Heat-Diffusion: Optimal Dynamic Routing for Multiclass Multihop Wireless Networks

Reza Banirazi, Edmond Jonckheere, Bhaskar Krishnamachari

***Abstract*—A new routing policy, named Heat-Diffusion (HD), is developed for multiclass, multihop wireless networks subject to stochastic arrivals, randomly varying topology, and channel interference. The new policy uses only current queue backlogs and current channel states, without requiring the knowledge of topology and arrivals. We show that HD is throughput-optimal in the sense that it stabilizes all stabilizable arrival rates. We also show that among all stabilizing routing policies, HD minimizes a quadratic routing penalty defined by endowing each channel with a time-varying cost-factor. Beyond this, we consider a class of routing policies which base decisions only on current queue backlogs and current channel states, possibly using the knowledge of arrival statistics and channel state probabilities. It is shown that in this class, HD minimizes average total queue congestion, and so average network delay. Introducing a control parameter to trade average delay for average routing cost, it is also shown that in the aforementioned class of routing policies, HD provides a Pareto optimal tradeoff between these two criteria. Moreover, we show that HD fluid limit follows graph combinatorial heat equation and so Kirchhoff's laws on electrical networks. This particularly opens a new way to analyze wireless networks using powerful tools from heat calculus and circuit theory.**

I. INTRODUCTION

Throughput optimality, i.e. utilizing the full capacity of a wireless network, is critical to respond to increasing demand for wireless applications. The seminal work in [1] showed that the queue-differential, channel-rate-based *Back-Pressure (BP)* algorithm is throughput optimal under very general conditions on arrival statistics and channel state probabilities. Follow-up works showed that the class of throughput-optimal policies is indeed large [4]–[7]. The challenge is then to develop one that, in addition, is optimal relative to some other objectives.

We propose the Heat-Diffusion (HD), a throughput-optimal policy which operates under the same general conditions and with the same complexity as BP, while holding the following important qualities: (i) HD minimizes average routing cost in the sense of Dirichlet. Endowing each wireless link with a cost-factor, we define Dirichlet routing cost as the product of the link cost-factor and the square of the link flow rate. This routing cost may reflect different topology-based penalties, e.g. channel quality, routing distance, power usage, etc. (ii) Consider the class of routing algorithms which use only current queue occupancies and current channel states, possibly together with the knowledge of arrival/channel probabilities. In this class, HD minimizes average total queue congestion, which is proportional to average network delay by Little's Theorem. (iii) In the above class, suppose that the performance region built on the average delay and the Dirichlet routing cost is convex. Then HD operates on the *Pareto boundary* of this region by changing a control parameter which compromises between the two objectives.

The authors are with Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089.

E-mail: {banirazi, jonckhee, bkrishna}@usc.edu

This work has been submitted in part to the IEEE INFOCOM conference, Toronto, Canada, April 2014.

A. Related Work

The study of BP schemes has been a very active research area with wide-ranging applications and many recent theoretical results. In packet switches, congestion-based scheduling was extended to admit more general functions of queue lengths with a particular interest on α-weighted schedulers using α-exponent of queue lengths [5]. There has been a non-proved conjecture that heavy traffic delay is minimized when $\alpha \to 0$. A discussion of this was given in [8] along with some counterexamples. As another extension in packet switches, [6] introduced Projective Cone Schedulers (PCS) to allow scheduling with non diagonal weight assignments. The work in [7] generalizes PCS using a tailored "patch-work" of localized piecewise quadratic Lyapunov functions. In wireless networks, shadow queues enabled BP to handle multicast sessions and reduce the number of actual queues that need to be maintained [9]. Replacing queue-length by packet-age, [10] introduced a delay-based BP policy. To improve BP delay performance, [11] proposed place-holders with Last-In-First-Out (LIFO) scheduling. Adaptive redundancy was used in [12] to reduce light traffic delay in intermittently connected mobile networks. Using graph embedding, [18] combined BP with greedy routing in hyperbolic coordinates to obtain a throughput-delay tradeoff. Some attempts have been done to adopt the original framework for handling finite buffer sizes [13]. Many researchers have focused on solving the centralized BP scheduling in a distributed fashion so as to make the implementation more convenient [14], [15]. There have also been several reductions of BP to practice in the form of distributed wireless protocols of pragmatically implemented and experimentally evaluated [16], [17].

A common theme to nearly all the papers on BP strategy, going back to the original work [1], is that the policy is derived via the minimization of a *bound* on the drift of an energy function. While this is sufficient to prove queue stability for all stabilizable arrival rates, it cannot secure an optimal performance with respect to either queue congestion or routing cost. Further, BP-based schemes ensure stability via maximizing the number of forwarding packets. While this is fine in networks without routing decision, such as packet switches and one-hop wireless networks, it can induce unnecessary packet transmissions on a multihop network, which in turn increases queue variation and routing cost.

B. Contribution

We derive HD from the combinatorial analogue of classic heat equation on smooth manifolds. Translating "queue occupancy measured in packets" to "heat quantity measured in calories," the *fluid limit* of interference HD flow mimics a suitably-weighted non-interference heat flow, in agreement with the Second Principle of Thermodynamics. The key contributions of this work are as follows.

First, we introduce a new paradigm that might be called "Wireless Network Thermodynamics." This builds a deep connection between wireless networking and well-studied domains of physics and mathematics. In particular, it opens a way to take advantage of powerful tools from circuit theory, such as effective resistance and graph Laplacian, and those from geometry, such as heat calculus and curvature, in the analysis and design of stochastic, packet-based, time-slotted queuing networks constrained by link interference.

Second, the new policy reduces the Dirichlet routing cost to its minimum feasible value. To the best of our knowledge, this is the first time a feasible routing algorithm asserts the strict minimization of a cost function subject to network stability, i.e. bounded average delay. We show that the drift-plus-penalty approach of [3] can get only close to the minimum of this routing cost at the expense of infinitely large network delay.

Third, the new policy minimizes average network delay in the class of all algorithms which make routing decision as a pure function of current queue congestion and current channel states, including the ones with perfect probability knowledge on arrivals and channel states. This important class contains stationary randomized algorithms [3], original BP policy [1], and most BP derivations [4]–[17].

Fourth, in the above-mentioned class of routing algorithms, through changing a control parameter, HD provides a Pareto optimal performance with respect to average delay and Dirichlet routing cost under the convexity assumption on Pareto boundary. This means that no other policy in this class can make a better tradeoff between these two performance criteria than what HD does. It equivalently means that any deviation from HD operating point will lead to the performance degradation in at least one of these criteria.

Last but not least, HD enjoys the same algorithmic structure and complexity as BP, giving them the same wide-reaching impact. Further, this paves the way to leverage all advanced improvements to BP to further enhance HD quality, and also to practice via a smooth software transition from BP to HD.

C. Outline of Paper

The next section provides preliminaries. We introduce HD policy in Sec. III followed by HD key property which is fundamental to other qualities, and then some illustrative examples. Using *Lyapunov drift,* Sec. IV shows that HD is throughput-optimal, and that in a class of routing algorithms, it minimizes average delay. In Sec. V we propose the physics-oriented model of heat process on a *directed graph.* Using *fluid limit theory,* Sec. VI shows that in limit, HD packet flow resembles combinatorial heat flow on the underlying directed graph. Using *heat calculus,* Sec. VII shows that HD strictly minimizes the Dirichlet routing cost. We discuss HD *Pareto optimality* in Sec. VIII. The paper is concluded in Sec. IX.

II. Preliminaries

We consider a wireless network operating in slotted time with normalized slots $n \in \{0, 1, 2, \cdots\}$. The network is described by a *simple, directed* connectivity graph with set of nodes $\mathcal{V}$ and directed edges $\mathcal{E}$. New packets with different destinations in a set $\mathcal{K} \subseteq \mathcal{V}$ randomly arrive into different nodes, requiring a multihop routing. Packets of the same destination form a *class.* Each node i holds a separate queue $q_i^{(d)}$ for each *d-class* to transmit over its outgoing links. Each link may transmit packets from only one class at each slot.

Wireless channels may change due to node mobility or surrounding conditions. Assuming the sets $\mathcal{V}$ and $\mathcal{E}$ change much slower than channel states, we fix them during the time of our interest. Then a temporarily unavailable link (due to, e.g., obstacle effect, channel fading, etc.) is characterized by zero link capacity. We assume that channel states remain fixed during a timeslot, while they may change across slots.

In wireless networks, transmission over a channel can happen only if certain constraints are imposed on transmissions over the other channels. An interference model specifies these restrictions on simultaneous transmissions. Given an interference model, a *maximal schedule* is a set of channels such that no two channels interfere with each other, and no more channel can be added to it without violating the model constraints. We describe a maximal schedule with a *scheduling vector* $\boldsymbol{\pi} \in \{0, 1\}^{|\mathcal{E}|}$ where $\pi_{ij} = 1$ if the channel ij is included. Given a connectivity graph $(\mathcal{V}, \mathcal{E})$, we define the *scheduling set* Π as the collection of all *maximal* scheduling vectors.

Consider a wireless link $ij \in \mathcal{E}$. The link *capacity* $\mu_{ij}(n)$, which is frequently called link *transmission rate* in literature, counts the maximum number of packets the link can transmit at the slot n. The link *actual-transmission* $f_{ij}^{(d)}(n)$, on the other hand, counts the number of d-class packets *genuinely* sent over the link at the slot n. Each link is also endowed with a *cost-factor* $\rho_{ij}(n) \geqslant 1$ that represents the cost of transmitting one packet over the link at the slot n.

A discrete-time stochastic process $x(n)$ is stable if

$$\overline{x} := \limsup_{\tau \to \infty} 1/\tau \sum\nolimits_{n=0}^{\tau-1} \mathbb{E}\{x(n)\} < \infty \qquad (1)$$

where $\mathbb{E}$ denotes expectation.[1] A queuing network is *stable* if all its queues are stable. A traffic rate matrix is *stabilizable* if there exists a routing policy to stably support it. For a routing policy, *stability region* is the set of all traffic rate matrices that it can stably support. Network layer *capacity region* $\mathcal{C}$ is the union of the stability regions achieved by all routing policies (possibly unfeasible). A routing policy is *throughput-optimal* if it stabilizes the entire capacity region.

A. Problem Statement

For a constrained network described above, we propose HD algorithm that solves the three *stochastic* optimization

[1] This definition of stability is frequently called *strong stability* in contrast with other weaker stability definitions such as rate stability. It is shown that strong stability entails all of the other forms of stability [3, Theorem 2.8].

problems as follows. It is important to note that these problems must be solved at the *network layer* alone. This is totally different from cross-layer optimization [22]–[25] which aims to control congestion by tuning arrival rates into the network layer. With no control on arrivals, our basic assumption is that the arrival rates lie within the network capacity region making the system stabilizable. Obviously, nothing prevents one to install a flow controller on top of HD or develop an HD-based Network Utility Maximization (NUM) protocol.

- *Dirichlet routing cost minimization problem:*

$$\begin{aligned} \text{Minimize:} \quad & \overline{R} := \sum_{ij\in\mathcal{E}} \sum_{d\in\mathcal{K}} \overline{\rho_{ij}\big(f_{ij}^{(d)}\big)^2} \\ \text{Subject to:} \quad & \text{Throughput optimality} \end{aligned} \tag{2}$$

where the overbar notation denotes the $\limsup$ expected time average as (1). The loss function $\overline{R}$ is important from two aspects. As a quadratic cost function, it connects routing cost minimization to fundamental optimization methods in linear regression theory and linear-quadratic optimal control. By concept, it spreads out traffic with a weighted bias towards lower penalty links that reminds the optimal diffusion processes in physics, such as heat flow and electrical current.

