c)} \leq \sum_{j=1}^N \overline{\mu}_{nj}^{(c)} \;, \forall (n, c), n \neq c \quad (9)$$

$$\alpha(t) \in \mathcal{A}_{\omega(t)} \;\; \forall t \in \{0, \ldots, t_{end} - 1\} \quad (10)$$

where $\phi(\overline{\boldsymbol{x}})$ is a continuous, concave, and entrywise non-decreasing function over the hyper-rectangle $\mathcal{R}$, defined:

$$\mathcal{R} \triangleq \{\boldsymbol{x} \in \mathbb{R}^M | 0 \leq x_m \leq A_m^{max} \; \forall m \in \{1, \ldots, M\}\} \quad (11)$$

Define ν_m^{min} and ν_m^{max} respectively as the *infimum* and *supremum* right partial derivative of $\phi(\boldsymbol{x})$ with respect to x_m over the interior of the rectangle $\mathcal{R}$, assumed to be finite.[3] An example is a *separable utility function* such as:

$$\phi(\boldsymbol{x}) = \sum_{m=1}^M \log(1 + \nu_m x_m)$$

for which $\nu_m^{max} = \nu_m$ and $\nu_m^{min} = \nu_m / (1 + \nu_m A_m^{max})$. Another example is:

$$\phi(\boldsymbol{x}) = \min[x_1, \ldots, x_M]$$

for which $\nu_m^{max} = 1$ and $\nu_m^{min} = 0$. Finally, in the linear case $\phi(\boldsymbol{x}) = \sum_{m=1}^M \nu_m x_m$ we have $\nu_m^{min} = \nu_m^{max} = \nu_m$.

[2] A more general optimization problem is treated in our technical report [24] (see also Section V).

[3] Continuous and concave functions over a hyper-rectangle $\mathcal{R}$ have well defined left and right partial derivatives over the interior of $\mathcal{R}$.

C. T-Slot Lookahead Policies

For a given $S(t)$ on slot t, define $\Omega(t)$ as the set of all vectors $[(b_{ij}), (\mu_{ij}^{(c)})]$ such that there exists an action $I \in \mathcal{I}_{S(t)}$ such that $b_{ij} = b_{ij}(I, S(t))$ for all i, j, and $\mu_{ij}^{(c)}$ satisfies $\mu_{ij}^{(c)} \geq 0$, $\mu_{ii}^{(c)} = \mu_{ij}^{(i)} = 0$ for all i, j, c, and:

$$\sum_{c=1}^{N} \mu_{ij}^{(c)} \leq b_{ij} \;\; \forall i, j \in \{1, \ldots, N\}$$

Rather than compare performance to the optimum of the problem (8)-(10) over the full interval $t \in \{0, \ldots, t_{end} - 1\}$, we compare to the maximum utility that can be achieved over successive frames of size T, assuming the time average constraints (9) must be achieved over each frame. Further, in this comparison we allow variables $[(b_{ij}(t)), (\mu_{ij}^{(c)}(t))]$ to be chosen in the extended set $Conv(\Omega(t))$, rather than just $\Omega(t)$.

Specifically, let $T > 0$ be an integer *frame size*. For each integer $r \geq 0$, define the *rth frame* as the interval $\tau \in \{rT, \ldots, rT+T-1\}$. Define F_r^* as the supremum value associated with the following problem (where $\boldsymbol{\gamma}^* \triangleq (\gamma_1^*, \ldots, \gamma_M^*)$):

$$\begin{aligned}
&\text{Maximize:} \quad \phi(\boldsymbol{\gamma}^*) \qquad (12)\\
&\text{Subject to:} \quad \gamma_m^* = \frac{1}{T}\sum_{\tau=rT}^{rT+T-1} x_m^*(\tau) \;\; \forall m \in \{1, \ldots, M\}\\
&\sum_{m \in \mathcal{M}_n^{(c)}} \gamma_m^* + \frac{1}{T}\sum_{\tau=rT}^{rT+T-1}\sum_{i=1}^{N} \mu_{in}^{*(c)}(\tau)\\
&\quad \leq \frac{1}{T}\sum_{\tau=rT}^{rT+T-1}\sum_{j=1}^{N} \mu_{nj}^{*(c)}(\tau) \;\; \forall (n,c), n \neq c\\
&[(b_{ij}^*(t)), (\mu_{ij}^{*(c)}(t))] \in Conv(\Omega(\tau)) \;\; , \forall \tau\\
&0 \leq x_m^*(\tau) \leq A_m(\tau) \;\; \forall m \in \{1, \ldots, M\}, \forall \tau
\end{aligned}$$

The value of F_r^* represents the supremum of the utility in (12) that can be achieved over the frame, considering all policies that satisfy the constraints and that have perfect knowledge of the future $\omega(\tau)$ values over the frame. Note that the trivial solution $\gamma_m = x_m = \mu_{in}^{(c)} = 0$ satisfies all constraints in the T-slot lookahead problem, and hence the problem is always feasible and $F_r^* \geq 0$ for all r. Our new goal is to design a non-anticipating control policy that is implemented over time $t_{end} = RT$ (for some positive integer R), and that satisfies all constraints of the original problem while achieving a total utility that is close to (or larger than) the value of:

$$\frac{1}{R}\sum_{r=0}^{R-1} F_r^* \qquad (13)$$

The problem (8)-(10) might have a strictly larger utility than (13) because it only requires the time average constraints to be met over the full time interval, rather than requiring them to be satisfied on each of the R frames. Nevertheless, when T is large, it is not trivial to achieve the utility value of (13), as this utility is defined over policies that have T-slot lookahead, whereas an actual policy does not have future lookahead capabilities and makes decisions within the smaller set $\Omega(t)$ rather than $Conv(\Omega(t))$.

