

A General Class of Throughput Optimal Routing Policies in Multi-hop Wireless Networks

M. Naghshvar, H. Zhuang, and T. Javidi

Department of Electrical and Computer Engineering,

University of California San Diego,

La Jolla, CA 92093 USA,

{naghshvar, hzhuang, tjavidi}@ucsd.edu

Abstract

This paper considers the problem of routing packets across a multi-hop wireless network while ensuring throughput optimality. One of the main challenges in the design of throughput optimal routing policies is identifying appropriate and universal Lyapunov functions with negative expected drift. The few well-known throughput optimal routing policies in the literature are constructed using simple quadratic or exponential Lyapunov functions of the queue backlogs and as such they do not use any metric of closeness to the destination. Consequently, these routing policies exhibit poor delay performance under many network topologies and traffic conditions.

By considering a class of continuous, differentiable, and piece-wise quadratic Lyapunov functions, this paper provides a large class of throughput optimal routing policies. The proposed class of Lyapunov functions allow for the routing policies to control the traffic along short paths for a large portion of state-space while ensuring a negative expected drift, hence, enabling the design of routing policies with much improved delay performances. In particular, and in addition to recovering the throughput optimality of the well known backpressure routing policy, an opportunistic routing policy with congestion diversity is proved to be throughput optimal.

I. INTRODUCTION

Opportunistic routing for multi-hop wireless ad-hoc networks has seen recent research interest to overcome deficiencies of conventional routing [1]–[5]. Opportunistic routing mitigates the impact of poor wireless links by exploiting the broadcast nature of wireless transmissions and the path diversity. More precisely, the routing decisions are made in an online manner by choosing the next relay based on the actual transmission outcomes as well as a rank ordering of neighboring nodes. The authors in [5] provided a Markov decision theoretic formulation for opportunistic routing. In particular, it is shown that for any given packet and at any relaying epoch, the optimal routing decision, in the sense of minimum cost or hop-count, is to select the next relay node based on an index. This index is equal to the expected cost or hop-count of relaying the packet along the least costly or the shortest feasible path to the destination. Furthermore, this index is computable in a distributed manner and with low complexity using a time-invariant probabilistic description of wireless links and the time-invariant transmission costs or transmission times. As such, [5] provides a unifying framework for almost all versions of opportunistic

routing [1]–[3], where the variations are due to the authors’ choices of costs; e.g. for ExOR [3], the cost to be minimized is the expected hop-counts (ETX).

When multiple streams of packets are to traverse the network, however, it might be necessary to route some packets along longer paths, if these paths eventually lead to links that are less congested. More precisely, and as noted in [6], [7], the above opportunistic routing schemes can potentially cause severe congestion and unbounded delays (see examples given in [7]). In other words, these routing schemes are said to fail to stabilize otherwise stabilizable traffic. In contrast, it is known that a simple routing policy, known as backpressure [8], ensures bounded expected total backlog for all stabilizable arrival rates. Most interestingly, this routing policy provides throughput optimality without knowledge of the network topology or the traffic rates. In the opportunistic context, diversity backpressure routing (DIVBAR) [6] provides an opportunistic generalization of backpressure which incorporates the wireless local transmission diversity.

Note that to ensure throughput optimality, backpressure-based algorithms [6], [8] do something very different from [1]–[5]; rather than any metric of closeness to the destination (or cost), they choose the receiver with the largest positive differential queue backlog (routing responsibility is retained by the transmitter if no such receiver exists). This very property of ignoring the cost to the destination, however, becomes the bane of this approach, leading to poor delay performance (see [7], [9]). In [7], the authors proposed a routing policy, known as Opportunistic Routing with Congestion Diversity (ORCD) with an improved delay performance. ORCD combines the congestion information with the shortest path calculations inherent in opportunistic routing [7]. The throughput optimality of ORCD was conjectured in [7] but was left unproven, due to the difficulty of identifying appropriate (and universal) Lyapunov functions with negative expected drift. In fact backpressure [8] and its variants [6], [10], [11] with quadratic Lyapunov function, and randomized strategies [12] with an exponential Lyapunov function remain to be the only known throughput optimal routing policies. The strict schur-convex structure of these Lyapunov functions, however, are such that their negative drift is ensured only at the cost of potentially large delays. In this paper, we provide a large class of throughput optimal policies by considering a class of piece-wise quadratic Lyapunov functions. The proposed class of Lyapunov functions allow for the routing policies to control the traffic along short paths for a large portion of state-space while ensuring a negative expected drift, hence, enabling the design of routing policies without many of the deficiencies of backpressure-based algorithms. We also specialize our result to recover the throughput optimality of two known routing policies, backpressure (already known to be throughput optimal) and ORCD (whose throughput optimality only was conjectured in [7]).

In this paper we assume each network node transmits over an orthogonal channel, so that there is no inter-channel interference. Furthermore, we assume that the network topology as well as probability of successful transmissions are fixed. These assumptions allow for a clear presentation of the routing problem and illuminate the main concepts in their simplest forms. However, we emphasize that the generalization to the networks with inter-channel interference and variable rate and power options follow easily as shown in [6]. In [6], the price of this generalization is shown to be the centralization of the routing/scheduling globally across the network or a constant factor performance loss of the distributed variants. The generalization to the case of 1) multi-destination scenario and 2) ergodic time-varying network topology and transmission probabilities are believed to be also straight forward but remain as future areas of work.

