Spatial-Temporal Routing for Supporting End-to-End Hard Deadlines in Multi-hop Networks

Xin Liu and Lei Ying
School of Electrical, Computer and Energy Engineering
Arizona State University
Tempe, AZ, United States, 85287
Email: {xliu272, lei.ying.2}@asu.edu

***Abstract*—We consider the problem of routing packets with end-to-end hard deadlines in communication networks. The problem is challenging due to the complex spatial-temporal correlation among flows with different deadlines. To tackle this challenging problem, we introduce the concepts of virtual links/routes to incorporate end-to-end deadline constraints into routing and propose a novel virtual queue architecture to guide the spatial-temporal routing which specifies where and when a packet should be routed. The proposed policy can support any periodic constant traffic within the network throughput region. Our simulations further show that the policy outperforms traditional policies such as backpressure and earliest-deadline-first (EDF) for general traffic flows.**

I. INTRODUCTION

Network applications, such as emergency messages, voice calls and video streaming, demand reliable and predictable transmissions over multi-hop networks. Despite remarkable progress on the design of communication networks with maximum throughput and low delays, communications with hard deadlines in multi-hop communication networks remains a challenging problem. When end-to-end hard deadlines are imposed on packets, routing/scheduling decisions at each hop are coupled with the decisions in the sequential hops. Each link faces the dilemma of allocating its capacity to flows with large backlogs or to those whose packets have short deadlines.

This challenging problem has been studied under several settings in the literature. [1] studied communication with per-packet deadline constraints in a wired multi-hop network and proposed a novel queue structure to dynamically adjust the service discipline to meet end-to-end deadlines. However, [1] does not establish any performance guarantee of the proposed resource allocation policies. [2] proposed a throughput optimal policy for tree networks and frame-based traffic flows such that the deadlines are the end of frames. [3] presented an online algorithm with throughput guarantees for general traffic flows and networks, but the competitive ratio is inversely proportional to the hops of the longest route. More recently, [4] developed a maximal throughput scheduling policy for multi-hop wireless networks in which links do not interfere with each other and the transmit power can be adjusted. One common assumption made in these work is that each flow is associated with a fixed route, so routing packets with end-to-end deadlines has not been studied in these papers.

In this paper, we focus on routing packets with end-to-end hard deadlines in multi-hop communication networks. In particular, no route is given for any flow and a packet is allowed to use any route along which it can reach its destination before the deadline expires. We introduce the concept of *virtual links* to model the link resource in spatial and temporal domains, and the concept of *virtual routes* to incorporate end-to-end deadline constraints in routing. These two novel concepts enable us to characterize the complicated spatial-temporal correlation of packets from different flows and with different deadlines. Based on the virtual-link/route concepts, we propose a new architecture which maintains per-destination and per-time virtual queues on each node to keep track of the congestion levels of each virtual link with respect to different flows. Based on that, a spatial-temporal routing policy is developed to dynamically balance traffic flows among virtual routes. The main contributions of this paper are summarized below.

- We introduce two novel concepts of *virtual links* and *virtual routes*, which generalize the traditional concepts of links and routes to the temporal domain and enable us to explicitly characterize the network throughput region under end-to-end deadline constraints.
- We develop distributed routing, called spatial-temporal backpressure, for traffic flows with end-to-end deadlines. For periodic traffic flows with a frame-structure, which will be defined in Section II, our proposed spatial-temporal backpressure can support any such traffic that is within the network throughput region. The routing is fully distributed and only requires information exchange among neighboring nodes.
- We evaluate the performance of the proposed policy by numerical simulations. From the simulations, the proposed policy significantly outperforms backpressure and EDF for periodic traffic flows and can also support a higher delivery ratio under general traffic flows.

II. NETWORK MODEL

We consider a multi-hop communication network denoted by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{L}$ is the set of links. The network is assumed to be a time-slotted system, where arrivals occur at the beginning of a time slot and departures occur at the end of a time slot. Denote by (a, b)

the link from node a to node b and $C_{(a,b)}$ the capacity of the link (i.e., the number of packets that can be transmitted in one time slot). In this paper, we assume transmissions at different links do not interfere with each other. So the communication network is either a wired network or a wireless network in which neighboring links operate over orthogonal channels. We further assume the capacity of a link is time invariant. Packets with hard deadlines are injected into the network, where a deadline τ for a packet means that the packet has to be delivered within τ time slots to its destination after its arrival; otherwise the packet will be dropped.