It is shown in [2], [3] that a stationary randomized algorithm can solve the optimization problem (2). While such a policy exists in theory, it is intractable in practice as it requires a full knowledge of channel state probabilities. Further, assuming all of the probabilities could be accurately estimated, the network controller would still need to solve a dynamic programming problem for each topology state, where the number of states grows geometrically with the number of channels. However, we show in Th. 8 that HD policy solves this problem without requiring the knowledge of arrival statistics or channel state probabilities, and without dealing with dynamic programming.

- *Average network delay minimization problem:*

$$\begin{aligned} \text{Minimize:} \quad & \overline{Q} := \sum_{i\in\mathcal{V}} \sum_{d\in\mathcal{K}} \overline{q_i^{(d)}} \\ \text{Subject to:} \quad & \text{Network constraints.} \end{aligned} \tag{3}$$

Solving this problem for a general case requires the Markov structure of topology process, plus arrival and channel state probabilities. Then in theory, the solution is obtained through dynamic programming for each possible topology along with solving a Markov decision problem. By even having all of the required information, the number of queue backlogs and channel states increase exponentially with the size of network, making dynamic programming and Markov decision theory prohibitive. In fact, even for the case of a single channel, it is difficult to implement the resulting stochastic algorithms [26]. While having a practical solution for a general case seems dubious, we show in Th. 3 that HD policy solves this problem within an important class of routing algorithms, without requiring any of the above-mentioned information, and without dealing with dynamic programing or Markov decision process.

- *Pareto optimization problem:*

$$\begin{aligned} \text{Minimize:} \quad & (1-\beta)\,\overline{Q} + \beta\,\overline{R} \\ \text{Subject to:} \quad & \text{1) Throughput optimality} \\ & \text{2) Network constraints} \end{aligned} \tag{4}$$

where $\beta \in [0,1]$ is a control parameter to determine relative importance between average delay and average routing cost. To the best of our knowledge, this is the first time such a multi-objective optimization problem is addressed in the level of network layer. While even the corresponding single-objective optimization problems are not easy to manage, we show in Th. 9 that within the class of routing algorithms defined in problem (3), HD policy solves problem (4) subject to convex Pareto boundary on the feasible $(\overline{Q}, \overline{R})$ region.

B. Original Back-Pressure Policy

The original BP [1] for network layer, at every timeslot n observes queue backlogs $q_i^{(d)}(n)$ and estimates channel capacities $\mu_{ij}(n)$ to make a routing decision as follows.

1) BP weighing: For every link ij and on each class d find $q_{ij}^{(d)}(n) := q_i^{(d)}(n) - q_j^{(d)}(n)$ and select the optimal class

$$d_{ij}^*(n) := \arg\max_{d\in\mathcal{K}} q_{ij}^{(d)}(n). \tag{5}$$

Then give a weight to the link using its estimated capacity as

$$w_{ij}(n) := \mu_{ij}(n)\, q_{ij}^{(d^*)}(n)^+ \tag{6}$$

where $x^+ := \max\{0, x\}$ for any x.

2) BP scheduling: Find the scheduling vector such that

$$\boldsymbol{\pi}(n) = \arg\max_{\boldsymbol{\pi}\in\Pi} \sum\nolimits_{ij\in\mathcal{E}} \pi_{ij} w_{ij}(n) \tag{7}$$

where ties are broken arbitrarily.

3) BP forwarding: Over each activated link with $w_{ij}(n) > 0$ transmit from the class $d_{ij}^*(n)$ at full capacity $\mu_{ij}(n)$. If there is no enough d^*-classes at node i, transmit null packets.

C. V-Parameter Back-Pressure Policy

To incorporate the Dirichlet routing cost $\overline{R}$ into the original BP, the drift-plus-penalty approach [2], [3], which we refer to as *V-parameter BP* hereafter, adds a usage cost to the link queue-differential via replacing the link weight (6) by

$$w_{ij}(n) := \mu_{ij}(n)\big(\, q_{ij}^{(d^*)}(n) - V\rho_{ij}(n)\,\mu_{ij}(n)\,\big)^+ \tag{8}$$

where $V \in [0, \infty)$ trades queue occupancy for routing penalty. Note that the original BP is recovered for $V = 0$.

The V-parameter BP yields a Dirichlet routing cost within $O(1/V)$ from its minimum feasible value to the detriment of growing average delay by $O(V)$ relative to that of the original BP [3]. Thus the policy is not able to achieve minimum routing cost subject to throughput optimality, i.e. finite delay.

Another issue is that the resulting tradeoff depends on both V and the network with two negative consequences: (i) The same V leads to different tradeoffs in different networks. (ii) The resulting tradeoff changes by network topology and arrival rates. Hence, finding a proper V is difficult in practice.

III. Pareto Optimal Heat-Diffusion Policy

To provide a convenient way of unifying the new scheme with the previous works on BP, we design HD with the same complexity, in both computation and implementation, as BP.

A. Heat-Diffusion Algorithm

Our proposed policy for the network layer, at every timeslot n observes queue backlogs $q_i^{(d)}(n)$ and estimates channel

TABLE I
COMPARING HD WITH V-PARAMETER BP IN A UNICLASS NETWORK.

Weighing	$\widehat{f_{ij}}(n)$	BP	$\mu_{ij}(n)$
		HD	$\min\{(1-\beta+\beta/\rho_{ij}(n))q_{ij}(n)^{+},\ \mu_{ij}(n)\}$
	$w_{ij}(n)$	BP	$\mu_{ij}(n)\left[q_{ij}(n)-V\rho_{ij}(n)\,\mu_{ij}(n)\right]^{+}$
		HD	$2(1-\beta+\beta/\rho_{ij}(n))q_{ij}(n)\widehat{f_{ij}}(n)-\widehat{f_{ij}}(n)^2$
Scheduling			$\boldsymbol{\pi}(n)=\arg\max_{\boldsymbol{\pi}\in\Pi}\sum_{ij\in\mathcal{E}}\pi_{ij}w_{ij}(n)$
Forwarding			$f_{ij}(n)=\begin{cases}\widehat{f_{ij}}(n) & \text{if } \pi_{ij}(n)=1\\ 0 & \text{otherwise}\end{cases}$

capacities $\mu_{ij}(n)$ and channel cost-factors $\rho_{ij}(n)$ to make a routing decision as follows.

1) HD weighing: For each link ij the optimal class is obtained, in the same way as BP, from

$$d^{*}_{ij}(n) := \arg\max_{d\in\mathcal{K}}\ q^{(d)}_{ij}(n). \tag{9}$$

To weigh the link, every link first calculates the number of packets it would transmit if it were activated:

$$\begin{aligned}\widehat{f_{ij}}(n) &:= \min\{\phi_{ij}(n)q^{(d^{*})}_{ij}(n)^{+},\ \mu_{ij}(n)\}\\ \phi_{ij}(n) &:= (1-\beta)+\beta/\rho_{ij}(n)\end{aligned} \tag{10}$$

where the control parameter β is as defined in (4), and the hat notation denotes a predicted value which would not necessarily be realized. Then the link weight is determined by

$$w_{ij}(n) := 2\,\phi_{ij}(n)q^{(d^{*})}_{ij}(n)\widehat{f_{ij}}(n)-\widehat{f_{ij}}(n)^2. \tag{11}$$

2) HD scheduling: Find the scheduling vector, in the same way as BP, such that

$$\boldsymbol{\pi}(n)=\arg\max_{\boldsymbol{\pi}\in\Pi}\sum\nolimits_{ij\in\mathcal{E}}\pi_{ij}w_{ij}(n) \tag{12}$$

where ties are broken arbitrarily.

3) HD forwarding: Over each activated link transmit $\widehat{f_{ij}}(n)$ number of packets from the class $d^{*}_{ij}(n)$, meaning that

$$f^{(d)}_{ij}(n)=\begin{cases}\widehat{f_{ij}}(n) & \text{if } \pi_{ij}(n)=1\ \&\ d=d^{*}_{ij}(n)\\ 0 & \text{otherwise}\end{cases} \tag{13}$$

where $f^{(d)}_{ij}(n)$ represents the number of d-classes which are actually transferred over the link ij at the slot n. It is important to discriminate between genuinely realized $f^{(d)}_{ij}(n)$ and predicted value $\widehat{f_{ij}}(n)$, also between actual transmission $f^{(d)}_{ij}(n)$ and link capacity $\mu_{ij}(n)$.

Table 1 compares HD and V-parameter BP algorithms, emphasizing the same structure and complexity. For simplicity, the table shows a single destination case.

Remark 1: (i) Since $\rho_{ij}(n)\geqslant 1$ by assumption, we have $0<\phi_{ij}(n)\leqslant 1$ for all $\beta\in[0,1]$. (ii) If $q^{(d^{*})}_{ij}(n)\leqslant 0$, i.e. $q_i(n)\leqslant q_j(n)$ for all classes, we get $\widehat{f_{ij}}(n)=0$ due to (10), and $w_{ij}(n)=0$ from (11). In this case, even if the link were scheduled by (12), still no packet would be transmitted on it. (iii) If $q^{(d^{*})}_{ij}(n)>0$, we have $q^{(d^{*})}_{ij}(n)^{+}=q^{(d^{*})}_{ij}(n)$, and since $\widehat{f_{ij}}(n)\leqslant\phi_{ij}(n)q^{(d^{*})}_{ij}(n)$ due to (10), thus the link weight (11) is always nonnegative. (iv) As $q_{ij}(n)^{+}\leqslant q_i(n)$ and $\phi_{ij}(n)\leqslant 1$, the value of $\widehat{f_{ij}}(n)$ never exceeds the number of packets in the transmitting node.

Remark 2: In a very special case that all link capacities are the same, i.e. $\mu_{ij}(n)=\mu(n)$, and all link queue-differentials are always less than it, i.e. $q_{ij}(n)<\mu(n)$, HD with $\beta=0$ and α-weighted policy of [5] with $\alpha=2$ become equivalent. Packet switches are well suited to this special case.

B. Highlights of Heat-Diffusion Design

H1: While BP is derived by link capacity $\mu_{ij}(n)$, HD emphasizes on actual number of transmittable packets $\widehat{f_{ij}}(n)$, though it also takes into account the link capacity through (10). Thus HD allocates resources only based on genuinely transmittable packets, without counting on null packets.

H2: The link weight (11), which itself directly controls the scheduling optimization problem, is taken quadratic in the queue-differential $q_{ij}(n)$, where for $\phi_{ij}(n)q_{ij}(n)\leqslant\mu_{ij}(n)$ is simplified into $w_{ij}(n)=\phi_{ij}(n)^2q_{ij}(n)^2$. This contrasts with BP weighing $w_{ij}(n)=\mu_{ij}(n)q_{ij}(n)$ which is linear in $q_{ij}(n)$. The quadratic weight is central to the HD key property (Th. 1) from which all other HD qualities originate.