We close this section with a note on the notations used. Let $[x]^+ = \max\{x, 0\}$. The indicator function $\mathbf{1}_{\{X\}}$ takes the value 1 whenever event X occurs, and 0 otherwise. For any set S , $|S|$ denotes the cardinality of S , while for any vector v , $\|v\|$ denotes the euclidean norm of v . For any set S , $int(S)$ is the set of all interior points of S . When dealing with

a sequence of sets C_1, C_2, \dots , we define $C^i = \cup_{j=1}^i C_j$. Lastly, we use bold letters to discriminate vectors from scalar quantities as well as their components.

II. PROBLEM FORMULATION AND OVERVIEW OF THE RESULTS

A. Problem Setup

We consider a time slotted system with slots indexed by $t \in \{0, 1, 2, \dots\}$ where slot t refers to the time interval $[t, t+1)$. There are $N+1$ nodes in the network labeled by $\Omega := \{0, 1, \dots, N\}$, where node 0 is assumed to be the destination.

Let random variable $A_i(t)$ represent the amount of data that exogenously arrives to node i during time slot t . Arrivals are assumed to be i.i.d. over time and bounded by a constant A_{max} . Let $\lambda_i = \mathbb{E}[A_i(t)]$ denote the exogenous arrival rate to node i . We define $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_N]$ to be the arrival rate vector. We assume packets that arrive exogenously at node i as well as packets routed to node i from other nodes are queued at node i in a buffer with infinite queuing space. Let $Q_i(t)$ denote the queue backlog of node i at time slot t . We assume any data that is successfully delivered to the destination will exit the network and hence, $Q_0(t) = 0$ for all time slots t . We define $\mathbf{Q}(t) = [Q_1(t), Q_2(t), \dots, Q_N(t)]$ to be the vector of queue backlogs of nodes $1, 2, \dots, N$.

We assume each node transmits at most one packet during a single time slot. Let $S_i(t)$ represent the (random) set of nodes that have received the packet transmitted by node i at time slot t . We refer to $S_i(t)$ as the set of *potential forwarders* for node i . Each node that receives a packet, sends an acknowledgment over a reliable control channel. Hence, we assume that node i has perfect knowledge of $S_i(t)$. Due to a perfect recall at any node i , we assume $i \in S_i(t)$ for all time t . We characterize the behavior of the wireless channel using a probabilistic *local broadcast model* [5]. The local broadcast model is defined using conditional probabilities $P(S|i) := \text{Prob}(\{S_i(t) = S \text{ when } i \text{ transmits a packet at time } t\})$, $S \subseteq \Omega$, $i \in \Omega$. Note that, by definition, for all $S \neq S'$, successful reception at S and S' are mutually exclusive and $\sum_{S \subseteq \Omega} P(S|i) = 1$. Node j is said to be *reachable* by node i (we write $i \rightarrow j$), if there exists a set of nodes $S \subseteq \Omega$ such that $j \in S$ and $P(S|i) > 0$.

We define a *routing decision* $\mu_{ij}(t)$ to be the number of packets whose relaying responsibility is shifted from node i to node j during time slot t . Note that $\mu_{ij}(t)$ forms the departure process from node i , while it is an element of the endogenous arrival to node j , and hence,

$$\mu_{ij}(t) \in \{0, 1\}, \mu_{ij}(t) \leq \mathbf{1}_{\{j \in S_i(t)\}}, \sum_{j=0}^N \mu_{ij}(t) \leq 1. \quad (1)$$

If $\mu_{ii}(t) = 1$, then node i retains the packet for future retransmissions. Without loss of generality, we assume that after a packet is successfully received at the destination, the packet would not be (re)transmitted by any other node, i.e. $\mu_{i0}(t) = 1$ if $0 \in S_i(t)$.

For a set C of nodes, we define $A_C(t) = \sum_{i \in C} A_i(t)$, $Q_C(t) = \sum_{i \in C} Q_i(t)$, $\mu_{C,in}(t) = \sum_{j \notin C} \sum_{k \in C} \mu_{jk}(t)$, and $\mu_{C,out}(t) = \sum_{j \in C} \sum_{k \notin C} \mu_{jk}(t)$.

The selection of routing decisions together with the exogenous arrivals impact the queue backlog of node i , $i \in \Omega$, in the following manner:

$$Q_i(t+1) = [Q_i(t) - \sum_{j \in \Omega} \mu_{ij}(t)]^+ + \sum_{j \in \Omega} \mu_{ji}(t) \mathbf{1}_{\{Q_j(t) \geq \mu_{ji}(t)\}} + A_i(t). \quad (2)$$

Definition. A *routing policy* is a collection of routing decisions $\cup_{i,j \in \Omega} \cup_{t=0}^{\infty} \{\mu_{ij}(t)\}$ where for all $i, j \in \Omega$ and $\theta \in \{0, 1\}$, the decision $\{\mu_{ij}(t) = \theta\}$ belongs to the σ -field generated by $\cup_{i,j \in \Omega} \{Q_i(0), S_i(0), \mu_{ij}(0), \dots, Q_i(t-1), S_i(t-1), \mu_{ij}(t-1), Q_i(t), S_i(t)\}$.

Definition. A routing policy Π is said to *stabilize* the network if the time average queue backlog of all nodes remain finite when packets are routed according to Π . The *stability region* of the network is the set of all arrival rate vectors λ for which there exists a routing policy that stabilizes the network.

Definition. A routing policy is said to be *throughput optimal* if it stabilizes the network for all arrival rate vectors that belong to the interior of the stability region.