D. Example for 1-Hop Networks

For intuition, here we describe F_r^* for a wireless downlink system with N queues, arrivals $(A_1(t), \ldots, A_N(t))$, and channel states $S(t) = (S_1(t), \ldots, S_N(t))$ (we have re-indexed the variables because of the one-hop structure of this example). There is a single server that can be allocated to at most one queue per slot, serving at rate $S_n(t)$ if it selects channel n on slot t. The decision vector is $I(t) = (I_1(t), \ldots, I_N(t))$ with $I_n(t) = 1$ if channel n is selected on slot t, and 0 else. The rates for each channel $n \in \{1, \ldots, N\}$ are $\mu_n(t) = b_n(I_n(t), S_n(t))$, where:

$$b_n(I_n(t), S_n(t)) = \begin{cases} S_n(t) & \text{if } I_n(t) = 1 \\ 0 & \text{if } I_n(t) = 0 \end{cases}$$

For a frame size T and frame r consisting of slots $\tau \in \{rT, \ldots, rT+T-1\}$ we have that F_r^* is the solution to:

$$\begin{aligned}
&\text{Maximize:} \quad \phi(\boldsymbol{\gamma}^*) && (14)\\
&\text{Subject to:} \quad \gamma_n^* = \frac{1}{T}\sum_{\tau=rT}^{rT+T-1} \mu_n^*(\tau) \;\; \forall n \in \{1, \ldots, N\} && (15)\\
&0 \leq \mu_n^*(\tau) \leq S_n(\tau) \;\; \forall n \in \{1, \ldots, N\}, \forall \tau && (16)\\
&\sum_{\{n | S_n(\tau) > 0\}} \frac{\mu_n^*(\tau)}{S_n(\tau)} \leq 1 \;\; , \forall \tau && (17)\\
&0 \leq x_n^*(\tau) \leq A_n(\tau) \;\; \forall (i,j), \forall \tau && (18)
\end{aligned}$$

The constraint (17) is a convexification of the more stringent constraint that $(\mu_1(\tau), \ldots, \mu_N(\tau))$ is non-zero in at most 1 component, and satisfies $\mu_n(\tau) = S_n(\tau)$ in any non-zero component n, which applies to the actual system.

III. Sample Path Lyapunov Optimization

The problem (8)-(10) is equivalent to the following, which introduces *auxiliary variables* $\boldsymbol{\gamma}(t) = (\gamma_1(t), \ldots, \gamma_M(t))$:

$$\begin{aligned}
&\text{Maximize:} \quad \phi(\overline{\boldsymbol{\gamma}})\\
&\text{Subject to:} \quad \overline{x}_m \geq \overline{\gamma}_m \;\; \forall m \in \{1, \ldots, M\}\\
&\sum_{m \in \mathcal{M}_n^{(c)}} \overline{x}_m + \sum_{i=1}^{N} \overline{\mu}_{in}^{(c)} \leq \sum_{j=1}^{N} \overline{\mu}_{nj}^{(c)} \;\; , \forall (n,c), n \neq c\\
&\boldsymbol{\gamma}(t) \in \mathcal{R} \; , \; \alpha(t) \in \mathcal{A}_{\omega(t)} \;\; \forall t \in \{0, \ldots, t_{end} - 1\}
\end{aligned}$$

where $\mathcal{R}$ is defined in (11), and $\overline{\boldsymbol{\gamma}}$ is defined:

$$\overline{\boldsymbol{\gamma}} \triangleq \frac{1}{t_{end}}\sum_{\tau=0}^{t_{end}-1} \boldsymbol{\gamma}(\tau)$$

The auxiliary variables $\boldsymbol{\gamma}(t)$ are useful because they are chosen in the rectangle set $\mathcal{R}$ every slot t, and decouple the nonlinear function $\phi(\cdot)$ from the variables $\boldsymbol{x}(t)$ that are constrained differently every slot based on the value of $A_m(t)$ (so that $0 \leq x_m(t) \leq A_m(t)$). The constraints $\overline{x}_m \geq \overline{\gamma}_m$ are enforced by *virtual queues* $G_m(t)$, with update:

$$G_m(t+1) = \max[G_m(t) + \gamma_m(t) - x_m(t), 0] \qquad (19)$$

The intuition is that if $G_m(t)$ is stable, then the time average of $x_m(t)$ is greater than or equal to that of $\gamma_m(t)$ [25][1]. This method of auxiliary variables was developed in [1][2].

Define $\boldsymbol{\Theta}(t) \triangleq [\boldsymbol{Q}(t), \boldsymbol{G}(t)]$ as a vector of all actual and virtual queues, and define the following *Lyapunov function*:

$$L(\boldsymbol{\Theta}(t)) \triangleq \frac{1}{2}\sum_{n,c} Q_n^{(c)}(t)^2 + \frac{1}{2}\sum_{m=1}^{M} G_m(t)^2$$

Let $\Delta_T(t)$ represent the T-*slot sample path Lyapunov drift* associated with particular actions implemented over the interval $\{t, \ldots, t+T-1\}$ when the queues have state $\boldsymbol{\Theta}(t)$ at the start of the interval:[4]

$$\Delta_T(t) \triangleq L(\boldsymbol{\Theta}(t+T)) - L(\boldsymbol{\Theta}(t)) \qquad (20)$$

This notion of T-slot drift differs from that given in [1] in that it does not involve an expectation.