H3: Varying the convexity factor β makes a *universal* tradeoff in performance that depends neither on the network nor on the arrivals with the following significant results:

- For all $\beta\in[0,1]$ the policy is throughput-optimal (Th. 2).
- For $\beta=0$ the average total queue congestion $\overline{Q}$, and so average delay, decreases to its minimum feasible value within the class of routing algorithms which rely only on current queue backlogs and current channel states (Th. 3).
- Raising β increases the average delay in return for a lower routing cost, where the exclusive merit of HD is to provide the best tradeoff between these two objectives (Th. 9).
- For $\beta=1$ the Dirichlet routing cost $\overline{R}$ reaches its least feasible value (Th. 8) through an optimal tradeoff with average delay. Note that in the V-parameter BP, delay grows to infinity as routing cost is pushed towards its minimum.

Figure 1 graphically compares the operation of HD for $\beta\in[0,1]$ with V-parameter BP for $V\in[0,\infty)$, assuming that the performance region has a convex Pareto boundary. The performance region is restricted to all $\overline{Q}$ achievable by the class of routing algorithms which act based only on current queue congestion and current channel states.

H4: Unlike BP that forwards the highest number of packets over activated links, HD controls the packet forwarding by limiting it to $\phi_{ij}(n)q_{ij}(n)$ with the maximum $\phi_{ij}=1$ at $\beta=0$ and the minimum $\phi_{ij}=1/\rho_{ij}$ at $\beta=1$. This reduces *queue oscillation* by the decrease of unnecessary packet forwarding across the links, which itself reduces power consumption and

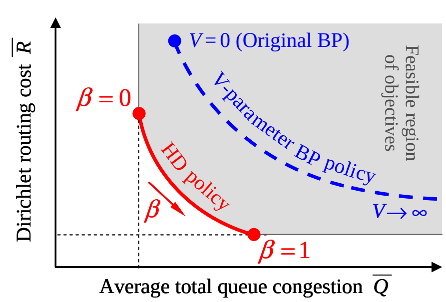

Fig. 1. Graphical description of HD Pareto optimality with respect to average queue congestion and Dirichlet routing cost, compared with BP performance.

routing penalty. Thus it is not surprising to see that ϕ_{ij} is decreasing, and so having a higher impact, by increasing β that means more emphasis on routing penalty. Forwarding a factor of queue differentials rather than filling up the link capacities also complies with mimicking *heat flow* on the underlying graph (Th. 5) that in effect minimizes time average routing cost via Dirichlet Principle (Th. 8).

C. The Key Property of Heat-Diffusion Policy

The next theorem formalizes the HD key property which is central to the proof of Th. 2 on HD throughput-optimality, Th. 3 on HD delay minimization, and Th. 5 on the connection between HD fluid limit and combinatorial heat equation leading to Th. 8 on HD routing cost minimization.

Theorem 1: At every timeslot n and for all $\beta \in [0,1]$, the HD policy solves the following optimization problem:

$$\begin{aligned} \text{Maximize:} \quad & \sum_{ij\in\mathcal{E}} \sum_{d\in\mathcal{K}} 2\,\phi_{ij}(n) q_{ij}^{(d)}(n) f_{ij}^{(d)}(n) - f_{ij}^{(d)}(n)^2 \\ \text{Subject to:} \quad & \text{Network constraints} \end{aligned} \tag{14}$$

where $\phi_{ij}(n) = (1-\beta) + \beta/\rho_{ij}(n)$ as defined in (10).

Proof: It follows directly from the HD algorithm (9)–(13) and is given in the appendix. ■

D. Illustrative Examples

To focus only on the policy itself, we take everything deterministic in our examples here. This assures us that the results purely show the policy performance not contaminated by stochastic effects. We however know that all HD properties are analytically proven for stochastic arrivals and random topologies under very general conditions.

Two-queue downlink: Consider a base station that transmits data to two downlink users, where at most one link can be activated at each timeslot. Let link 1 be of capacity $\mu_1 = 3$ (packets/slot) and link 2 of time-varying capacity $\mu_2 \geqslant 2$. Assume at every timeslot one packet arrives for each user. It is easy to verify that for $\mu_2 < 1.5$ the given arrival is beyond the capacity region. The performance of HD and BP are compared in Fig. 2 for $q_1(0) = q_2(0) = 0$. The leftmost panel depicts the timeslot evolution of $q_1(n)+q_2(n)$ for $\mu_2 = 18$. The rightmost panel depicts the steady-state average of total queue length as a function of μ_2. In BP the average total queue length increases linearly in μ_2, while HD holds the optimal performance for all admissible link capacities. This confirms H1 in the previous subsection, i.e., the efficiency of scheduling based on actual transmittable packets rather than link capacities.

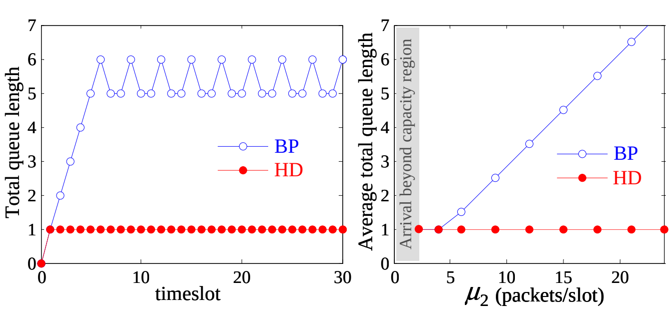

Fig. 2. Performance of HD versus BP in the two-queue downlink.

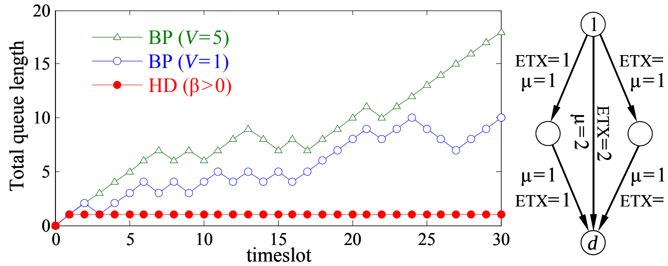

Fig. 3. Performance of HD versus BP in the lossy link network.

Lossy link network: Consider the 4-node network of Fig. 3 with lossy links and subject to 1-hop interference model. The links are labeled with both ETX and capacity, where ETX is a quality metric defined as the expected number of data transmissions required to send a packet without error over the link. Assume at every timeslot a single packet arrives at node 1 destined for node d. Following [16], let $\rho_{ij} = \text{ETX}_{ij}$. For zero initial conditions, Fig. 3 compares the performance of HD with BP. While HD easily stabilizes the total queued packets at one for any $\beta > 0$, trying with different values of V indicates the weakness of V-parameter BP in aptly supporting the arrival. This simplistically shows one of the impacts of entering the link cost-factor ρ_{ij} as a *multiplicand* in the HD formula (12) rather than an *addend* in the V-parameter BP formula (8).

Power minimization: Consider the sensor network of Fig. 4 subject to 1-hop interference model. Each link ij has a noise intensity $N_{ij} \in [1, 15]$ which is randomly assigned at first and keeps constant during the simulation. We adopt Shannon capacity $\mu_{ij} = B_{ij} \log_2(1 + P_{ij}/N_{ij})$ with P_{ij} the power transmission and B_{ij} the bandwidth. In each timeslot, two packets arrive at nodes 1, 2, 3, and 4 destined for the sink S. The aim is to minimize $\rho_{ij} f_{ij}^2$ where $\rho_{ij} = P_{ij}/\mu_{ij}$. For simplicity, we fix $P_{ij} = 15$ and $B_{ij} = 5$ for all links, so that the capacity of each link is determined by its noise intensity.

Figure 4 displays timeslot evolution of total queue lengths for HD with $\beta = 0$, and for the original BP where $V = 0$. Besides minimizing the average queue congestion, we notice little steady-state oscillations in HD contrary to large variations in BP that verifies H4. Figure 5 displays the tradeoff between queue congestion and power usage concurring with HD Pareto optimality displayed by Fig. 1. The attention is drawn on the rapid growth of queue lengths in V-parameter BP when average power usage is pushed downwards. Figure 6 displays timeslot evolution of total power usage for HD with $\beta = 1$, and for BP with $V = 1000$. Smaller oscillations in HD endorses both H1 and H4, showing the defect of capacity-based scheduling and maximum packet forwarding in BP.

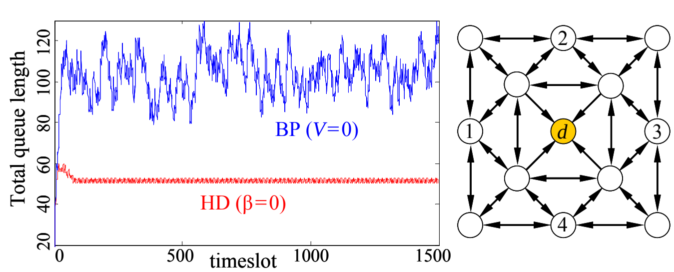

Fig. 4. Total queue backlog of HD versus BP in the power minimization.

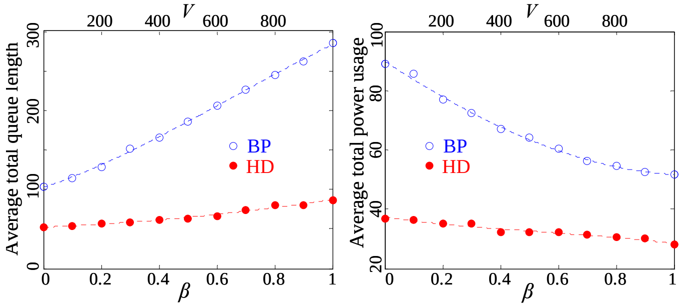

Fig. 5. Trading queue congestion for power consumption in HD as a function of β, and in BP as a function of V. Dashed lines represent interpolation.

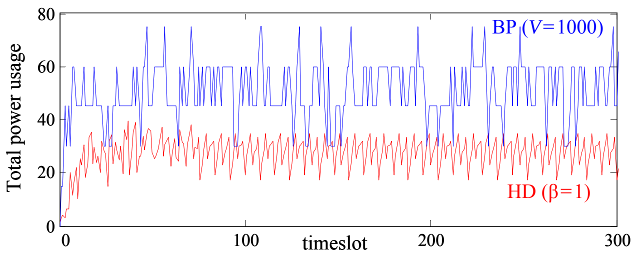

Fig. 6. Timeslot evolution of total power consumption in HD versus BP.

IV. Congestion Minimizing Routing Policy

A. State Space Representation of the System

Consider a general multiclass queuing network with $q_i^{(d)}(n)$ being the integer number of d-classes in the node i at the slot n as before. Assuming that the backlog of d-classes in the destination d is zero for all $d \in \mathcal{K}$, the state variables of the system are represented by the hyper-vector

$$\begin{aligned} \boldsymbol{q}_\circ(n) &:= \big[\, \boldsymbol{q}_\circ^{(1)}(n), \dots, \boldsymbol{q}_\circ^{(|\mathcal{K}|)}(n) \,\big]^\top \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|} \qquad (15)\\ \boldsymbol{q}_\circ^{(d)}(n) &:= \big[\, q_1^{(d)}(n), \dots, q_{d-1}^{(d)}(n), q_{d+1}^{(d)}(n), \dots, q_{|\mathcal{V}|}^{(d)}(n) \,\big]. \end{aligned}$$

where $q_d^{(d)}(n) \equiv 0$ is dropped from the set of state variables.