Fact 1. ([6, Corollary 1]) Let \mathfrak{S} denote the stability region of the network. An arrival rate vector λ is within the stability region \mathfrak{S} if and only if there exists a stationary randomized routing policy that makes routing decisions $\{\tilde{\mu}_{ij}(t)\}_{i,j \in \Omega}$, solely based on the collection of potential forwarders at time t , $\{S_i(t)\}_{i \in \Omega}$, and for which

$$\mathbb{E} \left[\sum_j \tilde{\mu}_{kj}(t) - \sum_i \tilde{\mu}_{ik}(t) \right] \geq \lambda_k.$$

In this paper we are interested in a class of routing policies which are throughput optimal but do not require knowledge of arrival rates.

B. Priority-Based Routing

In this section we introduce the class of *priority-based* routing policies. To define the priority-based routing policy we need the following definitions.

A *rank ordering* $R = (C_1, C_2, \dots, C_M)$ is an ordered list of non-empty sets C_1, C_2, \dots, C_M ($1 \leq M \leq N$), referred to as *ranking classes*, that make up a partition of $\{1, 2, \dots, N\}$ (all nodes except the destination node), i.e. $\cup_{i=1}^M C_i = \{1, 2, \dots, N\}$ and $C_i \cap C_j = \emptyset$, $i \neq j$. We denote the set of all possible rank orderings of $\{1, 2, \dots, N\}$ by \mathcal{R} . Note that when C_i 's are singleton, R reduces to a simple permutation of the nodes $\{1, 2, \dots, N\}$. Given a rank ordering $R = (C_1, C_2, \dots, C_M)$, we write $a \prec^R b$ to indicate that node $a \in C_i$ has a lower rank than $b \in C_j$, $i < j$. And we write $a \preceq^R b$, if $a \prec^R b$ or $a, b \in C_i$ for some i .

Definition. A *priority-based* routing policy $\Pi_{\{R(t)\}}$ is a routing policy under which node i , at time t and among its set of potential forwarders $S_i(t)$, selects a node with the lowest rank according to $R(t)$. In other words, under $\Pi_{\{R(t)\}}$, $\mu_{ij}(t) = 1$, only when $j \in S_i(t)$ and $j \preceq^{R(t)} k$ for all $k \in S_i(t)$.

Next we give a few definitions which allow us to compare rank orderings R and R' :

Definition. Let $R = (C_1, C_2, \dots, C_M)$ and $R' = (C'_1, C'_2, \dots, C'_M)$. We define a *mismatch* $m : \mathcal{R} \times \mathcal{R} \rightarrow \mathbb{N}$ as

$$m(R, R') = \min \{i \in \mathbb{N} : C_i \neq C'_i\}.$$

For two rank orderings R and R' , $m(R, R')$ compares ranking classes of R and R' from low to high and determines the index of the first ranking class in which they differ.

Definition. Given two rank orderings R and R' , we say R' is a *refinement* of R (and R is a *confinement* of R') if $i \prec^R j$ implies that $i \prec^{R'} j$ for any $i, j \in \Omega$.

Definition. Given two rank orderings $R = (C_1, C_2, \dots, C_M)$ and $R' = (C'_1, C'_2, \dots, C'_{M+1})$, we say R' is a *one-step refinement* of R (and R is a *one-step confinement* of R') with regard to ranking class C_i ($1 \leq i \leq M$) if

$$\begin{cases} C_k = C'_k & \text{if } 1 \leq k \leq i-1 \\ C_i = C'_i \cup C'_{i+1} & \\ C_k = C'_{k+1} & \text{if } i+1 \leq k \leq M \end{cases} .$$

The union of the sets of all one-step refinements and one-step confinements of R , denoted by $\mathcal{B}_1(R)$ and $\mathcal{B}_2(R)$ respectively, is referred to as *adjacency* of R and is denoted by $\mathcal{A}(R)$.

Next we introduce a class of priority-based routing policies under which $R(t)$ is chosen as a time-invariant function of $\mathbf{Q}(t)$, i.e. there exists a function $\pi : \mathbb{R}_+^N \rightarrow \mathcal{R}$ such that $R(t) = \pi(\mathbf{Q}(t))$. In section II-E we proceed to establish the throughput optimality of this class of routing policies.

C. f -policy

In this section we introduce a class of priority-based routing policies each of which is associated with a bivariate function f , hence referred to as an f -policy. Each such policy partitions the space of queue backlogs, \mathbb{R}_+^N , into $|\mathcal{R}|$ routing decision cones. The specific shape of each cone (and the set of its defining hyperplanes) is dictated by the corresponding function f and to each cone a unique rank ordering of nodes $R \in \mathcal{R}$ is assigned. Now define the mapping $\pi_f : \mathbb{R}_+^N \rightarrow \mathcal{R}$ such that at any time t and for all $\mathbf{Q}(t)$ in the cone associated with R , $\pi_f(\mathbf{Q}(t)) = R$. In order to give the precise description of f -policy, we need the following definitions.

Definition. Given a bivariate function f , a *penalty* function Λ_f is defined on backlog vector $\mathbf{Q} \in \mathbb{R}_+^N$, rank ordering $R = (C_1, C_2, \dots, C_M) \in \mathcal{R}$, and natural number n , $n \leq M$:

$$\Lambda_f(\mathbf{Q}, R, n) = \sum_{i=1}^n f(|C^{i-1}|, |C_i|) Q_{C_i}.$$

Definition. Consider two rank orderings R and R' and a bivariate function f . We say R *penalizes* \mathbf{Q} *less than* R' and write $R <_{\mathbf{Q}} R'$ if

- $\Lambda_f(\mathbf{Q}, R, m(R, R')) < \Lambda_f(\mathbf{Q}, R', m(R, R'))$
- or if
- $\Lambda_f(\mathbf{Q}, R, m(R, R')) = \Lambda_f(\mathbf{Q}, R', m(R, R'))$ and R is a one-step refinement of R' .