We assume a frame structure for traffic flows such that every T consecutive time slots are grouped into a frame. The deadline of a packet expires within the same frame in which the packet arrives. This frame-based traffic structure has been commonly assumed in the literature [2], [5], [6]. Traditionally, in a multi-commodity flow problem, a flow is defined by its source and destination. In this paper, we define a flow using 4-tuple: its source, destination, the relative arrival time in a frame and the relative expiration time in a frame. In particular, for flow f, let $s(f)$ denote its source, $d(f)$ denote its destination, $t_b(f)$ the relative arrival time (beginning time) of the packets belonging to flow f and $t_e(f)$ the relative expiration time (end time) of the packets of flow f. So the deadline of the packets of flow f is $t_e(f) - t_b(f) + 1$. Therefore, a flow is not only defined in space (by its source and destination) and also defined over time (by the beginning and end time slots). Denote $\mathcal{F}$ by the collection of flows.

Consider a simple example in Fig. 1, where the frame size is two. Two flows are defined in the figure, where packets of flow f_1 arrive at node a at the first time slot of a frame and need to be delivered to node c at the end of the frame, and packets of f_2 arrive at node a at the first time slot of a frame and needs to be delivered to node b at the end of the frame.

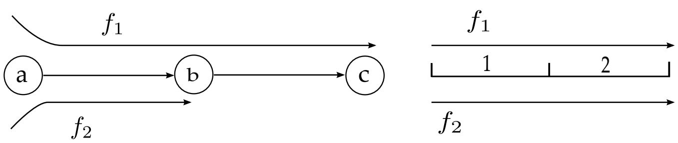

Fig. 1: Illustration of Flows

We further assume periodic constant traffic arrivals such that for each flow, the arrival pattern within each frame remains the same across frames. This models applications such as real-time surveillance systems where remote sensors collect and report data periodically.

III. Virtual Links/Routes and Spatial-Temporal Routing

Traditionally, the throughput region of a communication network is defined based on link capacity, and a set of flows are said to be within the network throughput region if there exists a resource allocation algorithm under which the long-term average throughput is equal to the arrival rate of each flow. So the traditional resource allocation in communication networks concerns the "spatial allocation" (the allocation of link capacity across flows). For packets with hard deadlines, resource allocation across the temporal domain becomes critical. In this paper, we aim at incorporating hard deadline constraints into the characterization of the throughput region. One key step is to introduce two novel ideas: virtual links and virtual routes. The concept of a virtual route expands the traditional route to the temporal domain and is defined to be an $N \times T$ matrix $\mathbf{R}$ such that $R(n,t) = 1$ denotes the route traverses node n at the tth time slot of a frame. A virtual route specifies not only which links to use but when to use them, so can be used to control the end-to-end deadline of transmitting a packet.

Like a virtual route, a virtual link expands the resource of a physical link to a spatial-temporal domain. Denote by $\{(a,b),t\}$ a virtual link, which represents link (a,b) at the tth time slot of a frame. Fig. 2 illustrates the concepts of virtual links and virtual routes using the toy example in the previous section. Since the frame size is two, each physical link is represented by two virtual links. For packets of flow f_1, it has to take the virtual route

$$\{(a,b),1\} \rightarrow \{(b,c),2\}$$

to reach destination c by the end of time slot 2. For packets belonging to flow f_2, both virtual routes $\{(a,b),1\}$ and $\{(a,b),2\}$ are feasible routes.

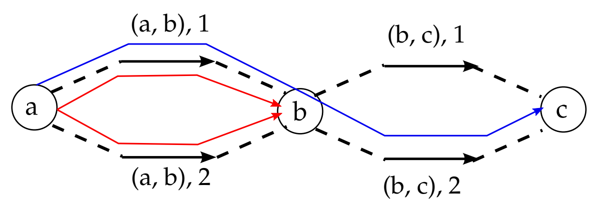

Fig. 2: Illustration of virtual routes and virtual links. Blue and red lines represent virtual routers for flows f_1 and f_2, respectively. Black lines represent virtual links.

Now to achieve the full throughput region in the network, a resource allocation algorithm needs to 1) identify the set of feasible virtual routes for each flow and 2) allocate the load appropriately on the set of routes under the network capacity constraints. Let us assume the capacity of both links are 10 packets/time slot in the toy example. We consider the following two different cases.