Throughout the paper, we use a $\circ$ subscript to denote a reduced vector or matrix obtained by discarding the entries corresponding to the destination node d.

Let a stochastic process $a_i^{(d)}(n)$ represent the integer number of exogenous d-classes arriving into the node i at the slot n. Discarding $a_d^{(d)}(n) \equiv 0$, the hyper-vector of node arrivals

$$\begin{aligned} \boldsymbol{a}_\circ(n) &:= \big[\, \boldsymbol{a}_\circ^{(1)}(n), \dots, \boldsymbol{a}_\circ^{(|\mathcal{K}|)}(n) \,\big]^\top \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|} \qquad (16)\\ \boldsymbol{a}_\circ^{(d)}(n) &:= \big[\, a_1^{(d)}(n), \dots, a_{d-1}^{(d)}(n), a_{d+1}^{(d)}(n), \dots, a_{|\mathcal{V}|}^{(d)}(n) \,\big]. \end{aligned}$$

Likewise, the hyper-vector of link actual-transmissions

$$\begin{aligned} \boldsymbol{f}(n) &:= \big[\, \boldsymbol{f}^{(1)}(n), \dots, \boldsymbol{f}^{(|\mathcal{K}|)}(n) \,\big]^\top \in \mathbb{R}^{|\mathcal{E}||\mathcal{K}|} \\ \boldsymbol{f}^{(d)}(n) &:= \big[\, f_1^{(d)}(n), \dots, f_{|\mathcal{E}|}^{(d)}(n) \,\big] \end{aligned} \qquad (17)$$

where as before, $f_{ij}^{(d)}(n)$ is the integer number of d-classes successfully transmitted over the link ij at the slot n. It is important to discriminate between the hyper-vector $\boldsymbol{f}(n)$ and the vector $\boldsymbol{\mu}(n) \in \mathbb{R}^{|\mathcal{E}|}$ representing link capacities at the slot n. In particular, while link capacities vary by topology, link actual-transmissions are assigned by a routing policy subject to $0 \leqslant f_{ij}^{(d)}(n) \leqslant \min\{q_i^{(d)}(n),\, \mu_{ij}(n)\}$.

Given a directed graph $(\mathcal{V}, \mathcal{E})$, let $\boldsymbol{B}$ denote the *node-edge* incidence matrix in which $B_{i\ell}$ is 1 if node i is the tail of directed edge ℓ, is -1 if i is the head, and is 0 otherwise. For a class d, let $\boldsymbol{B}_\circ^{(d)}$ denote a reduction of $\boldsymbol{B}$ through discarding the row corresponding to the destination node d. We refer to $\boldsymbol{B}_\circ^{(d)}$ as the *basis* incidence matrix with respect to the node d. Extending this notion to a multiclass framework, we build the *generalized* basis incidence matrix

$$\boldsymbol{B}_\circ := \mathrm{diag}\big(\, \big[\, \boldsymbol{B}_\circ^{(1)}, \dots, \boldsymbol{B}_\circ^{(|\mathcal{K}|)} \,\big] \,\big) \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}| \times |\mathcal{E}||\mathcal{K}|} \qquad (18)$$

where $\mathrm{diag}(\boldsymbol{v})$ is the diagonal matrix expansion of vector $\boldsymbol{v}$.

One can verify that $\boldsymbol{B}_\circ \boldsymbol{f}(n)$ is a hyper-vector in which the entry corresponding to node i and class d is given by

$$(\boldsymbol{B}_\circ \boldsymbol{f})_i^{(d)}(n) = \sum\nolimits_{b \in \mathrm{out}(i)} f_{ib}^{(d)}(n) - \sum\nolimits_{a \in \mathrm{in}(i)} f_{ai}^{(d)}(n)$$

where $\mathrm{in}(i)$ and $\mathrm{out}(i)$ respectively denote the set of incoming and outgoing neighbors of node i.

Using these ingredients, the $\boldsymbol{f}$-controlled, stochastic state dynamics of a multiclass queuing network is captured by

$$\boldsymbol{q}_\circ(n+1) = \boldsymbol{q}_\circ(n) + \boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ \boldsymbol{f}(n). \qquad (19)$$

Considering the difference between link capacity and link actual transmission explains why despite traditional notation in literature, we do not need any $(\cdot)^+$ operation in (19).

B. Characteristic of Network Capacity Region

Let a stochastic process $\boldsymbol{S}(n) = \big(S_1(n), \cdots, S_{|\mathcal{E}|}(n)\big)$ represent channel states at the slot n, describing conditions that affect channel capacities and link cost-factors. We assume that $\boldsymbol{S}(n)$ is an ergodic stationary process that takes values in a finite, but arbitrarily large, set $\mathcal{S}$. Thus by Birkhoff's ergodic theorem, each state $\boldsymbol{S} \in \mathcal{S}$ has a probability given by

$$s := \mathbb{P}\big\{\boldsymbol{S}(n) = \boldsymbol{S}\big\} = \limsup_{\tau \to \infty} 1/\tau \sum\nolimits_{n=0}^{\tau-1} \mathbb{I}_{\boldsymbol{S}(n) = \boldsymbol{S}} \qquad (20)$$

where $\sum_{\boldsymbol{S} \in \mathcal{S}} s = 1$ and $\mathbb{I}$ is the indicator function.

Consider a connectivity graph $(\mathcal{V}, \mathcal{E})$ together with a channel state process $\boldsymbol{S}(n)$. For an arrival *rate* vector $\overline{\boldsymbol{a}}$ to be in the network capacity region $\mathcal{C}$, the necessary and sufficient condition is the existence of a set of link actual-transmissions such that their expected time averages jointly satisfy node flow conservation and link capacity constraints, that is

$$\overline{a_i^{(d)}} = \sum\nolimits_{b \in \mathrm{out}(i)} \overline{f_{ib}^{(d)}} - \sum\nolimits_{a \in \mathrm{in}(i)} \overline{f_{ai}^{(d)}} \qquad (21)$$

$$\sum\nolimits_{i \in \mathcal{V}} \overline{a_i^{(d)}} = \sum\nolimits_{a \in \mathrm{in}(d)} \overline{f_{ad}^{(d)}} \qquad (22)$$

$$\sum\nolimits_{d \in \mathcal{K}} \overline{f_{ij}^{(d)}} \leqslant \sum\nolimits_{\boldsymbol{S} \in \mathcal{S}} s\, \mathbb{E}\big\{\mu_{ij}(n) \big| \boldsymbol{S}(n) = \boldsymbol{S}\big\} \qquad (23)$$

where the overbar notation is as defined in (1). Equalities (21) and (22) respectively secure flow conservation at intermediate nodes and at the destination d. Specifically, the matrix form of (21) becomes $\overline{\boldsymbol{a}_\circ} = \boldsymbol{B}_\circ \overline{\boldsymbol{f}}$ showing the time average of (19) subject to queue stability. The right hand side of (23) reads $\overline{\mu_{ij}}$ and so the inequality guarantees the link capacity constraint. Analytically important, the constraints (21)–(23) imply that the capacity region is convex, closed, and bounded [2].

Note that the link actual-transmissions are not fixed, but depend on the routing policy. Also observe that the number of routing policies which can potentially meet the constraints (21)–(23) is infinite. Among them are the ones that use the simple probability concept of distributing packets randomly so as the desired time averages (21)–(23) can be achieved. These stationary randomized policies typically require perfect knowledge of arrival and channel probabilities, making them an impractical solution. However, the fact that these *queue-independent* policies exist plays a crucial role in the analytical proof of HD properties in this section.

C. HD Throughput-Optimality for All β

To analyze the HD stability, we exploit the theory of Lyapunov stability for discrete-time systems. Unlike most of previous results in literature that use the sum of squares of queues as a Lyapunov candidate, we do not confine our attention only to classic quadratic functions derived from symmetric positive definite matrices. Let the Lyapunov candidate

$$\begin{aligned} W(n) &:= \boldsymbol{q}_\circ(n)^\top \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) \\ \boldsymbol{\Phi}(n) &:= \boldsymbol{I}_{|\mathcal{K}|} \otimes \operatorname{diag}\big(\boldsymbol{\phi}(n)\big) \end{aligned} \tag{24}$$

where $\boldsymbol{\phi}(n) \in \mathbb{R}^{|\mathcal{E}|}$ is the vector of $\phi_{ij}(n)$ as defined in (10), $\boldsymbol{I}_m$ denotes the identity matrix of size m, and $\otimes$ denotes tensor product. Note that $\big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ is not a symmetric matrix. Nevertheless, the next lemma ensures that the $W(n)$ function of (24) is indeed an energy function. Also note that for $\beta = 0$ where $\boldsymbol{\Phi}(n) = \boldsymbol{I}_{|\mathcal{E}||\mathcal{K}|}$, or for the case that all links are of the same cost-factor where $\boldsymbol{\Phi}(n) = \alpha \boldsymbol{I}_{|\mathcal{E}||\mathcal{K}|}$ for a scalar $\alpha > 0$, the $W(n)$ functional of (24) reduces to the sum of squares of queue lengths.

Lemma 1: Consider a wireless network with a connected topology. Then $\boldsymbol{M}_\circ(n) := \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ is *quasi-positive* for all n in the sense that $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} \geqslant 0$ for any vector $\boldsymbol{x}$, with equality if and only if $\boldsymbol{x} = \boldsymbol{0}$. Further,

$$\boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} = \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}\,, \quad \forall\, \boldsymbol{x} \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|}. \tag{25}$$

Proof: It is given in the appendix. ∎

Letting the Lyapunov drift $\Delta W(n) := W(n+1) - W(n)$ and substituting for $\boldsymbol{q}_\circ(n+1)$ from (19) lead to

$$\begin{aligned} \Delta W(n) &= \big(\boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ \boldsymbol{f}(n)\big)^\top \big(\boldsymbol{M}_\circ(n) + \boldsymbol{M}_\circ(n)^\top\big) \boldsymbol{q}_\circ(n) \\ &+ \boldsymbol{a}_\circ(n)^\top \boldsymbol{M}_\circ(n) \boldsymbol{a}_\circ(n) + \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n) \boldsymbol{B}_\circ \boldsymbol{f}(n) \\ &- \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \big(\boldsymbol{M}_\circ(n) + \boldsymbol{M}_\circ(n)^\top\big) \boldsymbol{a}_\circ(n) \end{aligned}$$

where $\boldsymbol{M}_\circ(n) = \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ as defined in Lem. 1.

Lemma 2: For arbitrary vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|}$,

$$\boldsymbol{x}^\top \big(\boldsymbol{M}_\circ(n)^\top + \boldsymbol{M}_\circ(n)\big)\, \boldsymbol{y} \leqslant \eta\, \boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}.$$

for a scalar η which takes the value 1 if $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y} \leqslant 0$ and the value 3 if $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y} > 0$.