Let $D_f(R)$, $R \in \mathcal{R}$, be a subset of \mathbb{R}_+^N such that for all $\mathbf{Q} \in D_f(R)$ and all $R' \in \mathcal{A}(R)$, $R <_{\mathbf{Q}} R'$, i.e.

$$D_f(R) = \{ \mathbf{Q} \in \mathbb{R}_+^N : R <_{\mathbf{Q}} R' \text{ for all } R' \in \mathcal{A}(R) \} . \quad (3)$$

Remark Let R and R' be two rank orderings and let $\eta \in \mathbb{R}^+$ be a constant. If $R <_{\mathbf{Q}} R'$ then $R <_{\eta \mathbf{Q}} R'$. In other words, $D_f(R)$ is a cone in \mathbb{R}_+^N .

Remark Due to the linearity of $\Lambda_f(\cdot, R, n)$ and finiteness of $\mathcal{A}(R)$, the boundaries of the cone corresponding to rank ordering R consists of finitely many hyperplanes of the form

$$\Lambda_f(\mathbf{Q}, R, m(R, R')) = \Lambda_f(\mathbf{Q}, R', m(R, R')),$$

where $R' \in \mathcal{A}(R)$.

Lemma 1. Let bivariate function f satisfy the following two conditions

- (C1) For all $m \geq 0$ and $n_1, n_2 > 0$

$$\frac{1}{f(m, n_1 + n_2)} = \frac{1}{f(m, n_1)} + \frac{1}{f(m + n_1, n_2)}.$$

- (C2) For all $m \geq 0$ and $n_1, n_2 > 0$

$$f(m, n_1) \geq f(m + n_1, n_2).$$


Then for any $\mathbf{Q} \in \mathbb{R}_+^N$, there exists a unique $R \in \mathcal{R}$ such that $\mathbf{Q} \in D_f(R)$.

Remark By Lemma 1, $\{D_f(R)\}_{R \in \mathcal{R}}$ forms a partition of \mathbb{R}_+^N . Hence, it is meaningful to define a function $\pi_f : \mathbb{R}_+^N \rightarrow \mathcal{R}$ such that $\pi_f(\mathbf{Q}) = R \Leftrightarrow \mathbf{Q} \in D_f(R)$.


Now we are ready to provide the precise definition of f -policy as discussed earlier.

Definition. f -policy is a priority-based routing policy $\Pi_{\{R(t)\}}$ where $R(t) = \pi_f(\mathbf{Q}(t))$.

Example 1. Consider a network of three nodes as given in Fig. 1(a). Let \mathcal{R} be the set of all rank orderings of $\{1, 2\}$, and $f(m, n) = \frac{1}{3^m(3^n - 1)}$. Fig. 1(b) shows the structure of the cones $\{D_f(R)\}_{R \in \mathcal{R}}$.



(a) A network of three nodes



(b) Structure of the cones

Fig. 1. Structure of the cones for a network of three nodes

Example 2. Consider a network of four nodes as given in Fig. 2(a). Let \mathcal{R} be the set of all rank orderings of $\{1, 2, 3\}$, and $f(m, n) = \frac{1}{3^m(3^n - 1)}$. Fig. 2(b) shows the structure of the cones $\{D_f(R)\}_{R \in \mathcal{R}}$.

By construction, f -policy orders the nodes based only on their queue backlogs using a bivariate function f independently of the topological characteristic of the network. In certain cases, this may cause packets to be routed away from the destination




Fig. 2. Structure of the cones for a network of four nodes

and hence, results in poor delay performance. In the next section, we introduce a modified version of f -policy, referred to as *critical f -policy*, with an improved delay performance. The main idea behind critical f -policy is that the rank orderings are limited to those under which there exists a path from any node i to the destination through the nodes with lower or the same rank as i . The precise description of critical f -policy is provided in the next section.

D. Critical f -policy

In order to give a detailed description of critical f -policy, we have to define a *critical* rank ordering.

Definition. A rank ordering R is referred to as *critical* if for each node i there exist distinct nodes j_1, j_2, \dots, j_l such that $i \rightarrow j_1 \rightarrow j_2 \rightarrow \dots \rightarrow j_l \rightarrow 0$ and $j_n \preceq^R i$ for all $1 \leq n \leq l$.

The set of all critical rank orderings is denoted by \mathcal{R}_c and note that $\mathcal{R}_c \subseteq \mathcal{R}$. Let $\mathcal{A}_c(R) \subseteq \mathcal{A}(R)$ denote the set of all critical one-step refinements and confinements of R . We define $D_f^c(R)$, $R \in \mathcal{R}_c$, as

$$D_f^c(R) = \{Q \in \mathbb{R}_+^N : R <_Q R' \text{ for all } R' \in \mathcal{A}_c(R)\}.$$

Definition. The network is said to be *connected* if for each node i there exist nodes j_1, j_2, \dots, j_l such that $i \rightarrow j_1 \rightarrow j_2 \rightarrow \dots \rightarrow j_l \rightarrow 0$.

Next lemma renders the set of cones as a partition of \mathbb{R}_+^N .

Lemma 2. Assume the network is connected. If bivariate function f satisfies conditions (C1) and (C2), then for all $Q \in \mathbb{R}_+^N$, there exists a unique $R \in \mathcal{R}_c$ such that $Q \in D_f^c(R)$.