- Case 1: Flow f_1 has a periodic arrival of 5 packets at the first time slot of each frame; and flow f_2 has a periodic arrival of 10 packets at the first time slot of each frame. One of the feasible routing solutions is shown in Fig. 3, where flow f_2 splits its traffic evenly among the two virtual routes $\{(a,b),1\}$ and $\{(a,b),2\}$.
- Case 2: Flow f_1 has a periodic arrival of 10 packets at the first time slot of each frame; and flow f_2 has a periodic arrival of 10 packets at the first time slot of each frame. In this case, the unique routing solution to support the traffic is to have flow f_2 use virtual route $\{(a,b),2\}$ and flow

f_1 use the unique virtual route $\{(a,b),1\} \rightarrow \{(b,c),2\}$ as shown in Fig. 4.

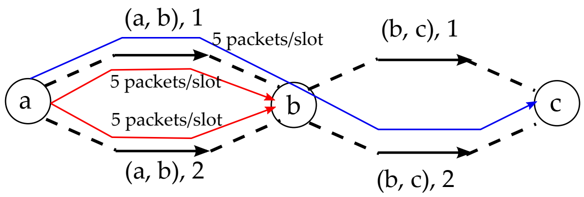

Fig. 3: Case 1

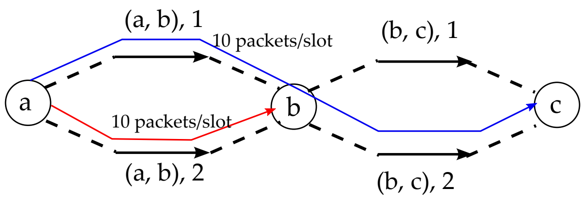

Fig. 4: Case 2

The challenge of routing packets with end-to-end hard deadline constraints is that even without introducing virtual routes, the number of possible routes in a network is exponential in the number of links. After representing each physical link with T virtual links, the number of virtual routes becomes even larger. To overcome this difficulty, we exploit the idea in [7], which includes both dynamical routing and load balancing in a communication network without requiring per route information.

One of the algorithms proposed in [7] achieves throughput optimality while guaranteeing hop constraints (a constraint on the number of hops a packet is allowed to travel before reaching its destination). A critical component of the algorithm is to maintain per-destination and per-hop queues and to only allow packets to be transmitted to neighboring nodes who can meet the hop constraints. For packets with hard deadlines, with the concept of virtual links and virtual routes, our proposed policy will maintain per-destination and per-time virtual queues, and only allow a packet to be transmitted to virtual links that guarantees the feasibility of delivering the packet before its deadline. Furthermore, our policy is completely distributed. The design of the policy and the performance analysis will be presented in the following sections.

IV. Throughput Region for communication networks with Hard Deadlines

For traffic flows with hard deadlines, the network throughput region depends on both link capacity and the distributions of traffic and deadlines. Assuming periodic constant traffic flows with a frame structure, we next characterize the network throughput region using flow conservation constraints and virtual link capacity constraints based on the necessary conditions for virtual commodities. We define $\mathcal{D}_{\{d,t_e\}} = \{f : d(f) = d, t_e(f) = t_e\}$ to be virtual commodity $\{d, t_e\}$ such that a packet of flow f in $\mathcal{D}_{\{d,t_e\}} = \{f : d(f) = d, t_e(f) = t_e\}$ has to reach the destination d by the end of the t_eth time slot in a frame. The following necessary condition has to be satisfied for each virtual commodity, each node and each time slot in a frame.

$$\sum_{f\in\mathcal{D}_{\{d,t_e\}}} a_f 1_{s(f)=n,t_b(f)=t} + \sum_{a:(a,n)\in\mathcal{L}} \sum_{i=1}^{t-1} u_{\{d,t_e\}}^{\{(a,i)\rightarrow(n,t-1)\}} = \sum_{b:(n,b)\in\mathcal{L}} \sum_{i=t}^{t_e-h_{b\rightarrow d}^{\min}} u_{\{d,t_e\}}^{\{(n,t)\rightarrow(b,i)\}}, \forall d, n\in\mathcal{N},\ t_e, t\in[1,T]. \tag{1}$$

In the condition above, $1_{s(f)=n,t_b(f)=t}=1$ indicates flow f is injected into node n at the beginning of the tth time slot and a_f is the arrival rate of flow f; $h_{b\rightarrow d}^{\min}$ denotes the minimum number of hops from node b to node d, so $t_e - h_{b\rightarrow d}^{\min}$ is the maximum number of time slots a packet can be hold at node n before transmitting to node b such that packet is still feasible to reach its destination before the deadline expires; $u_{\{d,t_e\}}^{\{(a,i)\rightarrow(n,t)\}}$ denotes the number of packets that arrive at node a at the beginning of the ith time slot and are transmitted to node n at the end of the tth time slot. The flow conservation constraint (1) basically states that the incoming packets of virtual commodity $\{d, t_e\}$ at node n at the beginning of the tth time slot should be sent to node n's neighbors in the subsequent time slots while guaranteeing the feasibility of delivering them before their deadlines.