Proof: It is given in the appendix. ∎

Exploiting Lem. 2 in the Lyapunov drift equation yields

$$\begin{aligned} \Delta W(n) &\leqslant \eta\, \big(\boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ \boldsymbol{f}(n)\big)^\top \boldsymbol{M}_\circ(n) \boldsymbol{q}_\circ(n) \\ &+ \boldsymbol{a}_\circ(n)^\top \boldsymbol{M}_\circ(n) \boldsymbol{a}_\circ(n) + \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n) \boldsymbol{B}_\circ \boldsymbol{f}(n) \\ &- \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \big(\boldsymbol{M}_\circ(n) + \boldsymbol{M}_\circ(n)^\top\big) \boldsymbol{a}_\circ(n). \end{aligned} \tag{26}$$

where η takes the value either 1 or 3 depending on if the functional $(\boldsymbol{a}_\circ - \boldsymbol{B}_\circ \boldsymbol{f})^\top \boldsymbol{M}_\circ \boldsymbol{q}_\circ$ is either negative or positive.

Turning back to the HD key property, we can formulate (14) in matrix form to obtain the following corollary to Th. 1.

Corollary 1: At every timeslot n and for all $\beta \in [0,1]$, the HD policy maximizes the $\boldsymbol{f}$-controlled functional

$$D(\boldsymbol{f}, \beta, n) := 2\, \boldsymbol{f}(n)^\top \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) - \boldsymbol{f}(n)^\top \boldsymbol{f}(n) \tag{27}$$

subject to network constraints, where $\boldsymbol{\Phi}(n)$ is as in (24). ◇

In the Lyapunov drift (26), let us replace $\boldsymbol{f}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{q}_\circ$ with $\boldsymbol{f}^\top \boldsymbol{\Phi} \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ$ in light of Lem. 1, add and subtract term $\frac{\eta}{2} \boldsymbol{f}^\top \boldsymbol{f}$, and use the $D(\boldsymbol{f}, \beta, n)$ expression of (27), to obtain

$$\begin{aligned} \Delta W(n) &\leqslant \eta\, \boldsymbol{a}_\circ(n)^\top \boldsymbol{M}_\circ(n) \boldsymbol{q}_\circ(n) - \frac{\eta}{2} D(\boldsymbol{f}, \beta, n) \\ &+ \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n) \Big(\boldsymbol{B}_\circ \boldsymbol{f}(n) - \boldsymbol{a}_\circ(n)\Big) \\ &+ \boldsymbol{a}_\circ(n)^\top \boldsymbol{M}_\circ(n) \boldsymbol{a}_\circ(n) - \frac{\eta}{2} \boldsymbol{f}(n)^\top \boldsymbol{f}(n). \end{aligned} \tag{28}$$

Now consider a traffic rate $\overline{\boldsymbol{a}_\circ}$ being interior to the capacity region $\mathcal{C}$, i.e. there exists a vector $\boldsymbol{\epsilon}$ with positive entries such that $\overline{\boldsymbol{a}_\circ} + \boldsymbol{\epsilon} \in \mathcal{C}$. Thus by condition (21), there exists a hyper-flow $\boldsymbol{f'}(n)$ such that $\overline{\boldsymbol{B}_\circ \boldsymbol{f'}} = \overline{\boldsymbol{a}_\circ} + \boldsymbol{\epsilon}$. At the same time, Corollary 1 guarantees that $D(\boldsymbol{f}^\star, \beta, n) \geqslant D(\boldsymbol{f'}, \beta, n)$ for all β and at each slot n, where $\boldsymbol{f}^\star(n)$ represents the link actual-transmissions provided by HD at the slot n. Then the next theorem is proven by showing that the expected value of Lyapunov drift (26) is bounded for all $\beta \in [0,1]$.

To simplify the proofs, throughout this section we assume both arrival and channel state processes are independently and identically distributed (i.i.d.) over timeslots. However, all the results can easily be extended to non-i.i.d. systems with stationary ergodic processes of finite mean and variance.

Theorem 2: Suppose that arrivals and channel states are i.i.d. over timeslots. The HD policy is throughput-optimal for all $\beta \in [0,1]$, meaning that it guarantees network stability under all stabilizable arrival rates.

Proof: For ease of notation, we will often drop the time variable (n). Taking conditional expectation from (28) given the current queue backlogs $\boldsymbol{q}_\circ(n)$ yields

$$\begin{aligned} \mathbb{E}\big\{\Delta W \,|\boldsymbol{q}_\circ\big\} &\leqslant \eta\, \mathbb{E}\big\{\boldsymbol{a}_\circ^\top \boldsymbol{M}_\circ \big|\boldsymbol{q}_\circ\big\}\, \boldsymbol{q}_\circ - \frac{\eta}{2}\, \mathbb{E}\big\{D(\boldsymbol{f}, \beta)\,\big|\boldsymbol{q}_\circ\big\} \\ &+ \mathbb{E}\Big\{\boldsymbol{a}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{a}_\circ + \boldsymbol{f}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{B}_\circ \boldsymbol{f} \\ &\qquad - \boldsymbol{f}^\top \boldsymbol{B}_\circ^\top (\boldsymbol{M}_\circ + \boldsymbol{M}_\circ^\top) \boldsymbol{a}_\circ - \frac{\eta}{2}\, \boldsymbol{f}^\top \boldsymbol{f}\, \Big|\boldsymbol{q}_\circ\Big\} \end{aligned} \tag{29}$$

where the conditional expectation is with respect to the randomness of arrivals, channel states, and routing decision—in case of a randomized routing algorithm.

Observe that $\boldsymbol{M}_\circ(n) = \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ is a function only of control parameter β and link cost-factors $\rho_{ij}(n)$. Since arrivals are independent of both β and ρ_{ij},

$$\mathbb{E}\big\{\boldsymbol{a}_\circ^\top \boldsymbol{M}_\circ \big|\boldsymbol{q}_\circ\big\} = \mathbb{E}\big\{\boldsymbol{a}_\circ^\top \big|\boldsymbol{q}_\circ\big\} + \mathbb{E}\big\{\boldsymbol{M}_\circ \big|\boldsymbol{q}_\circ\big\}.$$

Both β and ρ_{ij} are independent of $\boldsymbol{q}_\circ$, so is $\boldsymbol{M}_\circ$, which means that $\mathbb{E}\{\boldsymbol{M}_\circ|\boldsymbol{q}_\circ\} = \mathbb{E}\{\boldsymbol{M}_\circ\} = \overline{\boldsymbol{M}_\circ}$ where the last equality is due to the i.i.d. assumption on channel states. On the other hand, since the network layer routing controller has no impact on arrivals, $\boldsymbol{a}_\circ(n)$ is an independent system variable which is not influenced by anything. Therefore, $\boldsymbol{a}_\circ(n)$ is statistically uncorrelated with the random variable $\boldsymbol{q}_\circ(n)$. This results in $\mathbb{E}\{\boldsymbol{a}_\circ^\top|\boldsymbol{q}_\circ\} = \mathbb{E}\{\boldsymbol{a}_\circ^\top\} = \overline{\boldsymbol{a}_\circ}$ where the last equality is due to the i.i.d. assumption on arrivals. Putting these together yield

$$\mathbb{E}\big\{\boldsymbol{a}_\circ^\top \boldsymbol{M}_\circ \big|\boldsymbol{q}_\circ\big\}\, \boldsymbol{q}_\circ = \overline{\boldsymbol{a}_\circ}^\top\, \overline{\boldsymbol{M}_\circ}\, \boldsymbol{q}_\circ. \tag{30}$$

Let $\boldsymbol{f}(n)$ be the link actual-transmissions provided by HD with $\beta \in [0,1]$ at the slot n, and $\boldsymbol{f'}(n)$ the ones provided by any alternative policy. In light of Cor. 1, for each slot n

we have $D(\boldsymbol{f},\beta,n) \geqslant D(\boldsymbol{f'},\beta,n)$. Considering this with the equality (25) implies that

$$D(\boldsymbol{f},\beta,n) \geqslant 2\,\boldsymbol{f'}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\,\boldsymbol{q}_\circ(n) - \boldsymbol{f'}(n)^\top \boldsymbol{f'}(n).$$

Taking conditional expectation leads to

$$\mathbb{E}\big\{D(\boldsymbol{f},\beta)\big|\boldsymbol{q}_\circ\big\} \geqslant 2\,\mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ\big|\boldsymbol{q}_\circ\big\}\,\boldsymbol{q}_\circ - \mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{f'}\,\big|\boldsymbol{q}_\circ\big\}$$

for any alternative transmission decision $\boldsymbol{f'}(n)$. This includes the case when $\boldsymbol{f'}(n)$ is produced by a routing algorithm which makes independent, stationary, and randomized transmission decisions at each slot n based only on the current link capacities, and so independent of both queue backlogs and link cost-factors [2]. Let us fix $\boldsymbol{f'}(n)$ for such an algorithm. By independency of $\boldsymbol{f'}(n)$ from $\boldsymbol{q}_\circ$ and $\boldsymbol{M}_\circ$, $\mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ\big|\boldsymbol{q}_\circ\big\} = \mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{B}_\circ^\top\big\}\,\mathbb{E}\big\{\boldsymbol{M}_\circ\big|\boldsymbol{q}_\circ\big\}$. Since $\overline{\boldsymbol{a}_\circ}$ is interior to the capacity region $\mathcal{C}$, $\boldsymbol{f'}(n)$ can be arranged to stabilize an arrival rate $\overline{\boldsymbol{a}_\circ}+\boldsymbol{\epsilon} \in \mathcal{C}$. Then by the feasibility condition (21) and the i.i.d. assumption on arrivals and channel states, $\mathbb{E}\{\boldsymbol{B}_\circ \boldsymbol{f'}\} = \overline{\boldsymbol{a}_\circ}+\boldsymbol{\epsilon}$. Also we already showed that $\mathbb{E}\{\boldsymbol{M}_\circ|\boldsymbol{q}_\circ\} = \overline{\boldsymbol{M}}_\circ$. Therefore, $\mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ\big|\boldsymbol{q}_\circ\big\} = (\overline{\boldsymbol{a}_\circ}+\boldsymbol{\epsilon})^\top\,\overline{\boldsymbol{M}}_\circ$. Using this equality in conditional expectation of the D functional above,

$$\mathbb{E}\big\{D(\boldsymbol{f},\beta)\big|\boldsymbol{q}_\circ\big\} \geqslant 2\,(\overline{\boldsymbol{a}_\circ}+\boldsymbol{\epsilon})^\top\,\overline{\boldsymbol{M}}_\circ\,\boldsymbol{q}_\circ - \mathbb{E}\big\{\boldsymbol{f'}^\top \boldsymbol{f'}\,\big|\boldsymbol{q}_\circ\big\} \quad (31)$$

Exploiting the equality (30) and the inequality (31) in the drift inequality (29) leads to

$$\mathbb{E}\big\{\Delta W\,|\boldsymbol{q}_\circ\big\} \leqslant -\eta\,\boldsymbol{\epsilon}^\top\,\overline{\boldsymbol{M}}_\circ\,\boldsymbol{q}_\circ + \mathbb{E}\big\{\Gamma\,|\boldsymbol{q}_\circ\big\} \quad (32)$$

$$\Gamma := \boldsymbol{a}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{a}_\circ + \boldsymbol{f}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{B}_\circ \boldsymbol{f} - \boldsymbol{f}^\top \boldsymbol{B}_\circ^\top(\boldsymbol{M}_\circ + \boldsymbol{M}_\circ^\top)\boldsymbol{a}_\circ + \frac{\eta}{2}\,(\boldsymbol{f'}^\top \boldsymbol{f'} - \boldsymbol{f}^\top \boldsymbol{f}).$$