In other words, $\{D_f^c(R)\}_{R \in \mathcal{R}_c}$ is the set of cones that partition \mathbb{R}_+^N , and it is possible to define a function $\pi_f^c : \mathbb{R}_+^N \rightarrow \mathcal{R}_c$

such that $\pi_f^c(\mathbf{Q}) = R \Leftrightarrow \mathbf{Q} \in D_f^c(R)$.

Definition. A priority-based routing policy $\Pi_{\{R(t)\}}$ is said to be a *critical f -policy* if $R(t) = \pi_f^c(\mathbf{Q}(t))$.

Example 3. Consider the network of four nodes given in Example 2. Note that $(\{2\}, \{1\}, \{3\})$, $(\{2\}, \{3\}, \{1\})$, and $(\{2\}, \{1, 3\})$, are not critical. Fig. 3 shows the structure of the cones $\{D_f^c(R)\}_{R \in \mathcal{R}_c}$ where \mathcal{R}_c is the set of all critical rank orderings of $\{1, 2, 3\}$, and $f(m, n) = \frac{1}{3^m(3^n-1)}$. Note the difference with Fig. 2(b) depicting $\{D_f(R)\}_{R \in \mathcal{R}}$.




Fig. 3. Structure of the critical cones for the network of Example 2

Next we state the main results of this paper.

E. Overview of the Results

Theorem 1. Let f be a bivariate function that satisfies conditions (C1) and (C2). Then the associated f -policy (critical f -policy) is throughput optimal.

Theorem 1 introduces a new class of throughput optimal routing policies. The sketch of the proof is provided in section III, with the details provided in [13].

Definition. Let $\Pi_{\{R(t)\}}$ and $\Pi'_{\{R'(t)\}}$ be two priority-based routings. We say $\Pi'_{\{R'(t)\}}$ respects $\Pi_{\{R(t)\}}$ if $R'(t)$ is a refinement of $R(t)$ for all time slots t .

Theorem 2. Suppose $\Pi_{\{R(t)\}}$ is a priority-based routing policy that is throughput optimal. Any priority-based routing policy that respects $\Pi_{\{R(t)\}}$ is also throughput optimal.

Note that Theorem 2 enables the proof of throughput optimality of specific routing policies. For example, in section IV-B, Theorems 1 and 2 are used to prove the throughput optimality of two known routing policies, backpressure [8] and ORCD [7]. The proof of Theorem 2 is simple and is given in [13].

III. THROUGHPUT OPTIMALITY OF f -POLICY

In this section we assume that routing decisions $\{\mu_{ij}^*(t)\}_{i,j \in \Omega}$, are made under an f -policy for which f is a bivariate function satisfying conditions (C1) and (C2). In this setting we prove that f -policy is throughput optimal. The proof is based on the following corollary to *Foster-Lyapunov Theorem*.

Fact 2. ([14, Lemma4.1]) Let $L^* : \mathbb{R}_+^N \rightarrow \mathbb{R}_+$ be a Lyapunov function. If there exist constants $B > 0$, $\epsilon > 0$, such that for all time slots t we have:

$$\mathbb{E}[L^*(\mathbf{Q}(t+1)) - L^*(\mathbf{Q}(t)) | \mathbf{Q}(t)] \leq B - \epsilon \sum_{k=1}^N Q_k(t),$$

then the network is stable, i.e, the time average queue backlog of all nodes remain finite.

To prove Theorem 1, we identify a class of Lyapunov functions that under the corresponding f -policy satisfy the conditions of Fact 2 for all arrival rate vectors $\lambda \in \text{int}(\mathfrak{S})$. In particular, we construct a piece-wise Lyapunov function, $L_f^* : \mathbb{R}_+^N \rightarrow \mathbb{R}_+$, by assigning to each cone $D_f(R)$, $R = (C_1, C_2, \dots, C_M)$, a quadratic function of the queue backlogs:

$$L_f(\mathbf{Q}, R) = \sum_{i=1}^M f(|C^{i-1}|, |C_i|) Q_{C_i}^2.$$

Since the collection of cones form a partition of \mathbb{R}_+^N , we can combine the above quadratic functions to arrive at a piece-wise quadratic function

$$L_f^*(\mathbf{Q}) = L_f(\mathbf{Q}, \pi_f(\mathbf{Q})) = \sum_{R \in \mathcal{R}} L_f(\mathbf{Q}, R) \mathbf{1}_{\{\mathbf{Q} \in D_f(R)\}}. \quad (4)$$

Lemma 3. $L^*(\cdot)$ is continuous and differentiable.

Note that the continuity and differentiability of $L^*(\cdot)$ follow 1) the continuity and differentiability of the construction of $L_f(\cdot, R)$ inside the cone corresponding to R , as well as 2) the construction of penalty function on the separating hyperplanes at the boundary of $D_f(R)$. The details are given in [13].

Next we provide the main steps in showing L^* has a negative expected drift.