We further have the following necessary condition due to link capacity constraint. Note that all traffics $\{(a,i) \rightarrow (n,t)\}$ uses virtual link $\{(a,n),t\}$, so we have the following capacity constraint

$$\sum_{\{d,t_e\}\in\mathcal{D}} \sum_{i=1}^{t} u_{\{d,t_e\}}^{\{(a,i)\rightarrow(n,t)\}} \leq C_{\{(a,n),t\}},\ (a,n)\in\mathcal{L},\ t\in[1,T], \tag{2}$$

where $C_{\{(a,n),t\}} = C_{(a,n)}$ for any t. Moreover, we define the link capacity region to be $\mathcal{C} = \{\mathbf{u} \mid \mathbf{u}$ satisfies (2) $\}$, where $\mathbf{u}$ is the vector version of transmission rates.

Now given traffic $\mathbf{A} = \{a_f\}_{f\in\mathcal{F}}$ with deadline constraint $\mathbf{D} = \{t_b(f), t_e(f)\}_{f\in\mathcal{F}}$, we can define throughput region as

$$\Omega = \{(\mathbf{A},\mathbf{D}) \mid \text{there exists } \mathbf{u} \text{ that satisfies both (1) and (2)}\}$$

Then we give the following definition and theorem.

Definition 1. The arrival traffic $(\mathbf{A},\mathbf{D})$ is *supportable* by a routing policy if the packet dropping rate converges to zero as $t \rightarrow \infty$.

Theorem 1. No routing policy can support an arrival traffic $(\mathbf{A},\mathbf{D}) \notin \Omega$.

Proof. This proof is deferred to our technical report [8]. □

V. Optimization Framework

After defining throughput region Ω, we consider the following optimization problem whose goal is to reduce the average latency while guaranteeing hard deadline constraints

$$\min \sum_{\{d,t_e\}\in\mathcal{D}} \sum_{a=1}^{N}\sum_{b=1}^{N}\sum_{t=1}^{T}\sum_{i=t}^{T}(i-t)U\left(u_{\{d,t_e\}}^{\{(a,t)\to(b,i)\}}\right) \tag{3a}$$

$$\text{s.t. } \mathbf{u}\in\Omega. \tag{3b}$$

In the optimization problem, $(i-t)U(\cdot)$ is the utility function of transmission rate $u_{\{d,t_e\}}^{\{(a,t)\to(b,i)\}}$. When $U(\cdot) = 0$, the optimization problem (3) is formulated for pure stability. When $U(\cdot)$ is a strictly increasing and convex function, the optimization problem (3) has a unique optimal solution. In our optimization framework, we add a weight $(i-t)$ to the utility on $u_{\{d,t_e\}}^{\{(a,t)\to(b,i)\}}$ so the policy routes more packets of virtual commodity $\{d,t_e\}$ along $\{t,i\}$ instead of $\{t,i+1\}$ whenever possible to further reduce end-to-end latency. In this paper, we only consider the case of $U(\cdot)=0$. For the case $U(\cdot)$ is a strictly increasing and convex function, the interested reader is referred to our technical report [8].

To derive a distributed policy from the optimization problem (3), we keep link constraints (2) and define $\lambda_{\{n,t,d,t_e\}}$ to be the Lagrangian multiplier associated with flow conservation constraint (1). We obtain the following partial Lagrangian function [9]:

$$\mathcal{L}(\boldsymbol{\lambda},\mathbf{u}) = \sum_{\{d,t_e\}\in\mathcal{D}}\sum_{n=1}^{N}\sum_{t=1}^{T}\lambda_{\{n,t,d,t_e\}}\Bigg(\sum_{f\in\mathcal{D}_{\{d,t_e\}}} a_f 1_{s(f)=n,t_b(f)=t} + \sum_{a:(a,n)\in\mathcal{L}}\sum_{i=1}^{t-1}u_{\{d,t_e\}}^{\{(a,i)\to(n,t-1)\}} - \sum_{b:(n,b)\in\mathcal{L}}\sum_{i=t}^{t_e-h_{b\to d}^{\min}}u_{\{d,t_e\}}^{\{(n,t)\to(b,i)\}}\Bigg).$$