Investigating $\Gamma(n)$, note that (i) all arrivals have finite mean and variance, (ii) each link actual-transmission is at most equal to the link capacity which is finite, and thus both $\boldsymbol{f}(n)$ and $\boldsymbol{f'}(n)$ have a finite upper bound, and (iii) matrix $\boldsymbol{M}_\circ(n)$ is quasi-positive in the sense of Lem. 1, and so all its eigenvalues are positive and bounded (recall that $\phi_{ij}(n) \leqslant 1$). Therefore, the expected value of $\Gamma(n)$ is finite at each slot n, and so there exists a scalar $\Gamma_{\max}$ such that $\mathbb{E}\{\Gamma\,|\boldsymbol{q}_\circ\} \leqslant \Gamma_{\max}$. Now investigating $\boldsymbol{\epsilon}^\top\,\overline{\boldsymbol{M}}_\circ\,\boldsymbol{q}_\circ$, note that (i) $\boldsymbol{\epsilon}$ is a vector of positive entries, (ii) $\boldsymbol{q}_\circ(n)$ is a vector of nonnegative entries, and (iii) $\boldsymbol{M}_\circ(n)$ has all of its eigenvalues positive. Therefore, there exists a positive scalar $\lambda_{\min}$ for which $\boldsymbol{\epsilon}^\top\,\overline{\boldsymbol{M}}_\circ\,\boldsymbol{q}_\circ \geqslant \lambda_{\min}\boldsymbol{\epsilon}^\top \boldsymbol{q}_\circ$. Using these two inequalities in (32) leads to

$$\mathbb{E}\big\{\Delta W\,|\boldsymbol{q}_\circ\big\} \leqslant -\eta\,\lambda_{\min}\,\boldsymbol{\epsilon}^\top \boldsymbol{q}_\circ + \Gamma_{\max}.$$

Define $\epsilon_{\min}$ as the smallest entry of the $\boldsymbol{\epsilon}$ vector. Then for $\|\boldsymbol{q}_\circ\| > \Gamma_{\max}/(\eta\,\lambda_{\min}\epsilon_{\min})$ we get $\mathbb{E}\{\Delta W\,|\boldsymbol{q}_\circ\} < 0$. Then by Theorem 2 in [33], the queuing system is stable and so $\overline{\boldsymbol{a}_\circ}$ is in the HD stability region. As $\boldsymbol{\epsilon}$ can be arbitrarily small, this implies that any traffic rate $\overline{\boldsymbol{a}_\circ}$ being interior to the capacity region $\mathcal{C}$ is stabilized by HD with any $\beta \in [0,1]$, meaning that HD is throughput optimal for all $\beta \in [0,1]$. ■

D. Minimum Network Delay with $\beta = 0$

The HD Pareto optimality stands on two pillars: minimization of average delay $\overline{Q}$ for $\beta = 0$, and minimization of Dirichlet routing cost $\overline{R}$ for $\beta = 1$. In this subsection we establish the first pillar. Prior to stating the main result in Th. 3, we propose two lemmas which are used in the proof.

Lemma 3: The HD policy with $\beta = 0$ maximizes

$$2\,\boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) - \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{B}_\circ \boldsymbol{f}(n) \quad (33)$$

at every timeslot n subject to network constraints.

Proof: It is given in the appendix. ■

Considering Lem. 3 and Cor. 1 together, it turns out that both the functional (27) at $\beta = 0$, where $\boldsymbol{\Phi}(n)$ becomes the identity matrix, and the functional (33) are maximized for the same control action $\boldsymbol{f}(n)$. It is worth to remark that the claim is not about the same maximum value for these two functionals, but about the same maximizing solution $\boldsymbol{f}(n)$.

Lemma 4: Suppose that a general routing policy stabilizes an arrival rate vector $\overline{\boldsymbol{a}_\circ}$ resulting in timeslot queue occupancies $\boldsymbol{q}_\circ(n)$ and link actual-transmissions $\boldsymbol{f}(n)$. Then

$$2\,\overline{\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ, \boldsymbol{f}\}} - \overline{\mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ \boldsymbol{f}\}} = \overline{\mathbb{V}\mathrm{ar}\{\boldsymbol{a}_\circ\}} \quad (34)$$

where for two arbitrary random variables $\boldsymbol{X}$ and $\boldsymbol{Y}$ we define $\mathbb{C}\mathrm{ov}\{\boldsymbol{X},\boldsymbol{Y}\} := \mathbb{E}\{\boldsymbol{X}^\top \boldsymbol{Y}\} - \mathbb{E}\{\boldsymbol{X}\}^\top \mathbb{E}\{\boldsymbol{Y}\}$ and $\mathbb{V}\mathrm{ar}\{\boldsymbol{X}\} := \mathbb{C}\mathrm{ov}\{\boldsymbol{X},\boldsymbol{X}\}$, and where the overbar notation denotes the $\limsup$ expected time average as (1).

Proof: It is given in the appendix. ■

The equality (34) implies that a stabilizing routing decision with a higher average total variance of link forwardings necessarily results in a higher average total covariance between link forwardings and link queue differentials.

Having Lem. 3 and Lem. 4, the next theorem is proven in the class of all routing algorithms whose routing decision is a function only of current queue congestion and current channel states. This important class includes all opportunistic max-weight schedulers which do not incorporate the Markov structure of topology process into their decisions, among which is BP [1] and most its derivations [4]–[17]. It also includes all stationary randomized algorithms which make a routing decision as a pure (possibly randomized) function only of current channel states, typically using the perfect knowledge of arrival statistics and channel state probabilities.

Theorem 3: Suppose that arrivals and channel states are i.i.d. over timeslots. Consider a class of routing algorithms which act based only on current queue backlogs and current channel states. Within this class, the HD policy with $\beta = 0$ solves the average network delay minimization problem in (3).

Proof: For ease of notation, we will often drop the time variable (n). Consider a stabilizable traffic rate $\overline{\boldsymbol{a}_\circ}$ interior to the stability region of a "generic" routing policy which acts based only on current queue backlogs and current channel states. Also consider the functional (33) in Lem. 3, which is maximized by HD with $\beta = 0$, and let us denote it with $G(\boldsymbol{f},\boldsymbol{q}_\circ,n)$. If the generic policy maximizes G every timeslot, it results in the same time average performance as that of HD with $\beta = 0$. Thus we suppose that G obtained by the generic routing policy is not maximal. This implies that there exist a routing algorithm (possibly unfeasible), which also acts based only on current queue backlogs and current channel states, and a sufficiently small scalar $\delta > 0$ such that every timeslot the algorithm can stabilize the traffic rate $\overline{\boldsymbol{a}_\circ} + \delta\boldsymbol{1}$ while making G greater than or equal to that of the generic routing policy, where $\boldsymbol{1}$ denotes the vector of all ones with the same size as $\boldsymbol{a}_\circ$. We refer to this algorithm as "fictitious" as we do not intend to know how it really works. To assure that such an algorithm exists, one may endow it with the ability of perfect precognition to predict all future events without uncertainty.

Let $\boldsymbol{f}(n)$ and $\boldsymbol{q}_\circ(n)$ denote the timeslot quantities provided by the generic routing policy. Also let $\boldsymbol{f'}(n)$ be provided by the above-defined fictitious algorithm at the slot n given the current queue backlogs $\boldsymbol{q}_\circ(n)$. In view of the fact that $G(\boldsymbol{f'}, \boldsymbol{q}_\circ, n) \geqslant G(\boldsymbol{f}, \boldsymbol{q}_\circ, n)$ for every timeslot, we have

$$\begin{aligned} &2\,\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\}^\top\,\mathbb{E}\{\boldsymbol{q}_\circ\} + 2\,\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f'}\} \\ &\qquad - \mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\}^\top\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\} - \mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f'}\} \geqslant \\ &2\,\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\}^\top\,\mathbb{E}\{\boldsymbol{q}_\circ\} + 2\,\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\} \\ &\qquad - \mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\}^\top\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\} - \mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\} \end{aligned}$$

where the expectation is with respect to the randomness of arrivals, channel states, and routing decision—in case of randomized algorithms. We also have $\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\} = \mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\} + \delta$ due to the feasibility condition (21) and the i.i.d. assumption on arrivals and channel states. Plugging this equation in the above inequality, we obtain

$$\begin{aligned} 2\,\delta\mathbf{1}^\top\mathbb{E}\{\boldsymbol{q}_\circ\} &\geqslant 2\delta\mathbf{1}^\top\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\} - \delta^2\mathbf{1}^\top\mathbf{1} \\ &\quad + \big(2\,\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\} - \mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\}\big) \\ &\quad - \big(2\,\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f'}\} - \mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f'}\}\big) \end{aligned}$$

which holds for each timeslot. Summing over timeslots 0 until $\tau - 1$, dividing the sum by τ, and taking a $\limsup$ of $\tau \to \infty$ lead to the following time average result:

$$\begin{aligned} 2\,\delta\mathbf{1}^\top\,\overline{\boldsymbol{q}_\circ} &\geqslant 2\,\delta\mathbf{1}^\top(\overline{\boldsymbol{B}_\circ\boldsymbol{f'}}) - \delta^2\mathbf{1}^\top\mathbf{1} \\ &\quad + \big(2\,\overline{\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\}} - \overline{\mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\}}\big) \\ &\quad - \big(2\,\overline{\mathbb{C}\mathrm{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f'}\}} - \overline{\mathbb{V}\mathrm{ar}\{\boldsymbol{B}_\circ\boldsymbol{f'}\}}\big). \end{aligned}$$

Due to Lem. 4, the second and the third lines are canceling each other, which leaves the inequality with

$$2\,\delta\mathbf{1}^\top\,\overline{\boldsymbol{q}_\circ} \;\geqslant\; 2\,\delta\mathbf{1}^\top(\overline{\boldsymbol{B}_\circ\boldsymbol{f'}}) - \delta^2\mathbf{1}^\top\mathbf{1}. \tag{35}$$

Now consider the HD policy with $\beta = 0$ and suppose that it provides $\boldsymbol{f}^\star(n)$ and $\boldsymbol{q}_\circ^\star(n)$ at each slot n. Again, let $\boldsymbol{f'}(n)$ be provided by the fictitious algorithm at each slot n, but this time given the current queue backlogs $\boldsymbol{q}_\circ^\star(n)$. Due to Lem. 3, $G(\boldsymbol{f'}, \boldsymbol{q}_\circ^\star, n) \leqslant G(\boldsymbol{f}^\star, \boldsymbol{q}_\circ^\star, n)$ at every slot n. Performing the similar steps of taking expectation and then translating the results into the time average form, and using the equality $\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f'}\} = \mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}^\star\} + \delta$, we obtain

$$2\,\delta\mathbf{1}^\top\,\overline{\boldsymbol{q}_\circ^\star} \;\leqslant\; 2\,\delta\mathbf{1}^\top(\overline{\boldsymbol{B}_\circ\boldsymbol{f'}}) - \delta^2\mathbf{1}^\top\mathbf{1} \tag{36}$$

Comparing (35) and (36) leads to $\delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ^\star} \leqslant \delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ}$. This means that for any generic routing policy which acts based only on current queue backlogs and current channel states and stabilizes an arrival rate $\overline{\boldsymbol{a}_\circ}$ via creating an average queue congestion $\overline{\boldsymbol{q}_\circ}$, there exists a $\delta > 0$ such that $\delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ^\star} \leqslant \delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ}$ where $\overline{\boldsymbol{q}_\circ^\star}$ is the average queue congestion if HD with $\beta = 0$ is employed to stabilize the arrival rate $\overline{\boldsymbol{a}_\circ}$. Then observing that $\delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ^\star} \leqslant \delta\mathbf{1}^\top\overline{\boldsymbol{q}_\circ}$ is an equivalent expression for $\overline{Q^\star} \leqslant \overline{Q}$ with $\overline{Q}$ as defined in (3) concludes the proof. ■

V. Classical versus Combinatorial Heat Process

To design a routing policy within the spirit of heat diffusion, we need to bring the heat equation from ordinary smooth geometry to the purely combinatorial domain of a network. To formulate the heat diffusion process on a graph, we use the theory of *combinatorial geometry* where the notion of *chains-cochains* provides a genuine counterpart for *differential forms* in geometry. In this fashion, values on vertices, edges, and faces become analogues of pointwise functions, line integrals, and surface integrals, respectively. Having these fundamental elements of arbitrary degree, important structures and invariants of the smooth geometry are directly transferred to the discrete setting, culminating in a discrete *Hodge theory* which leads to *Laplace-deRham* operator. Interested readers can find the details in [27] and references therein.