Let us consider the Lyapunov drift when $\mathbf{Q}(t) \in D_f(R)$ for some $R = (C_1, C_2, \dots, C_M) \in \mathcal{R}$. By Lemma 3, $L_f^*(\cdot)$ is continuous and differentiable. Thus, we can write $L_f^*(\mathbf{Q}(t+1))$ in terms of its first-order Taylor expansion around $L_f^*(\mathbf{Q}(t))$ and we obtain

$$\begin{aligned} & L_f^*(\mathbf{Q}(t+1)) - L_f^*(\mathbf{Q}(t)) \\ &= (\mathbf{Q}(t+1) - \mathbf{Q}(t)) \cdot \nabla L_f^*(\mathbf{Q}(t)) + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|) \\ &= \sum_{i=1}^M f(|C^{i-1}|, |C_i|) 2Q_{C_i}(t) (Q_{C_i}(t+1) - Q_{C_i}(t)) + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|) \\ &= \sum_{i=1}^M f(|C^{i-1}|, |C_i|) [Q_{C_i}^2(t+1) - Q_{C_i}^2(t) - (Q_{C_i}(t+1) - Q_{C_i}(t))^2] + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|) \\ &= \sum_{i=1}^M f(|C^{i-1}|, |C_i|) [Q_{C_i}^2(t+1) - Q_{C_i}^2(t)] + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|). \end{aligned} \quad (5)$$

Lemma 4. Let $R = (C_1, C_2, \dots, C_M) \in \mathcal{R}$ and $\mathbf{Q}(t) \in D_f(R)$. We have

$$Q_{C_i}^2(t+1) - Q_{C_i}^2(t) \leq \beta_f - 2Q_{C_i}(t)(\mu_{C_i, \text{out}}^*(t) - \mu_{C_i, \text{in}}^*(t) - A_{C_i}(t)),$$

where β_f is a constant bounded real number.

Now taking expectation from both sides of (5) and using Lemma 4 we obtain,

$$\begin{aligned} & \mathbb{E} [L_f^*(\mathbf{Q}(t+1)) - L_f^*(\mathbf{Q}(t)) | \mathbf{Q}(t)] \\ & \leq B_f - 2 \sum_{i=1}^M f(|C^{i-1}|, |C_i|) Q_{C_i}(t) \mathbb{E} [\mu_{C_i, out}^*(t) - \mu_{C_i, in}^*(t) - A_{C_i}(t) | \mathbf{Q}(t)] + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|), \end{aligned} \quad (6)$$

where B_f is a constant bounded real number.

Lemma 5 below shows that under an f -policy the negative drift term in (6) is bounded by the negative drift under any other set of routing decisions, including the stabilizing randomized rule $\{\tilde{\mu}_{ij}(t)\}_{i,j \in \Omega}$ given in Fact 1.

Lemma 5. *Let $R = (C_1, C_2, \dots, C_M) \in \mathcal{R}$, $\mathbf{Q}(t) \in D_f(R)$, and let $\{\mu_{ij}^*(t)\}_{i,j \in \Omega}$ represent routing decisions made under an f -policy. For any collection of routing decisions $\{\mu_{ij}(t)\}_{i,j \in \Omega}$, we have*

$$\sum_{i=1}^M f(|C^{i-1}|, |C_i|) Q_{C_i}(t) (\mu_{C_i, out}^*(t) - \mu_{C_i, in}^*(t)) \geq \sum_{i=1}^M f(|C^{i-1}|, |C_i|) Q_{C_i}(t) (\mu_{C_i, out}(t) - \mu_{C_i, in}(t)). \quad (7)$$

However, since $\lambda \in \text{int}(\mathfrak{S})$, there exists a positive vector ϵ (vector of length N with all elements equal to ϵ , $\epsilon > 0$) such that $\lambda + \epsilon \in \mathfrak{S}$. Thus, from Fact 1

$$\mathbb{E} [\tilde{\mu}_{C_i, out}(t) - \tilde{\mu}_{C_i, in}(t) - A_{C_i}(t) | \mathbf{Q}(t)] \geq \epsilon. \quad (8)$$

Combining (7) and (8) with (6), we have

$$\mathbb{E} [L_f^*(\mathbf{Q}(t+1)) - L_f^*(\mathbf{Q}(t)) | \mathbf{Q}(t)] \leq B_f - 2\epsilon \sum_{i=1}^M f(|C^{i-1}|, |C_i|) Q_{C_i}(t) + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|). \quad (9)$$

Property (C2) of function f , however, implies that

$$f(0, |C_1|) \geq f(|C^1|, |C_2|) \geq \dots \geq f(|C^{M-1}|, |C_M|) \geq f(|C^M|, 1) = f(N, 1). \quad (10)$$

Combining (9) and (10) gives that

$$\mathbb{E} [L_f^*(\mathbf{Q}(t+1)) - L_f^*(\mathbf{Q}(t)) | \mathbf{Q}(t)] \leq B_f - 2\epsilon f(N, 1) \sum_{k=1}^N Q_k(t) + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|). \quad (11)$$

Since $\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|$ is bounded, there exists a constant, say B'_f , such that $B_f + o(\|\mathbf{Q}(t+1) - \mathbf{Q}(t)\|) \leq B'_f$ for all time slots t . Therefore, we can rewrite (11) as

$$\mathbb{E} [L_f^*(\mathbf{Q}(t+1)) - L_f^*(\mathbf{Q}(t)) | \mathbf{Q}(t)] \leq B'_f - \epsilon' \sum_{k=1}^N Q_k(t),$$

where $\epsilon' = 2\epsilon f(N, 1)$. Now from Fact 2, the proof of Theorem 1 is complete.

Note that the proof of throughput optimality for critical f -policy follows similar lines above and is given in [13].

IV. f -POLICY AND THE DESIGN OF THROUGHPUT OPTIMAL ROUTING POLICIES

A. The structure of the Lyapunov function for f -policy

In Examples 2 and 3 we considered a network of four nodes and showed the structure of (critical) cones for that network. In this section we study the Lyapunov function as defined in (4) for the same network. Fig. 4 illustrates routing decision cones $D_f(R)$ and the associated quadratic function $L(\cdot, R)$.