Given $\boldsymbol{\lambda}$, we then consider the following problem:

$$\min\ \mathcal{L}(\boldsymbol{\lambda},\mathbf{u}) \tag{4a}$$

$$\text{s.t. } \mathbf{u}\in\mathcal{C}. \tag{4b}$$

We decouple objective (4a) into individual optimization problems for each link (a,b) as follows (the details can be found in our technical report [8])

$$\max \sum_{\{d,t_e\}\in\mathcal{D}}\sum_{t=1}^{t_e-h_{b\to d}^{\min}}\sum_{i=1}^{t}\left(\lambda_{\{a,i,d,t_e\}}-\lambda_{\{b,t+1,d,t_e\}}\right)u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}} \tag{5a}$$

$$\text{s.t. } \sum_{\{d,t_e\}\in\mathcal{D}}\sum_{i=1}^{t}u_{\{d,t_e\}}^{\{(a,i)\to(n,t)\}} \le C_{\{(a,n),t\}},\quad \forall\ (a,b)\in\mathcal{L},\ t\in[1,T]. \tag{5b}$$

Further, (5) is decoupled into sub-optimization problems at each time slot in a frame

$$\max \sum_{\{d,t_e\}\in\mathcal{D}}\sum_{i=1}^{t}\left(\lambda_{\{a,i,d,t_e\}}-\lambda_{\{b,t+1,d,t_e\}}\right)u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}} \tag{6a}$$

$$\text{s.t. } \sum_{\{d,t_e\}\in\mathcal{D}}\sum_{i=1}^{t}u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}} \le C_{\{(a,b),t\}},\quad \forall\ (a,b)\in\mathcal{L},\ t\in[1,T]. \tag{6b}$$

Now the original problem (3) has been decoupled into individual sub-optimizaiton problem (6) for each link at each time slot in a frame. Based on the reduced formulation (6), we then develop a fully distributed routing policy.

VI. Spatial-Temporal Routing

Motivated by the connection between Lagrangian dual variables and queue lengths [10], we first introduce the virtual queue architecture.

Let $d_{\{n,t,d,t_e\}}[k]$ denote virtual queue associated with the Lagrangian dual variable $\lambda_{\{n,d,t_e,t\}}$ which updates as

$$d_{\{n,t,d,t_e\}}[k+1] = \Bigg(d_{\{n,t,d,t_e\}}[k] + \sum_{f\in\mathcal{D}_{\{d,t_e\}}} a_f 1_{s(f)=n,t_b(f)=t} + \sum_{a:(a,n)\in\mathcal{L}}\sum_{i=1}^{t-1}u_{\{d,t_e\}}^{\{(a,i)\to(n,t-1)\}}[k] - \sum_{b:(n,b)\in\mathcal{L}}\sum_{i=t}^{t_e-h_{b\to d}^{\min}}u_{\{d,t_e\}}^{\{(n,t)\to(b,i)\}}[k]\Bigg)^{+}. \tag{7}$$

Note that $d_{\{n,t,d,t_e\}}[k]$ measures the congestion level of transmitting the commodity $\{d,t_e\}$ at node n at the tth time slot during frame k. We remark that k is the frame index, so the virtual queues are updated once every frame (instead of every time slot). If $d_{\{n,t,d,t_e\}}[k]$ is large, packets belonging to $\{d,t_e\}$ are less likely to get through if routed to node n at the tth time slot, so should be routed to a different time slot t in the frame. As convention, we assume $d_{\{n,\cdot,n,\cdot\}}[k]=0$ for $\forall n\in\mathcal{N}$ since packets will be moved to the upper layer after arriving at their destinations. We still use a toy example to illustrate the virtual queue architecture in Fig. 5, where packets of commodities $\{c,2\}$ (blue) and $\{b,2\}$ (red) are injected into the common source a at the beginning of the first time slot, so node a maintains virtual queue $d_{\{a,1,b,2\}}$ for $\{b,2\}$ and $d_{\{a,1,c,2\}}$ for $\{c,2\}$. The virtual queues are counters and do not hold real packets. Node b maintains $d_{\{b,2,b,2\}}$ and $d_{\{b,3,b,2\}}$ for $\{b,2\}$ and $d_{\{b,2,c,2\}}$ for $\{c,2\}$ since for $\{b,2\}$, packets can be routed along either $d_{\{a,1,b,2\}}\to d_{\{b,2,b,2\}}$ or $d_{\{a,1,b,2\}}\to d_{\{b,3,b,2\}}$ and for $\{c,2\}$, $d_{\{a,1,c,2\}}\to d_{\{b,2,c,2\}}$ is the only feasible route in its first hop. Node c maintains $d_{\{c,3,c,2\}}$ for the commodity $\{c,2\}$. As defined, $d_{\{b,2,b,2\}}$, $d_{\{b,3,b,2\}}$ and $d_{\{c,3,c,2\}}$ are always 0.