A. Continuous Heat Diffusion on Manifolds

On a smooth manifold $\mathcal{M}$ charted in local coordinates $\boldsymbol{x}$, let $Q(\boldsymbol{x}, t)$ be the spatial distribution of temperature, $\boldsymbol{F}(\boldsymbol{x}, t)$ be the heat flux, and $A(\boldsymbol{x}, t)$ be the scalar field of heat sources (with minus for sinks). The law of heat conservation entails

$$\frac{\partial Q(\boldsymbol{x}, t)}{\partial t} = -\mathrm{div}\boldsymbol{F}(\boldsymbol{x}, t) + A(\boldsymbol{x}, t). \tag{37}$$

Fick's law states that heat flows from warm to cold regions with the heat flow proportional to the temperature gradient,

$$\boldsymbol{F}(\boldsymbol{x}, t) = -\sigma(\boldsymbol{x})\,\nabla Q(\boldsymbol{x}, t) \tag{38}$$

where $\sigma(\boldsymbol{x})$ is *thermal diffusivity* quantifying how fast heat moves through the material. Putting (37) and (38) together,

$$\frac{\partial Q(\boldsymbol{x}, t)}{\partial t} = \mathrm{div}\big(\sigma(\boldsymbol{x})\,\nabla Q(\boldsymbol{x}, t)\big) + A(\boldsymbol{x}, t). \tag{39}$$

To solve this equation uniquely, besides time initial condition, one needs to prescribe conditions of Q on a *boundary* $\partial\mathcal{M}$.

B. Continuous Heat Diffusion on Undirected Graphs

In the context of combinatorial geometry, we view a graph as a *simplicial 1-complex* and transfer the elements of classic heat equations to this cell complex as a discrete domain. In doing so, the smooth manifold $\mathcal{M}$ is replaced by a 0-*chain* vector, pointwise functions $Q(\boldsymbol{x}, t)$ and $A(\boldsymbol{x}, t)$ are respectively replaced by 0-*cochain* vectors $\boldsymbol{q}(t)$ and $\boldsymbol{a}(t)$ (node variables), line integral $\boldsymbol{F}(\boldsymbol{x}, t)$ is replaced by a 1-*cochain* vector $\boldsymbol{f}(t)$ (edge variable), and thermal diffusivity σ is replaced by an edge weight vector $\boldsymbol{\sigma}$.

The structure of a p-complex is fully described by a collection of k incidence matrices for $1 \leqslant k \leqslant p$. Specifically, the k-incidence matrix provides the algebraic *boundary* operator which transforms a k-chain into its *oriented* set of boundary cells. Conversely, the adjoint of the incidence matrix represents the *coboundary* operator which acts on cochains resembling the *exterior derivative* in classic geometry.

As a 1-complex, a graph has only the 1-incidence relation given by the *node-edge* incidence matrix $\boldsymbol{B}$ of Sec. IV, except that we substitute edge direction for algebraic topological edge orientation. Note that on a cell complex, the "orientation" gives an *arbitrary* direction to each cell. Specifically, edge orientation on a graph specifies which direction of flow is taken positive. This is *totally different* from edge direction on a directed graph discussed in the next subsection.

Using the above setup, on an *undirected* graph with node d as the single sink, the combinatorial analogue of the classic heat equations (37)–(39) are given by

$$\dot{\boldsymbol{q}}^{(d)}(t) = -\boldsymbol{B}\,\boldsymbol{f}^{(d)}(t) + \boldsymbol{a}^{(d)}(t)\,,\;\; q_d(t) = 0 \tag{40}$$

$$\boldsymbol{f}^{(d)}(t) = \operatorname{diag}(\boldsymbol{\sigma})\,\boldsymbol{B}^\top \boldsymbol{q}^{(d)}(t) \tag{41}$$

$$\dot{\boldsymbol{q}}^{(d)}(t) = -\boldsymbol{B}\operatorname{diag}(\boldsymbol{\sigma})\boldsymbol{B}^\top \boldsymbol{q}^{(d)}(t) + \boldsymbol{a}^{(d)}(t)\,,\;\; q_d(t) = 0 \tag{42}$$

where a dot on the top represents time derivative. Notice the boundary $\partial\mathcal{M}$ on the manifold is reduced to a single node d on the graph. Emphasizing this, the superscript (d) indicates that all generated heat in the graph goes into the node d.

Imposing the boundary condition $q_d(t)=0$, one can eliminate the sink d from (40)–(42). This yields the reduced set of *continuous-time* graph heat equations with the sink d as

$$\boldsymbol{f}^{(d)}(t) = \operatorname{diag}(\boldsymbol{\sigma})\,\boldsymbol{B}_\circ^{(d)\top} \boldsymbol{q}_\circ^{(d)}(t) \tag{43}$$

$$\begin{aligned} \dot{\boldsymbol{q}}_\circ^{(d)}(t) &= -\boldsymbol{L}_\circ^{(d)} \boldsymbol{q}_\circ^{(d)}(t) + \boldsymbol{a}_\circ^{(d)}(t) \\ \boldsymbol{L}_\circ^{(d)} &:= \boldsymbol{B}_\circ^{(d)} \operatorname{diag}(\boldsymbol{\sigma})\,\boldsymbol{B}_\circ^{(d)\top} \end{aligned} \tag{44}$$

where the $\circ$ subscript denotes a reduced quantity as before. The linear operator $\boldsymbol{L}_\circ^{(d)}$ is called *Dirichlet Laplacian* with respect to the node d, which is a *positive definite* matrix.

C. Continuous Heat Diffusion on Directed Graphs

On a directed graph, the combinatorial heat conservation (40) remains unchanged, but the Fick's law (41) must be modified to allow the flow in only one direction. Let the edge orientation concur with the edge direction. Like the undirected case, one can drop the sink d from the equations by fixing the boundary condition $q_d(t) = 0$. Then we obtain the reduced set of *continuous-time* heat equations on a *directed* graph as

$$\boldsymbol{f}^{(d)}(t) = \operatorname{diag}(\boldsymbol{\sigma}) \max\bigl\{\boldsymbol{0},\; \boldsymbol{B}_\circ^{(d)\top} \boldsymbol{q}_\circ^{(d)}(t)\bigr\} \tag{45}$$

$$\begin{aligned} \dot{\boldsymbol{q}}_\circ^{(d)}(t) &= -\vec{\boldsymbol{L}}_\circ^{(d)} \boldsymbol{q}_\circ^{(d)}(t) + \boldsymbol{a}_\circ^{(d)}(t) \\ \vec{\boldsymbol{L}}_\circ^{(d)} &:= \boldsymbol{B}_\circ^{(d)} \operatorname{diag}(\boldsymbol{\sigma})\,\boldsymbol{B}_\circ^{(d)\top} \operatorname{diag}\bigl(\mathbb{I}_{\boldsymbol{B}_\circ^{(d)\top} \boldsymbol{q}_\circ^{(d)}(t) \succ \boldsymbol{0}}\bigr) \end{aligned} \tag{46}$$

where $\boldsymbol{0}$ denotes the zero vector, $\max$ is taken entrywise, and $\mathbb{I}_{\boldsymbol{v}\succ\boldsymbol{0}}$ is the entrywise indicator *vector* function that its entry i is one if $v_i > 0$, and zero otherwise. We call $\vec{\boldsymbol{L}}_\circ^{(d)}$ the *nonlinear Dirichlet Laplacian* acting on a directed graph.

Remark 3: For the first time, heat diffusion on directed graphs is formulated via a nonlinear Laplacian. This is in agreement with the recent work in [28] showing that heat diffusion on Finsler manifolds, the natural counterparts of directed graphs in continuous domain, leads to a *nonlinear* Laplacian. In the graph literature, different *linear* Laplacians have been proposed for directed graphs [29, Sec. 3]. While successful to address some purely graphical issues, they are not able to convey the physics of the diffusion process, nor the intrinsic nonlinearity due to the one-way flow restrictions.

VI. Wireless Network Thermodynamics

Note that (45)–(46) still represent a deterministic, continuous, uniclass process with no link interference. The latter, particularly, makes the wireless problem quite intractable. Nevertheless, in this section we advocate a genuine diffusion process on stochastic, time-slotted, multiclass, interference networks by showing that the HD *fluid limit* at $\beta = 1$ follows graph heat equations on a suitably-weighted directed graph.

A. Fluid Limit of Heat-Diffusion Policy

Fluid limit of a stochastic process is the limiting dynamics obtained by *scaling* in time and amplitude. Under very mild conditions, it is shown that these scaled trajectories converge to a set of deterministic equations called *fluid model.* Using this deterministic model, one can analyze the *rate-level,* rather than *packet-level,* behavior of the original stochastic process. For the details, refer to [20], [21] and references therein.

Fluid limit: Let $\boldsymbol{X}(\omega,t)$ be a realization of a continuous-time stochastic process $\boldsymbol{X}$ along a sample path ω. Define the scaled process $\boldsymbol{X}^r(\omega,t) := \boldsymbol{X}(\omega,rt)/r$ for any $r > 0$. A deterministic function $\tilde{\boldsymbol{X}}(t)$ is a *fluid limit* if there exist a sequence r and a sample path ω such that $\lim_{r\to\infty} \boldsymbol{X}^r(\omega,t) \to \tilde{\boldsymbol{X}}(t)$ uniformly on compact sets. For a stable flow network, the existence of fluid limits is guaranteed if exogenous arrivals are of finite variance. It is further shown that each fluid limit is Lipschitz-continuous, and so differentiable, almost everywhere with respect to Lebesgue measure on $[0,\infty)$ [20], [21].