[4]The value $\Delta_T(t)$ depends on $\boldsymbol{\Theta}(t)$, the random events $\{\omega(t), \ldots, \omega(t+T-1)\}$, and the control actions $\{\alpha(t), \ldots, \alpha(t+T-1)\}$.

A. The Drift-Plus-Penalty Method

It is difficult to know the T-slot drift because it depends on future (and hence unknown) $\omega(t)$ values. Thus, following the approach [1], our policy every slot t observes the current $\omega(t)$ and $\boldsymbol{\Theta}(t)$ and chooses a control action $\alpha(t) \in \mathcal{A}_{\omega(t)}$ and auxiliary variables $\boldsymbol{\gamma}(t) \in \mathcal{R}$ to come within an additive constant of minimizing an upper bound on the following 1-slot *drift-plus-penalty* expression:

$$\Delta_1(t) - V\phi(\boldsymbol{\gamma}(t))$$

where $V \geq 0$ is a control parameter chosen in advance to affect a performance tradeoff.

Lemma 1: The 1-slot drift $\Delta_1(t)$ satisfies:

$$\Delta_1(t) - V\phi(\boldsymbol{\gamma}(t)) \leq B - V\phi(\boldsymbol{\gamma}(t)) + \sum_{m=1}^{M} G_m(t)[\gamma_m(t) - x_m(t)]$$
$$+ \sum_{n,c} Q_n^{(c)}(t)\left[\sum_{m\in\mathcal{M}_n^{(c)}} x_m(t) + \sum_{i=1}^{N}\mu_{in}^{(c)}(t) - \sum_{j=1}^{N}\mu_{nj}^{(c)}(t)\right]$$

where the constant B is defined:

$$B \triangleq \frac{1}{2}\sum_{m=1}^{M}(A_m^{max})^2 + \frac{1}{2}\sum_{n=1}^{N}[(A_n^{max,in})^2 + (b_n^{max,out})^2]$$

where $A_n^{max,in}$ and $b_n^{max,out}$ are bounds on the maximum arrivals to and departures from node n on a given slot, including all commodities (where the arrival bound includes both exogenous and endogenous arrivals).

Proof: The proof involves squaring the queue update equations for $Q_n^{(c)}(t)$, $G_m(t)$ in (2), (19) and then adding the "penalty" $-V\phi(\boldsymbol{\gamma}(t))$ to both sides. Details are in [24]. □

B. The Universal Network Scheduling Algorithm

Our algorithm makes decisions about auxiliary variables $\boldsymbol{\gamma}(t)$ and flow control variables $\boldsymbol{x}(t)$ to minimize the right-hand-side of the drift-plus-penalty expression in Lemma 1. It also chooses routing and resource allocation variables $(\mu_{ab}^{(c)}(t))$, $I(t)$ to "approximately" minimize the right-hand-side of this expression (to within an additive constant), where the approximation enables a deterministic queue backlog bound.

- (Auxiliary Variables) For each slot t, the $\boldsymbol{G}(t)$ queues are observed and $\boldsymbol{\gamma}(t)$ is chosen to solve:

  $$\text{Maximize:} \quad V\phi(\boldsymbol{\gamma}(t)) - \textstyle\sum_{m=1}^{M} G_m(t)\gamma_m(t) \quad (21)$$
  $$\text{Subject to:} \quad 0 \leq \gamma_m(t) \leq A_m^{max} \;\; \forall m \in \{1,\ldots,M\} \quad (22)$$

  This amounts maximization of M separate single-variable concave functions in the case when $\phi(\boldsymbol{\gamma})$ has the separable structure $\phi(\boldsymbol{\gamma}) = \sum_{m=1}^{M}\phi_m(\gamma_m)$.
- (Flow Control) For each slot t, each session m observes $A_m(t)$ and the queue values $G_m(t)$, $Q_{n_m}^{(c_m)}(t)$ (where n_m denotes the source node of session m, and c_m represents its destination). Note that these queues are all local to the source node of the session, and hence can be observed easily. It then chooses $x_m(t)$ to solve:

  $$\text{Maximize:} \quad G_m(t)x_m(t) - Q_{n_m}^{(c_m)}(t)x_m(t) \quad (23)$$
  $$\text{Subject to:} \quad 0 \leq x_m(t) \leq A_m(t)$$

  This reduces to the "bang-bang" flow control decision of choosing $x_m(t) = A_m(t)$ if $Q_{n_m}^{(c_m)}(t) \leq G_m(t)$, and $x_m(t) = 0$ otherwise.
- (Resource Allocation and Transmission) For each slot t, the network controller observes queue backlogs $\{Q_n^{(c)}(t)\}$ and the topology state $S(t)$ and chooses $I(t) \in \mathcal{I}_{S(t)}$ and $\{\mu_{ij}^{(c)}(t)\}$ subject to (3)-(6) to *approximately* solve:

  $$\text{Max:} \; \textstyle\sum_{n,c} Q_n^{(c)}(t)[\sum_{j=1}^{N}\mu_{nj}^{(c)}(t) - \sum_{i=1}^{N}\mu_{in}^{(c)}(t)] \quad (24)$$
  $$\text{S.t.:} \quad I(t) \in \mathcal{I}_{S(t)} \text{ and (3)-(6)}$$

  The specific choices that are made in the approximation are detailed in the next subsection.
- (Queue Updates) Update the queues $G_m(t)$ and $Q_n^{(c)}(t)$ according to (19) and (2).