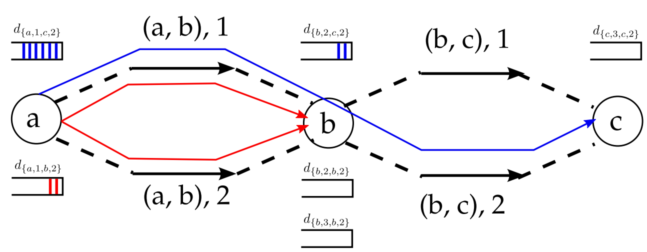

Fig. 5: Illustration of the Virtual Queue Architecture

Based on the virtual queue architecture, we propose the following spatial-temporal routing policy, as presented in Policy 1.

Spatial-temporal backpressure $w_{\{d,t_e\}}^{\{(a,i)\to(b,t+1)\}}[k]$ is the difference between the values of virtual queues of (a,i) and $(b,t+1)$ for commodity $\{d,t_e\}$, which extends the traditional backpressure to the spatial-temporal domain. In policy 1, virtual link $\{(a,b),t\}$ increases the capacity allocation to the tuple $\{d,t_e,i\}^*$ with the maximum backpressure $w_{\{d,t_e\}^*}^{\{(a,i^*)\to(b,t+1)\}}$ (breaking ties arbitrarily). If the number of real packets of $\{d,t_e\}$ is less than $\bar{u}_{\{d,t_e\}}^{\{(a,b),t\}}[k]$, then all real packets are transmitted. Again, we consider the example in Fig. 5. Suppose $a_{f_1} = a_{f_2} = 10$, $d_{\{a,1,c,2\}}[k] = 30$, $d_{\{b,2,c,2\}}[k] = 10$ and $d_{\{a,1,b,2\}}[k] = 10$, then $w_{\{c,2\}}^{\{(a,1)\to(b,2)\}}[k] = 20$, $w_{\{b,2\}}^{\{(a,1)\to(b,2)\}}[k] = 10$ and $w_{\{b,2\}}^{\{(a,1)\to(b,3)\}}[k] = 10$. During the first time slot at the kth frame, as $w_{\{c,2\}}^{\{(a,1)\to(b,2)\}}[k] > w_{\{b,2\}}^{\{(a,1)\to(b,2)\}}[k]$, virtual link $\{(a,b),1\}$ will increases its capacity allocation to the commodity $\{c,2\}$. Similarly in the second time slot, virtual link $\{(a,b),2\}$ and $\{(b,c),2\}$ will increases its capacity allocation to $\{b,2\}$ and $\{c,2\}$, respectively. Then at the beginning of the $k+1$th frame, the values of virtual queues are updated to be $d_{\{a,1,c,2\}}[k+1] = 30$, $d_{\{b,2,c,2\}}[k+1] = 10$ and $d_{\{a,1,b,2\}}[k+1] = 10$, which remain unchanged in the kth frame and real transmission rates $\bar{\mathbf{u}}$ converges to a feasible solution for $a_{f_1} = a_{f_2} = 10$ as shown in Fig. 4.