Fig. 4. Structure of the Lyapunov function for the network of Example 2

As shown in Fig. 4, f -policy groups the queues based on their backlogs. In the central cone, all nodes belong to the same ranking class. The Lyapunov function in this cone is the squared sum of all queue backlogs and it is clear that the Lyapunov drift in this case is the same (and negative from Theorem 1) for all non-idling routing policies. This property allows a routing policy to potentially deviate from backpressure decisions in order to arrive at a better delay performance. However, when one of the queues becomes relatively large in comparison to the other nodes' backlogs, the backlog vector falls in a cone in which the node with large backlog is in a separate ranking class. The Lyapunov function for this cone, now, is the squared queue backlog of the node with large backlog plus the squared sum of other queue backlogs and its negative drift is ensured only when packets are routed from the node with disproportionately large backlog to other nodes. Similarly one can analyze the behavior of the Lyapunov function in other cones.

It is important to observe that the Lyapunov function in cones 1, 3, 5, 7, 9, and 11, is a weighted quadratic function, closely related to the quadratic Lyapunov function associated with backpressure routing. As discussed in [15], this ensures the negative direction when leaving the subspaces of the form $\{Q_i = 0\}_{i \in I}$, for some $I \subseteq \{1, 2, \dots, N-1\}$. What is significant in the structure of the Lyapunov function is this behavior together with the flexibility in choice of function f . In effect, function f determines the size of cones associated with various rank orderings, i.e. it provides flexibility in the relative occupancy time spent in non-idling cone 13 versus backpressure-like cones 1, 3, 5, 7, 9, 11.

B. Two Known Examples

In this section we use Theorems 1 and 2 to prove the throughput optimality of two known routing policies, backpressure [8] and ORCD [7].

In the opportunistic variant of backpressure routing, DIVBAR [6], among the set of nodes that have received a packet transmitted by node i , one of the nodes with the largest positive differential queue backlog is selected as the next forwarder. Therefore, backpressure is a priority-based routing policy $\Pi_{\{R_b(t)\}}$ where $R_b(t)$ is a partitioning of the nodes based on their queue backlogs, with smaller backlog meaning lower rank, i.e. $Q_i(t) < Q_j(t)$ implies that $i \prec^{R_b(t)} j$. This policy is shown to provide throughput optimality [6]. Next we give an alternative proof. This is done by showing that for any bivariate function f that satisfies conditions (C1) and (C2), backpressure respects f -policy. In other words, we show that if node j has a lower rank than node k under any f -policy, then $Q_j < Q_k$. The proof is immediate using the following Lemma.

Lemma 6. *Let $R = (C_1, C_2, \dots, C_M) \in \mathcal{R}$ and $\mathbf{Q} \in D_f(R)$. Then for any node k in ranking class C_i ,*

$$Q_k > \frac{f(0, |C^{i-1}|)}{f(|C^{i-1}|, 1)} Q_{C^{i-1}} \geq Q_{C^{i-1}}.$$

In the rest of this section, we give a brief description of another congestion-based routing policy, known as ORCD [7] and prove its throughput optimality. In [7], ORCD was introduced as an alternative to backpressure routing to improve its delay performance. However, the throughput optimality of ORCD remained unproven.

ORCD is a priority-based routing policy $\Pi_{\{R_{CD}(t)\}}$ in which nodes are ordered according to a cost measure of congestion “down the stream” from each node i denoted by $V_i(t)$. In other words, $i \prec^{R_{CD}(t)} j$ if $V_i(t) < V_j(t)$. The congestion cost measure for node i , $V_i(t)$, can be calculated for each time slot t in a recursive manner using a stochastic variant of Dijkstra algorithm. The recursive procedure results in a vector $[V_0(t), V_1(t), \dots, V_N(t)]$ that satisfies the following fixed point equation for all time slots t :

$$V_0(t) = 0, \tag{12}$$

$$V_i(t) = \frac{Q_i(t)}{p(i, t)} + \sum_{k \in U_i(t)} \frac{p(i, k, t)}{p(i, t)} V_k(t), \tag{13}$$

where $U_i(t) = \{k : i \rightarrow k \text{ and } V_k(t) < V_i(t)\}$ is the set of nodes with costs strictly lower than node i at time t , and

$$p(i, t) = \sum_{S: S \cap U_i(t) \neq \emptyset} P(S|i), \tag{14}$$

$$p(i, k, t) = \sum_{S: k \in S \cap U_i(t), V_k(t) < V_j(t) \forall j \in S} P(S|i). \tag{15}$$

Next we prove the throughput optimality of ORCD by showing that ORCD respects critical f -policy corresponding to any bivariate function f that satisfies condition (C1) and for all $m \geq 0$ and $n_1, n_2 > 0$

$$\frac{f(m, n_1)}{f(m + n_1, n_2)} \geq \frac{1}{p_{min}}, \tag{16}$$

where $p_{min} = \min \{P(S|i) : i \in \Omega, S \subseteq \Omega, P(S|i) > 0\}$. Note that, for instance, function $f(m, n) = \frac{1}{K^m(K^n - 1)}$, $K \geq 1 + \frac{1}{p_{min}}$, is such a function. In other words, we show that ORCD respects the critical f -policy for all such f . Mathematically, for all $j, k \in \Omega$ such that $j \prec^{\pi_f^c(\mathbf{Q}(t))} k$, then $j \prec^{R_{CD}(t)} k$ as well. Let $\pi_f^c(\mathbf{Q}(t)) = (C_1, C_2, \dots, C_M) \in \mathcal{R}_c$, and let $k \in C_i$ and $j \in C^{i-1}$. We consider two cases.