C. Resource Allocation and Transmission

There are two modifications to the max-weight rule (24) that we use below (similar to [25][10]): The first modifies the backpressure problem to ensure bounded queues. The second allows for approximate implementations of a max-weight rule. Specifically, define *differential backlogs* $W_{ij}^{(c)}(t)$ as follows:

$$W_{ij}^{(c)}(t) \triangleq Q_i^{(c)}(t) - Q_j^{(c)}(t)$$

Define $\hat{W}_{ij}^{(c)}(t)$ as follows:

$$\hat{W}_{ij}^{(c)}(t) \triangleq \begin{cases} W_{ij}^{(c)}(t) + \theta_i^{(c)} - \theta_j^{(c)} & \text{if } Q_j^{(c)}(t) \leq Q^{max} - \beta_j \\ -1 & \text{otherwise} \end{cases} \quad (25)$$

where for each $j \in \{1,\ldots,N\}$, β_j is the largest amount of any commodity that can enter node j, considering both exogenous and endogenous arrivals (this is finite by the boundedness assumptions), and where Q^{max} is defined:

$$Q^{max} \triangleq V\nu^{max} + A^{max} + \beta^{max}$$

where ν^{max}, A^{max}, β^{max} are given by:

$$\nu^{max} \triangleq \max_{m\in\{1,\ldots,M\}} \nu_m^{max} \;\; , \;\; A^{max} \triangleq \max_{m\in\{1,\ldots,M\}} A_m^{max}$$
$$\beta^{max} \triangleq \max_{n\in\{1,\ldots,N\}} \beta_n$$

Finally, the values $\theta_i^{(c)}$ are any non-negative weights that represent some type of estimate of the distance from node i to destination c (possibly being zero if there is no such estimate available). Such weights are known to experimentally improve delay by biasing routing decisions towards favorable directions [6][26][10]. Then define $\hat{W}_{ij}(t)$ as:

$$\hat{W}_{ij}(t) \triangleq \max_{c\in\{1,\ldots,N\}|(i,j)\in\mathcal{L}^{(c)}} \max[\hat{W}_{ij}^{(c)}(t), 0]$$

and define $\hat{c}^*_{ij}(t)$ as the maximizing commodity. Choose $I(t) \in \mathcal{I}_{S(t)}$ to *come within an additive constant* $C \geq 0$ *of maximizing*:

$$\sum_{i=1}^N \sum_{j=1}^N b_{ij}(I(t), S(t))\hat{W}_{ij}(t) \quad (26)$$

and choose transmission variables:

$$\mu_{ij}^{(c)}(t) = \begin{cases} b_{ij}(I(t), S(t)) & \text{if } c = \hat{c}^*_{ij}(t) \text{ and } \hat{W}_{ij}^{(c)}(t) \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

Theorem 1: Suppose the following: We implement the above universal network scheduling algorithm using any parameter $V \geq 0$ and any weights $\theta_i^{(c)} \geq 0$. Every slot we choose $I(t) \in \mathcal{I}_{S(t)}$ to come within an additive constant $C \geq 0$ of maximizing the expression (26). Suppose all actual queues are initially empty, so that $\boldsymbol{Q}(0) = \boldsymbol{0}$, and that $G_m(t)$ queues are initialized to $G_m(0) = \max[V\nu_m^{min} - A_m^{max}, 0]$. Then:

(a) The queues $Q_n^{(c)}(t)$ and $G_m(t)$ are *deterministically bounded* for all t, so that for all $(n, c) \in \{1, \ldots, N\}^2$ and all $m \in \{1, \ldots, M\}$:

$$Q_n^{(c)}(t) \leq Q^{max} \triangleq V\nu^{max} + A^{max} + \beta^{max} \quad \forall t \geq 0 \quad (28)$$

$$G_m^{min} \leq G_m(t) \leq G_m^{max} \quad \forall t \geq 0 \quad (29)$$

where constants G_m^{min} and G_m^{max} are defined:

$$\begin{aligned} G_m^{min} &\triangleq \max[V\nu_m^{min} - A_m^{max}, 0] \\ G_m^{max} &\triangleq V\nu_m^{max} + A_m^{max} \end{aligned}$$

(b) For any integer frame size $T > 0$ and any integer number of frames $R > 0$, define $t_{end} \triangleq \{0, 1, \ldots, RT - 1\}$. Then time average admission rates $\overline{\boldsymbol{x}}$, defined as time averages over the first t_{end} slots as in (7), satisfy:

$$\begin{aligned} \phi(\overline{\boldsymbol{x}}) \geq & \frac{1}{R}\sum_{r=0}^{R-1} F_r^* - \frac{\tilde{B}T + C}{V} \\ & - \sum_{m=1}^M \frac{\nu_m^{max}(V(\nu_m^{max} - \nu_m^{min}) + 2A_m^{max})}{RT} \end{aligned} \quad (30)$$

where $\frac{1}{R}\sum_{r=0}^{R-1} F_r^*$ represents the target utility (13) associated with implementing the T-slot lookahead policy over R frames of size T. The constant $\tilde{B}$ is equal to B plus an additional term related to worst case b^{max}, A_m^{max} parameters and the $\theta_i^{(c)}$ values (and is independent of V, R, T). The constant $\tilde{B}$ is computed explicitly in [24].