Policy 1 Spatial-Temporal Backpressure Policy

1: For virtual link $\{(a,b),t\}$ at the beginning of the kth frame, calculate spatial-temporal backpressure

$$w_{\{d,t_e\}}^{\{(a,i)\to(b,t+1)\}}[k] = \Big(d_{\{a,i,d,t_e\}}[k] - d_{\{b,t+1,d,t_e\}}[k]\Big)^+$$

2: Compute virtual transmission rate at the beginning of the kth frame

$$u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}[k] = \arg\max_{\mathbf{u}\in\mathcal{C}} \sum_{\{d,t_e\}\in\mathcal{D}} \sum_i w_{\{d,t_e\}}^{\{(a,i)\to(b,t+1)\}}[k] \cdot u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}$$

3: Compute real transmission rate for $\{d,t_e\}$ on virtual link $\{(a,b),t\}$ by

$$\bar{u}_{\{d,t_e\}}^{\{(a,b),t\}}[k] = \sum_{i=1}^{t} \bar{u}_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}[k],$$

where $\bar{u}_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}[k]$ is the time average of the virtual transmission rate up to the kth frame, i.e.,

$$\bar{u}_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}[k] = \frac{1}{k}\sum_{i=1}^{k} u_{\{d,t_e\}}^{\{(a,i)\to(b,t)\}}[i].$$

4: Transmit packets of virtual commodity $\{d,t_e\}$ at node a with rate $\bar{u}_{\{d,t_e\}}^{\{(a,b),t\}}[k]$ to node b at the tth time slot of the kth frame.

5: Update virtual queues at the end of the kth frame as (7).

Next, we present the main theorem of this paper.

Theorem 2. Given any traffic with end-to-end deadline constraints such that $\big((1+\varepsilon)\mathbf{A},\mathbf{D}\big) \in \Omega$ for some $\varepsilon > 0$, $\bar{\mathbf{u}}[k]$ generated by policy 1 converges to $\bar{\mathbf{u}}^*$ such that $\bar{\mathbf{u}}^* \in \Omega$ and $\bar{\mathbf{u}}^*$ satisfies both conditions (1) and (2) for the given traffic.

Proof. This proof is deferred to our technical report [8]. □

Spatial-temporal backpressure policy will result in a feasible routing solution whenever possible. The policy exploits all feasible virtual routes. As a result, packets might traverse routes with unnecessary large end-to-end delay. To further improve delay performance, we propose a water-filling policy to reduce end-to-end latency while guaranteeing throughput optimality. The algorithm and the performance analysis can be found in our technical report in [8].

VII. Simulations

In this section, we evaluate the performance of the proposed policy via simulations. We consider a network topology shown in Fig. 6, where all links have the same capacity of $C = 10$ packets/time slot and we have three flows in the network. Flow 1 is from node 1 to 6, flow 2 is from node 2 to 7, and flow 3 is from node 3 to 5. We call our spatial-temporal policy "ST Policy" and compare its performance with the backpressure and EDF policies. In our simulations, EDF adopts randomized routing to route each packet to a randomized chosen neighbor, and backpressure adopts random packet scheduling for the selected commodity.

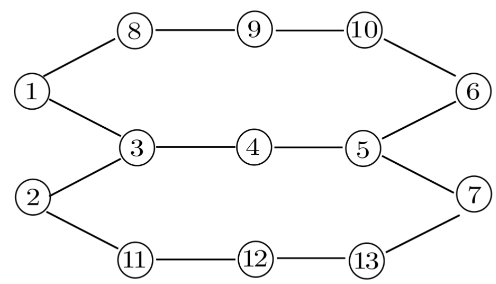

Fig. 6: Network Topology

Periodic constant traffics: We assume frame size $T = 6$. Packets of flow 1 and flow 2 arrive at the beginning of the frame and expire at the end of the frame, Packets of flow 3 arrive at time slot 2 and expire at time slot 4 of a frame. Packet arrivals per frame increase from $(20,10,0)$ to $(40,30,20)$ with the step size $(2,2,2)$. We observe in Fig. 7 that our proposed spatial-temporal policy outperforms both backpressure and EDF at all incoming rates. EDF has a higher throughput than backpressure at the low traffic regime and backpressure outperforms EDF at the high traffic regime. It can also be easily verified that (40, 30, 20) is at the boundary of the network capacity given the periodic constant traffic pattern. The figure shows the ST policy can support (40,30,20) as we proved in the main theorem.