Case I. Node k reaches a node in $C^{i-1} \cup \{0\}$.

We next show

$$V_k(t) \geq Q_k(t) > \frac{Q_{C^{i-1}}(t)}{p_{min}} \geq V_j(t). \quad (17)$$

The first inequality in (17) is immediate from (13). The second and third inequality follow from the arguments below. In [13], it is shown that

$$Q_k(t) > \frac{f(0, |C^{i-1}|)}{f(|C^{i-1}|, 1)} Q_{C^{i-1}}(t) \geq \frac{Q_{C^{i-1}}(t)}{p_{min}}. \quad (18)$$

The following Lemma is useful in finding an upper bound for V_j , $j \in C^{i-1}$.

Lemma 7. *For any two nodes a and b , if $a \rightarrow b$, then*

$$V_a(t) \leq \frac{Q_a(t)}{p_{min}} + V_b(t).$$

Since $\pi_f^c(\mathbf{Q}(t))$ is critical, there exist distinct intermediate nodes $j_1, j_2, \dots, j_l \in C^{i-1}$ such that $j \rightarrow j_1 \rightarrow j_2 \rightarrow \dots \rightarrow j_l \rightarrow 0$. Using Lemma 7 recursively and noting that $V_0(t) = 0$, we have the following upper bound of $V_j(t)$,

$$V_j(t) \leq \frac{Q_j(t)}{p_{min}} + \frac{Q_{j_1}(t)}{p_{min}} + \dots + \frac{Q_{j_l}(t)}{p_{min}} \leq \frac{Q_{C^{i-1}}(t)}{p_{min}}. \quad (19)$$

Combining (18) and (19) gives (17).

Case II. Node k does not reach any node in $C^{i-1} \cup \{0\}$.

Let \hat{C}_i be the set of nodes in C_i that reach a node in $C^{i-1} \cup \{0\}$. All the paths from node k to the destination are through the nodes in \hat{C}_i , hence, $V_k(t) \geq \min_{m \in \hat{C}_i} V_m(t)$. However, from Case I, for each node $j \in C^{i-1}$ and $m \in \hat{C}_i$, $V_m(t) \geq V_j(t)$. This completes the proof.

V. DISCUSSION AND FUTURE WORK

In this paper, we provided a large class of throughput optimal policies by considering a class of piece-wise quadratic Lyapunov functions. We also specialized our result to recover and prove the throughput optimality of two known routing policies, backpressure and ORCD.

In this paper we considered a single destination scenario and we assumed that the network topology as well as probability of successful transmissions were fixed. The generalization to the case of 1) multi-destination scenario and 2) ergodic time-varying network topology and transmission probabilities are believed to be also straight forward but remain as future areas of work.

In this paper we provided a proof for throughput optimality of ORCD with centralized controller. Proving the throughput optimality of distributed version of ORCD is also an area of future research.

REFERENCES

- [1] P. Larsson, "Selection Diversity Forwarding in a Multihop Packet Radio Network with Fading channel and Capture," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 47–54, October 2001.
- [2] M. Zorzi and R. R. Rao, "Geographic Random Forwarding (GeRaF) for Ad Hoc and Sensor Networks: Multihop Performance," *IEEE Transactions on Mobile Computing*, vol. 2, no. 4, pp. 337–348, 2003.
- [3] S. Biswas and R. Morris, "ExOR: Opportunistic Multi-hop Routing for Wireless Networks," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4, pp. 133–144, October 2005.
- [4] S.R. Das S. Jain, "Exploiting Path Diversity in the Link Layer in Wireless Ad hoc Networks," *Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, 2005. WoWMoM 2005.*, pp. 22–30, June 2005.

- [5] C. Lott and D. Teneketzis, "Stochastic Routing in Ad-hoc Networks," *IEEE Transactions on Automatic Control*, vol. 51, no. 1, pp. 52–72, January 2006.
- [6] M. J. Neely and R. Ugaonkar, "Optimal Backpressure Routing for Wireless Networks with Multi-Receiver Diversity," *Ad Hoc Networks*, vol. 7, no. 5, pp. 862–881, July 2009.
- [7] P. Gupta and T. Javidi, "Towards Throughput and Delay Optimal Routing for Wireless Ad-Hoc Networks," in *Asilomar Conference*, November 2007, pp. 249–254.
- [8] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1949, 1992.
- [9] S. Subramanian, S. Shakkottai, and P. Gupta, "On Optimal Geographic Routing in Wireless Networks with Holes and Non-Uniform Traffic," in *INFOCOM*, 2007, pp. 1019–1027.
- [10] S. Sarkar and S. Ray, "Arbitrary Throughput Versus Complexity Tradeoffs in Wireless Networks using Graph Partitioning," *IEEE Transactions on Automatic Control*, vol. 53, no. 10, pp. 2307–2323, November 2008.
- [11] L. Ying and S. Shakkottai, "On Throughput-Optimal Scheduling with Delayed Channel State Feedback," in *Information Theory and Applications Workshop*, 2008, pp. 339–344.
- [12] B. Smith and B. Hassibi, "Wireless erasure networks with feedback," 2008, CoRR abs/0804.4298.
- [13] M. Naghshvar, H. Zhuang, and T. Javidi, "A General Class of Throughput Optimal Routing Policies in Multi-hop Wireless Networks," *preprint*.
- [14] 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. 1–144, 2006.
- [15] S. Meyn, *Control Techniques for Complex Networks*, Cambridge University Press, 2008.