D. Discussion of Theorem 1

Note that the final term on the right-hand-side of (30) is *independent of* V in the special case when $\phi(\boldsymbol{x})$ is linear, so that $\nu_m^{max} = \nu_m^{min}$ for all m. When $\phi(\boldsymbol{x})$ is nonlinear, the above theorem shows that the "fudge factor" between the achieved utility $\phi(\overline{\boldsymbol{x}})$ and the target $\frac{1}{R}\sum_{r=0}^{T-1} F_r^*$ has the form:

$$\textit{fudge factor} = \frac{B_1 T}{V} + \frac{B_2 V}{RT}$$

where constants B_1 and B_2 are defined:

$$\begin{aligned} B_1 &\triangleq \tilde{B} + C/T \\ B_2 &\triangleq \sum_{m=1}^M [\nu_m^{max}(\nu_m^{max} - \nu_m^{min}) + 2\nu_m^{max} A_m^{max}/V] \end{aligned}$$

Finally, we note that the initial conditions for $Q_n^{(c)}(t)$ and $G_m(t)$ given in Theorem 1 are not crucial. The same inequalities (28)-(29) hold for all slots $t \geq 0$ provided that they also hold for $t = 0$, and (32) of the proof shows an additional fudge-factor that quickly vanishes is added to the utility bound.

E. Proof of Theorem 1 part (a)

Lemma 2: If for a given $m \in \{1, \ldots, M\}$ and a given slot t, we have $G_m(t) > V\nu_m^{max}$, then the auxiliary variable selection rule (21)-(22) chooses $\gamma_m(t) = 0$, and hence $G_m(t)$ cannot increase on the next slot. Alternatively, if $G_m(t) < V\nu_m^{min}$, then $G_m(t)$ cannot decrease on the next slot.

Proof: (Lemma 2) Suppose $G_m(t) > V\nu_m^{max}$ for a particular $m \in \{1, \ldots, M\}$. Let $\boldsymbol{\gamma} \in \mathcal{R}$, and define $\boldsymbol{\gamma}_m^{(0)}$ as $\boldsymbol{\gamma}$ with the mth entry set to 0. Because $\phi(\boldsymbol{\gamma})$ is non-decreasing with maximum mth partial derivative ν_m^{max}, we have for any $\boldsymbol{\gamma}$ such that $\boldsymbol{\gamma} \in \mathcal{R}$:

$$\begin{aligned} V\phi(\boldsymbol{\gamma}) - \sum_{i=1}^M G_i(t)\gamma_i &\leq V\phi(\boldsymbol{\gamma}_m^{(0)}) - \sum_{i \neq m} G_i(t)\gamma_i \\ &\quad + \gamma_m(V\nu_m^{max} - G_m(t)) \\ &\leq V\phi_m(\boldsymbol{\gamma}_m^{(0)}) - \sum_{i \neq m} G_i(t)\gamma_i \end{aligned}$$

Equality in the above holds if and only if $\gamma_m = 0$, because $(V\nu_m^{max} - G_m(t)) < 0$. Thus, the rule (21)-(22) must choose $\gamma_m(t) = 0$. It follows by the $G_m(t)$ update rule (19) that $G_m(t)$ cannot increase on the next slot. It can similarly be shown that $G_m(t)$ cannot decrease if it is less than $V\nu_m^{min}$. □

Proof: (Theorem 1 part (a)) Fix $m \in \{1, \ldots, M\}$. We first show that $G_m(t)$ is bounded by $V\nu_m^{max} + A_m^{max}$ for all t. Suppose this holds for a given slot t (we know it holds for $t = 0$ by assumption). We prove it also holds for $t+1$. If $G_m(t) \leq V\nu_m^{max}$, then $G_m(t+1) \leq V\nu_m^{max} + A_m^{max}$, since $G_m(t)$ can increase by at most A_m^{max} on each slot (see dynamics in (19)). On the other hand, if $V\nu_m^{max} < G_m(t) \leq V\nu_m^{max} + A_m^{max}$, then by Lemma 2 we know $G_m(t)$ cannot increase on the next slot, and so $G_m(t+1) \leq G_m(t) \leq V\nu_m^{max} + A_m^{max}$. It follows by induction that the bound holds for all $t \geq 0$.

Similarly, because Lemma 2 also ensures $G_m(t)$ cannot decrease when it is less than $V\nu_m^{min}$, it follows easily that $G_m(t) \geq \max[V\nu_m^{min} - A_m^{max}, 0]$ for all t.

We now prove the $Q_n^{(c)}(t)$ bound. Suppose that $Q_n^{(c)}(t) \leq Q^{max}$ for all (n, c) for a given slot t (this clearly holds for $t = 0$). We show that it also holds for slot $t+1$. Take any particular (n, c). If $Q_n^{(c)}(t) \leq Q^{max} - \beta_n$, then the desired bound must hold on slot $t+1$ because at most β_n units of new commodity c data (considering both exogenous and endogenous arrivals) can enter node n in one slot (by definition of β_n).

Now suppose that $Q^{max} - \beta_n < Q_n^{(c)}(t) \leq Q^{max}$. Thus (because $G_m(t) \leq V\nu^{max} + A^{max}$ for all m and all t):

$$G_m(t) \leq V\nu^{max} + A^{max} \leq Q^{max} - \beta_n < Q_n^{(c)}(t)$$

Thus, $Q_n^{(c)}(t) > G_m(t)$ for all $m \in \{1, \ldots, M\}$. It follows by the flow control algorithm (23) that $x_m(t) = 0$ for all sources m for which $n_m = n$ and $c_m = c$, that is, all sources that have source node n and have commodity c. Hence, there can be no exogenous arrivals on slot t. We next show that there can be no endogenous arrivals from other nodes either, and so $Q_n^{(c)}(t)$ cannot increase on the next slot (proving the result). To this end, note that because $Q_n^{(c)}(t) > Q^{max} - \beta_n$, we have $\hat{W}_{in}^{(c)}(t) = -1$ for all other nodes i that can send data to node n (recall the definition of $\hat{W}_{ij}^{(c)}(t)$ in (25)). It follows by (27) that $\mu_{in}^{(c)}(t) = 0$ for all i, and so node n can receive no new commodity c data on slot t. □