General traffic: We consider a traffic pattern such that each of three flows has packet arrivals at every time slot. Each

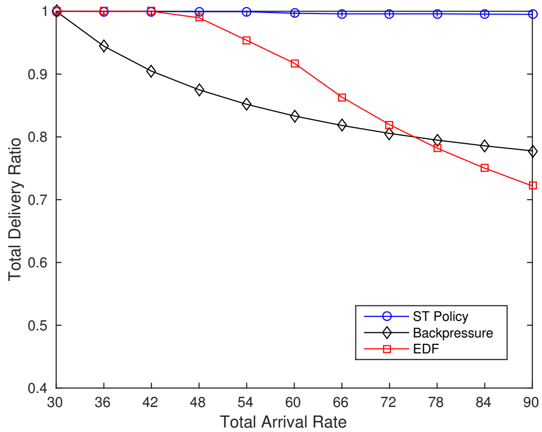

Fig. 7: Performance of different routing policies for a periodic constant traffic pattern

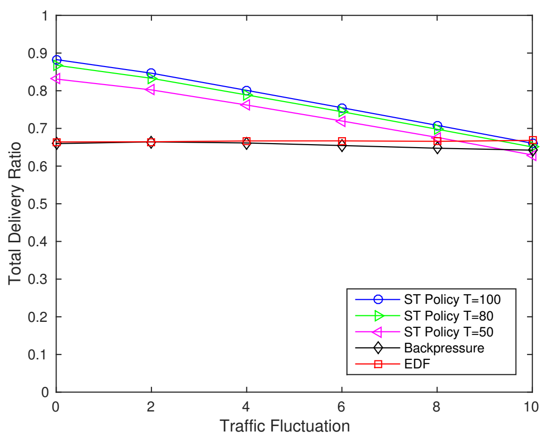

Fig. 8: Performance of different routing policies for a general traffic pattern

packet has a deadline of 6 time slots. The number of packet arrivals for each flow at each time slot is a random variable that takes integer values from $10 - i$ to $10 + i$, equally likely. We varied i from 0 to 10. A larger value of i implies larger variance of the incoming traffic. We impose virtual frame structure on this general traffic pattern and evaluated the ST policy as the frame size increases, $T = 50$, 80, and 100. From Fig. 8, we can see that the delivery ratio of the ST policy increases as frame T increases. Furthermore, the ST policy outperforms both backpressure and EDF for almost all values of i even when T is small. The delivery ratio increases as the variance of the incoming traffic decreases. This is intuitive since it is more difficult to guarantee end-to-end deadlines when the randomness in the traffic increases.

VIII. Conclusions

In this paper, we developed spatial-temporal routing for supporting end-to-end hard deadlines in communication networks. We first introduced the concepts of virtual links and virtual routes by modeling each physical link as multiple virtual links across time slots for each time frame, by which we were able to characterize the deadline constrained network throughput region with some additional assumption on the incoming traffic. Then we proposed a virtual queue architecture and developed a spatial-temporal routing policy, called spatial-temporal backpressure. Our policy can support any periodic constant traffic within the network throughput region, and for the general traffic patterns, numerical simulations verified that the proposed policy outperform existing routing policies.

Acknowledgment

This work was supported in part by the U.S. Office of Naval Research (ONR Grant No. N00014-15-1-2169) and by the National Science Foundation (NSF) under Grants CNS-1264012 and CNS-1262329.

References

[1] R. Li and A. Eryilmaz, "Scheduling for end-to-end deadline-constrained traffic with reliability requirements in multihop networks," *IEEE/ACM Trans. Netw.*, vol. 20, no. 5, pp. 1649–1662, Oct 2012.

[2] I.-H. Hou, "Packet scheduling for real-time surveillance in multihop wireless sensor networks with lossy channels," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 1071–1079, Feb 2015.

[3] Z. Mao, C. E. Koksal, and N. B. Shroff, "Optimal online scheduling with arbitrary hard deadlines in multihop communication networks," *IEEE/ACM Trans. Netw.*, vol. PP, no. 99, Nov 2014.

[4] R. Singh and P. R. Kumar, "Decentralized throughput maximizing policies for deadline-constrained wireless networks," *arXiv preprint arXiv:1509.05293*, 2015.

[5] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," in *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*, Rio de Janeiro, Brazil, 2009, pp. 486–494.

[6] J. J. Jaramillo and R. Srikant, "Optimal scheduling for fair resource allocation in ad hoc networks with elastic and inelastic traffic," *IEEE/ACM Trans. Netw.*, vol. 19, pp. 1125–1136, Aug. 2011.

[7] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, "On combining shortest-path and back-pressure routing over multihop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 841–854, June 2011.

[8] X. Liu and L. Ying, "Spatial-temporal routing for supporting end-to-end deadlines in multi-hop networks," Arizona State University, Tech. Rep., Jan 2016.

[9] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY: Cambridge Unversity Press, 2004.

[10] R. Srikant and L. Ying, *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, 2014.