IV. Utility Analysis — Proof of Theorem 1(b)

Lemma 3: The universal scheduling algorithm with any parameter $V \geq 0$, any weights $\theta_i^{(c)}$, and any approximate implementation that solves (26) to within a constant $C \geq 0$ every slot, satisfies for any slot $t \geq 0$ and any integer $T > 0$:

$$\begin{aligned}
\Delta_T(t) - V\sum_{\tau=t}^{t+T-1}\phi(\boldsymbol{\gamma}(\tau)) \leq \tilde{B}T^2 + CT - VT\phi(\boldsymbol{\gamma}^*) \\
+ \sum_{m=1}^{M} G_m(t)\sum_{\tau=t}^{t+T-1}[\gamma_m^* - x_m^*(\tau)] \\
+ \sum_{n,c} Q_n^{(c)}(t)\sum_{\tau=t}^{t+T-1}\left[\sum_{i=1}^{N}\mu_{in}^{*(c)}(\tau) - \sum_{j=1}^{N}\mu_{nj}^{*(c)}(\tau)\right] \\
+ \sum_{n,c} Q_n^{(c)}(t)\sum_{\tau=t}^{t+T-1}\sum_{m\in\mathcal{M}_n^{(c)}} x_m^*(\tau)
\end{aligned}$$

where $\boldsymbol{\gamma}^*$ is any vector in $\mathcal{R}$, $x_m^*(\tau)$ are any values that satisfy $0 \leq x_m^*(\tau) \leq A_m(\tau)$, and $\mu_{ab}^{*(c)}(\tau)$ are any values in $Conv(\Omega(\tau))$. The constant $\tilde{B}$ is equal to B plus an additional term that depends on the worst case b^{max} and A_m^{max} values.

Proof: The proof follows from the fact that decision variables $\mu_{ab}^{(c)}(t)$ are chosen in $\Omega(t)$ to minimize the right-hand-side of the drift-plus-penalty expression in Lemma 1 to within an additive constant on every slot t. Because they appear in linear terms on the right-hand-side, they also minimize over all variables in $Conv(\Omega(t))$. See [24] for details, where $\tilde{B}$ is explicitly computed. □

Now fix a frame size T and define $t_r \triangleq rT$ as the start of frame r, for $r \in \{0, 1, \ldots, R-1\}$. We apply the above lemma with $t = t_r$, and with γ_m^*, $x_m^*(\tau)$, $b_{ij}^*(\tau)$, $\mu_{ab}^{*(c)}(\tau)$ defined as the actions that solve the T-slot lookahead problem (12) over the frame $\tau \in \{t_r, \ldots, t_r + T - 1\}$.[5] Plugging these decisions into the right-hand-side of the inequality in Lemma 3, we find that the constraints of problem (12) greatly simplify the inequality to:

$$\Delta_T(t_r) - V\sum_{\tau=t_r}^{t+T-1}\phi(\boldsymbol{\gamma}(\tau)) \leq \tilde{B}T^2 + CT - VTF_r^* \quad (31)$$

[5]If the supremum utility in (12) is not achievable, we can recover (31) simply by taking a limit over policies that approach the supremum.

where we have used the fact that $\phi(\boldsymbol{\gamma}^*) = F_r^*$, where F_r^* is the supremum utility in the problem (12).

Summing (31) over $r \in \{0, \ldots, R-1\}$ and using the definition of $\Delta_T(t_r)$ in (20) yields:

$$\begin{aligned}
L(\boldsymbol{\Theta}(RT)) - L(\boldsymbol{\Theta}(0)) - V\textstyle\sum_{\tau=0}^{RT-1}\phi(\boldsymbol{\gamma}(\tau)) \\
\leq \tilde{B}RT^2 + CRT - VT\textstyle\sum_{r=0}^{R-1}F_r^* \quad (32)
\end{aligned}$$

Rearranging terms, dividing by VRT, and using the fact that $L(\boldsymbol{\Theta}(RT)) - L(\boldsymbol{\Theta}(0)) \geq 0$ (since $Q_n^{(c)}(RT) \geq 0 = Q_n^{(c)}(0)$ and $G_m(RT) \geq G_m^{min} = G_m(0)$) yields:

$$\tfrac{1}{RT}\textstyle\sum_{\tau=0}^{RT-1}\phi(\boldsymbol{\gamma}(\tau)) \geq \tfrac{1}{R}\sum_{r=0}^{R-1}F_r^* - \tfrac{(\tilde{B}T+C)}{V}$$

However, by concavity of $\phi(\boldsymbol{\gamma})$ and Jensen's inequality, the left-hand-side of the above inequality can be modified to yield:

$$\phi(\overline{\boldsymbol{\gamma}}) \geq \tfrac{1}{R}\textstyle\sum_{r=0}^{R-1}F_r^* - \tfrac{(\tilde{B}T+C)}{V}$$

where $\overline{\boldsymbol{\gamma}}$ is defined:

$$\overline{\boldsymbol{\gamma}} \triangleq \tfrac{1}{RT}\textstyle\sum_{\tau=0}^{RT-1}\boldsymbol{\gamma}(\tau) \quad (33)$$

The result of Theorem 1 part (b) then follows by relating $\phi(\overline{\boldsymbol{\gamma}})$ to $\phi(\overline{\boldsymbol{x}})$, as described in the next lemma.

Lemma 4: Define $\overline{\boldsymbol{\gamma}}$ as a time average of $\boldsymbol{\gamma}(\tau)$ over RT slots, given by (33), and define $\overline{\boldsymbol{x}}$ as a time average of $\boldsymbol{x}(\tau)$ over the same interval. Then:

$$\phi(\overline{\boldsymbol{x}}) \geq \phi(\overline{\boldsymbol{\gamma}}) - \sum_{m=1}^{M}\frac{\nu_m^{max}(G_m(RT) - G_m(0))}{RT}$$

Assuming that $G_m(0) = G_m^{min}$ for all $m \in \{1, \ldots, M\}$, then:

$$\phi(\overline{\boldsymbol{x}}) \geq \phi(\overline{\boldsymbol{\gamma}}) - \sum_{m=1}^{M}\frac{\nu_m^{max}[V(\nu_m^{max} - \nu_m^{min}) + 2A_m^{max}]}{RT}$$

Proof: (Lemma 4) From the update rule for $G_m(t)$ in (19) we have for any slot τ:

$$G_m(\tau+1) \geq G_m(\tau) + \gamma_m(\tau) - x_m(\tau)$$

Summing the above over $\tau \in \{0, \ldots, RT-1\}$ yields:

$$G_m(RT) - G_m(0) \geq \textstyle\sum_{\tau=0}^{RT-1}\gamma_m(\tau) - \sum_{\tau=0}^{RT-1}x_m(\tau)$$

Dividing by RT yields:

$$\overline{x}_m \geq \overline{\gamma}_m - [G_m(RT) - G_m(0)]/(RT)$$

This can be written in vector form as:

$$\overline{\boldsymbol{x}} \geq \max[\overline{\boldsymbol{\gamma}} - (\boldsymbol{G}(RT) - \boldsymbol{G}(0))/(RT), \boldsymbol{0}]$$

where we define $\boldsymbol{G}(RT)$ as an M-dimensional vector with entries $G_m(RT)$, and $\boldsymbol{G}(0)$ is defined similarly. The $\max[\cdot, 0]$ in the above vector inequality is due to the fact that $\overline{x}_m \geq 0$ for all m. Because $\phi(\boldsymbol{x})$ is entrywise non-decreasing with maximum mth partial derivatives ν_m^{max}, we have:

$$\begin{aligned}
\phi(\overline{\boldsymbol{x}}) &\geq \phi(\max[\overline{\boldsymbol{\gamma}} - (\boldsymbol{G}(RT) - \boldsymbol{G}(0))/(RT), \boldsymbol{0}]) \\
&\geq \phi(\overline{\boldsymbol{\gamma}}) - \sum_{m=1}^{M}\frac{\nu_m^{max}(G_m(RT) - G_m(0))}{RT}
\end{aligned}$$

This proves the first part of the lemma. Assuming that $G_m(0) = G_m^{min}$ for all m, we know by part (a) of Theorem 1 that $G_m(RT) - G_m(0) \leq G_m^{max} - G_m^{min}$, and:

$$G_m^{max} - G_m^{min} \leq V(\nu_m^{max} - \nu_m^{min}) + 2A_m^{max}$$

□

V. General Optimization of Time Averages

We conclude with a recipe for general optimization of time averages in a universal context (see [24] for analysis against a T-slot lookahead policy). Consider a slotted time system with $\omega(t)$ being a general *random event* and $\alpha(t) \in \mathcal{A}_{\omega(t)}$ representing a general control action, which affects *system attributes* $\boldsymbol{x}(t) = (x_1(t), \ldots, x_M(t))$, $\boldsymbol{y}(t) = (y_0(t), y_1(t), \ldots, y_L(t))$ via arbitrary bounded functions $x_m(t) = \hat{x}_m(\alpha(t), \omega(t))$, $y_l(t) = \hat{y}_l(\alpha(t), \omega(t))$ for $m \in \{1, \ldots, M\}$, $l \in \{0, 1, \ldots, L\}$. There is no probability distribution for the events $\omega(t)$. The problem is to solve:

$$\begin{aligned} \text{Minimize:} \quad & \overline{y}_0 + f(\overline{\boldsymbol{x}}) \\ \text{Subject to:} \quad & \overline{y}_l + g_l(\overline{\boldsymbol{x}}) \leq 0 \;\; \forall l \in \{1, \ldots, L\} \\ & \alpha(t) \in \mathcal{A}_{\omega(t)} \;\; \forall t \geq 0 \end{aligned}$$

where $f(\boldsymbol{x})$, $g_l(\boldsymbol{x})$ are general convex functions. Assume that $x_{m,min} \leq x_m(t) \leq x_{m,max}$ for all t, for some finite constants $x_{m,min}$, $x_{m,max}$. We introduce virtual queues $Z_l(t)$ and $H_m(t)$ as follows:

$$Z_l(t+1) = \max[Z_l(t) + y_l(t) + g_l(\boldsymbol{\gamma}(t)), 0] \quad (34)$$

$$H_m(t+1) = H_m(t) + \gamma_m(t) - x_m(t) \quad (35)$$

where auxiliary variables $\boldsymbol{\gamma}(t) = (\gamma_1(t), \ldots, \gamma_M(t))$ are chosen every slot t subject to:

$$x_{m,min} \leq \gamma_m(t) \leq x_{m,max} \; \forall m \in \{1, \ldots, M\} \quad (36)$$

The virtual queues $H_m(t)$ have a different structure than $G_m(t)$ in (19) because they enforce the equality constraint $\overline{\gamma}_m = \overline{x}_m$, which is needed because $f(\boldsymbol{x})$, $g_l(\boldsymbol{x})$ are not necessarily entrywise non-decreasing or non-increasing.

Define $\boldsymbol{\Theta}(t) \triangleq [\boldsymbol{Z}(t), \boldsymbol{H}(t)]$, and define $L(\boldsymbol{\Theta}(t))$ as:

$$L(\boldsymbol{\Theta}(t)) \triangleq \frac{1}{2}\sum_{l=1}^{L} Z_l(t)^2 + \frac{1}{2}\sum_{m=1}^{M} H_m(t)^2$$

Define the 1-slot drift $\Delta_1(t)$ as before. Every slot t, the algorithm observes $\omega(t)$ and chooses $\alpha(t) \in \mathcal{A}_{\omega(t)}$ and $\boldsymbol{\gamma}(t)$ subject to (36) to minimize (to within an additive constant):

$$\Delta_1(t) + V\hat{y}_0(\alpha(t), \omega(t)) + Vf(\boldsymbol{\gamma}(t))$$

This reduces to the following simple policy. Every slot t:

- Observe $\boldsymbol{\Theta}(t)$. Choose $\boldsymbol{\gamma}(t)$ subject to (36) to minimize:
$$Vf(\boldsymbol{\gamma}(t)) + \sum_{l=1}^{L} Z_l(t) g_l(\boldsymbol{\gamma}(t)) + \sum_{m=1}^{M} H_m(t)\gamma_m(t)$$
- Observe $\boldsymbol{\Theta}(t)$, $\omega(t)$. Choose $\alpha(t) \in \mathcal{A}_{\omega(t)}$ to minimize:
$$V\hat{y}_0(\alpha(t), \omega(t)) + \sum_{l=1}^{L} Z_l(t)\hat{y}_l(\alpha(t), \omega(t)) - \sum_{m=1}^{M} H_m(t)\hat{x}_m(\alpha(t), \omega(t))$$
- Update $Z_l(t)$ and $H_m(t)$ by (34) and (35).

References

[1] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.

[2] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *Proc. IEEE INFOCOM*, March 2005.

[3] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, pp. 396-409, April 2008.

[4] M. J. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time Varying Channels*. PhD thesis, Massachusetts Institute of Technology, LIDS, 2003.

[5] L. Tassiulas. Scheduling and performance limits of networks with constantly changing topology. *IEEE Transactions on Information Theory*, vol. 43, no. 3, pp. 1067-1073, May 1997.

[6] M. J. Neely, E. Modiano, and C. E Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 89-103, January 2005.

[7] R. Urgaonkar and M. J. Neely. Opportunistic scheduling with reliability guarantees in cognitive radio networks. *IEEE Transactions on Mobile Computing*, vol. 8, no. 6, pp. 766-777, June 2009.

[8] L. Huang and M. J. Neely. The optimality of two prices: Maximizing revenue in a stochastic communication system. *IEEE/ACM Transactions on Networking*, vol. 18, no. 2, pp. 406-419, April 2010.

[9] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, vol. 50, no. 4, pp. 401-457, 2005.

[10] M. J. Neely and R. Urgaonkar. Cross layer adaptive control for wireless mesh networks. *Ad Hoc Networks (Elsevier)*, vol. 5, no. 6, pp. 719-743, August 2007.

[11] M. J. Neely, E. Modiano, and C. E. Rohrs. Tradeoffs in delay guarantees and computation complexity for $n \times n$ packet switches. *Proc. of Conf. on Information Sciences and Systems (CISS), Princeton*, March 2002.

[12] J. Andrews, S. Shakkottai, R. Heath, N. Jindal, M. Haenggi, R. Berry, D. Guo, M. Neely, S. Weber, S. Jafar, and A. Yener. Rethinking information theory for mobile ad hoc networks. *IEEE Communications Magazine*, vol. 46, no. 12, pp. 94-101, Dec. 2008.

[13] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, IT-23, no. 3, pp. 337-343, May 1977.

[14] T. M. Cover. Universal portfolios. *Mathematical Finance*, vol. 1, no. 1, pp. 1-29, Jan. 1991.

[15] T. M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, vol. 42, no. 2, 1996.

[16] N. Merhav and M. Feder. Universal schemes for sequential decision from individual data sequences. *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1280-1292, July 1993.

[17] M. J. Neely. Stock market trading via stochastic network optimization. *ArXiv Technical Report*, arXiv:0909.3891v1, Sept. 2009.

[18] M. J. Neely. Stock market trading via stochastic network optimization. *Proc. IEEE Conference on Decision and Control (CDC)*, Atlanta, GA, Dec. 2010.

[19] J. A. Garay and I. S. Gopal. Call preemption in communication networks. *Proc. IEEE INFOCOM*, vol. 44, pp. 1043-1050, Florence, Italy, 1992.

[20] S. Plotkin. Competitive routing of virtual circuits in atm networks. *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 1128-1136, Aug. 1995.

[21] L. Lin, N. B. Shroff, and R. Srikant. Asymptotically optimal power-aware routing for multihop wireless networks with renewable energy sources. *Proc. IEEE INFOCOM*, March 2005.

[22] J. J. Jaramillo and R. Srikant. Admission control and routing in multi-hop wireless networks. *Proc. IEEE Conf. on Decision and Control (CDC)*, Dec. 2008.

[23] M. Andrews. Maximizing profit in overloaded networks. *Proc. IEEE INFOCOM*, March 2005.

[24] M. J. Neely. Universal scheduling for networks with arbitrary traffic, channels, and mobility. *ArXiv technical report*, arXiv:1001.0960v1, Jan. 2010.

[25] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 2915-2934, July 2006.

[26] M. J. Neely. Optimal backpressure routing for wireless networks with multi-receiver diversity. *Proc. of Conf. on Information Sciences and Systems (CISS), Invited paper on Optimization of Communication Networks*, March 2006.