Cumulative process: To develop a continuous-time approximation of the system, we model the network by its *cumulative* processes. Let $\boldsymbol{a}_\circ^{\rm tot}(n)$ and $\boldsymbol{f}^{\rm tot}(n)$ be the hyper-vector of cumulative node arrivals and cumulative link transmissions up to the slot n. Assuming $\boldsymbol{a}_\circ^{\rm tot}(0)=\boldsymbol{0}$ and $\boldsymbol{f}^{\rm tot}(0)=\boldsymbol{0}$,

$$\boldsymbol{q}_\circ(n) = \boldsymbol{q}_\circ(0) + \boldsymbol{a}_\circ^{\rm tot}(n) - \boldsymbol{B}_\circ \boldsymbol{f}^{\rm tot}(n).$$

Let $\widehat{f_{ij}^{(d)}}(n)$ be the predicted number of d-classes that the link ij would transmit if it were activated in the slot n, and form the hyper-vector $\widehat{\boldsymbol{f}}(n)$ conformably structured as $\boldsymbol{f}(n)$ in (17). Also let $T_{\boldsymbol{\pi}}(n)$ be the cumulative number of timeslots in which the scheduling vector $\boldsymbol{\pi} \in \Pi$ has been selected. Assuming $T_{\boldsymbol{\pi}}(0) = 0$, it is not difficult to verify that

$$\boldsymbol{f}^{\rm tot}(n) = \sum_{\boldsymbol{\pi}\in\Pi} \sum_{k=1}^{n} \Bigl(T_{\boldsymbol{\pi}}(k) - T_{\boldsymbol{\pi}}(k-1)\Bigr)\Bigl((\boldsymbol{1}_{|\mathcal{K}|} \otimes \boldsymbol{\pi}) \odot \widehat{\boldsymbol{f}}(k)\Bigr)$$

where $\boldsymbol{1}_m$ is the vector of all ones with the size m, and $\odot$ denotes the entrywise product. The first parenthesis in the sum is equal to one if the scheduling vector $\boldsymbol{\pi}$ has been selected at the slot k, and zero otherwise. The term $(\boldsymbol{1}_{|\mathcal{K}|} \otimes \boldsymbol{\pi})$ extends the scheduling vector $\boldsymbol{\pi} \in \mathbb{R}^{|\mathcal{E}|}$ to be used in a multiclass fashion. Then its entrywise product with $\widehat{\boldsymbol{f}}(k)$ represents the number of packets that could be transferred from each class over each link if the scheduling vector $\boldsymbol{\pi}$ was selected. It is important to note that at every timeslot, the routing policy determines each entry of $\widehat{\boldsymbol{f}}(k)$ and selects a scheduling vector $\boldsymbol{\pi} \in \Pi$.

General equations: Given a sample path ω, we extend a time-slotted process to be continuous-time via linear interpolation in each interval $(n, n+1)$. Let exogenous arrivals occur at the beginning of each timeslot, so that $\boldsymbol{a}_\circ^{\rm tot}(t)$ represents the cumulative arrivals by the time t. Assuming normalized timeslots with the period of time unit, we obtain a set of continuous-time, stochastic, basic equations as

$$\boldsymbol{q}_\circ(t) = \boldsymbol{q}_\circ(0) + \boldsymbol{a}_\circ^{\rm tot}(t) - \boldsymbol{B}_\circ \boldsymbol{f}^{\rm tot}(t) \tag{47}$$

$$\dot{\boldsymbol{f}}^{\rm tot}(t) = \sum\nolimits_{\boldsymbol{\pi}\in\Pi} \dot{T}_{\boldsymbol{\pi}}(t)\bigl((\boldsymbol{1}_{|\mathcal{K}|} \otimes \boldsymbol{\pi}) \odot \widehat{\boldsymbol{f}}(t)\bigr) \tag{48}$$

$$\dot{T}_{\substack{\boldsymbol{\pi}(t) \\ \boldsymbol{\pi}\in\Pi}} = \begin{cases} 1 & \text{if } \boldsymbol{\pi} \text{ is chosen at time } t \\ 0 & \text{otherwise} \end{cases} \tag{49}$$

$$\sum\nolimits_{\boldsymbol{\pi}\in\Pi} T_{\boldsymbol{\pi}}(t) = t \;\text{ with } T_{\boldsymbol{\pi}}(t) \text{ nondecreasing.} \tag{50}$$

The equality (48) entails the existence of a $\delta > 0$ such that

$$f_{ij}^{\text{tot}(d)}(t') - f_{ij}^{\text{tot}(d)}(t) = \sum\nolimits_{\boldsymbol{\pi}\in\Pi} \pi_{ij}\, \widehat{f_{ij}^{(d)}}(t)\big(T_{\boldsymbol{\pi}}(t') - T_{\boldsymbol{\pi}}(t)\big)$$

for any $t' \in [t, t+\delta]$. This expresses the fact that if a link has a positive flow of d-classes at time t, the number of d-classes transmitted by the link in the interval $[t,t'] \subset [t,t+\delta]$ is equal to the amount of time it has been activated during $[t,t']$ multiplied by its transmission rate prediction at time t.

Particular equations: While (47)–(50) hold for any stable network operating under an arbitrary non-idling routing policy, each policy determines $\widehat{\boldsymbol{f}}(t)$ and $T_{\boldsymbol{\pi}}(t)$ in its own particular way. Specifically, referring to (9)–(13), HD imposes

$$\widehat{f_{ij}^{(d)}}(t) \overset{\text{HD}}{=} \begin{cases} \min\{\phi_{ij}(t) q_{ij}^{(d)}(t)^{+},\, \mu_{ij}(t)\} & \text{if } d = d_{ij}^{*}(t) \\ 0 & \text{otherwise} \end{cases} \tag{51}$$

$$\boldsymbol{w}(t) \overset{\text{HD}}{=} \widehat{\boldsymbol{f}(t)} \odot \big(\, 2\, \boldsymbol{\Phi}(t) \boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}(t) - \widehat{\boldsymbol{f}(t)}\,\big) \tag{52}$$

$$\boldsymbol{\pi}(t) = \arg\max\nolimits_{\boldsymbol{\pi}\in\Pi}\, (\mathbf{1}_{|\mathcal{K}|} \otimes \boldsymbol{\pi})^{\top} \boldsymbol{w}(t) \tag{53}$$

where $\boldsymbol{\phi}(t) \in \mathbb{R}^{|\mathcal{E}|}$ is the vector of $\phi_{ij}(t) = (1-\beta) + \beta/\rho_{ij}(t)$ as defined in (10), $\boldsymbol{\Phi}(t) = \boldsymbol{I}_{|\mathcal{K}|} \otimes \operatorname{diag}\big(\boldsymbol{\phi}(t)\big)$ as defined in (24), and $\boldsymbol{w}(t)) \in \mathbb{R}^{|\mathcal{E}||\mathcal{K}|}$ is a hyper-vector consists of the weights assigned by HD to each class on each link at time t. The immediate conclusion from the link weighing (52) and the link scheduling (53) is that the HD key property stated by Th. 1 in packet-level can equally be asserted in rate-level.

Corollary 2: At any given time t and for all $\beta \in [0,1]$, the HD policy maximizes the $\boldsymbol{f}$-controlled functional

$$D(\boldsymbol{f}, \beta, t) = 2\, \boldsymbol{f}(t)^{\top} \boldsymbol{\Phi}(t) \boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}(t) - \boldsymbol{f}(t)^{\top} \boldsymbol{f}(t) \tag{54}$$

subject to network constraints. $\diamondsuit$

For a comparison, observe that the original BP imposes

$$\widehat{f_{ij}^{(d)}}(t) \overset{\text{BP}}{=} \begin{cases} \min\{q_{i}^{(d)}(t), \mu_{ij}(t)\} & \text{if } d = d_{ij}^{*}(t) \;\&\; q_{ij}^{(d)}(t) > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\boldsymbol{w}(t) \overset{\text{BP}}{=} \big(\mathbf{1}_{|\mathcal{K}|} \otimes \boldsymbol{\mu}(t)\big) \odot \big(\boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}(t)\big)^{+}$$

where $\boldsymbol{\mu}(t) \in \mathbb{R}^{|\mathcal{E}|}$ is the vector of link capacities at time t, and $\boldsymbol{v}^{+} := \max\{\mathbf{0}, \boldsymbol{v}\}$ for any vector $\boldsymbol{v}$ with $\max$ taken entrywise. Accordingly, at any time t the policy maximizes $\big(\mathbf{1}_{|\mathcal{K}|} \otimes \boldsymbol{\mu}(t)\big)^{\top} \big(\boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}(t)\big)^{+}$ subject to network constraints.

Arrival and topology: Assume that every stochastic exogenous arrival $a_i^{(d)}(n)$ is of finite mean and variance. To derive a deterministic continuous-time fluid model, all stochastic input variables and all time-varying system parameters are replaced by their expected time average values. This provides us with another set of equations as

$$\boldsymbol{a}_{\circ}^{\text{tot}}(t) = \overline{\boldsymbol{a}_{\circ}}\, t \quad , \quad \boldsymbol{\rho}(t) = \overline{\boldsymbol{\rho}} \quad , \quad \boldsymbol{\mu}(t) = \overline{\boldsymbol{\mu}}\,. \tag{55}$$

where $\boldsymbol{\rho}(t) \in \mathbb{R}^{|\mathcal{E}|}$ is the vector of link cost-factors at time t. The existence of $\overline{\boldsymbol{\rho}}$ and $\overline{\boldsymbol{\mu}}$ is assured by (20), precisely as

$$\overline{\boldsymbol{\mu}} = \sum\nolimits_{\boldsymbol{S}\in\mathcal{S}} s\, \mathbb{E}\{\boldsymbol{\mu}(n) | \boldsymbol{S}(n) = \boldsymbol{S}\}$$
$$\overline{\boldsymbol{\rho}} = \sum\nolimits_{\boldsymbol{S}\in\mathcal{S}} s\, \mathbb{E}\{\boldsymbol{\rho}(n) | \boldsymbol{S}(n) = \boldsymbol{S}\}.$$

Though the HD routing decision does not depend on the average values in (55), the existence and finiteness of these values are important in the proof of the next theorem.

Theorem 4: Define the *HD fluid model* as the collection of deterministic continuous-time equations (47)–(55). Then under the HD policy with $\beta \in [0,1]$, each fluid limit $\tilde{\boldsymbol{X}}(t) = \big(\tilde{\boldsymbol{q}}_{\circ}(t), \tilde{\boldsymbol{f}}^{\text{tot}}(t), \tilde{T}_{\boldsymbol{\pi}}(t)\big)$ satisfies the HD fluid model.

Proof: The proof follows the same line of argument proposed by [20, Theorem 2.3.2] or [21, Proposition 4.12] and is omitted for brevity. ∎

B. Thermodynamic-Like Packet Routing

Consider a wireless network of $|\mathcal{K}|$ classes subject to an arrival rate stabilized by a routing policy. In steady-state condition, for the flow of each class d there exists a fluid limit $\tilde{\boldsymbol{X}}^{(d)}(t)$. We claim that under HD routing with $\beta \in [0,1]$, every fluid limit $\tilde{\boldsymbol{X}}^{(d)}(t)$ takes the form of the heat process on the underlying directed graph with suitably-weighted edges.

To conceptualize this claim, think of $|\mathcal{K}|$ fictitious graphs, all of the same incidence matrix as that of the wireless network and with the edge weights $\sigma_{ij} = \overline{\phi_{ij}}$. Associate with each arrival $a^{(d)}(n)$ on the real network a corresponding heat source with intensity $\overline{a^{(d)}}$ on the related fictitious graph for which the node d is the single sink. In this fashion, we create $|\mathcal{K}|$ isolated deterministic graphs, on each of which the flow of heat is governed by (45)–(46). Joining these $|\mathcal{K}|$ *decoupled* systems, we obtain our target trajectory with the dynamics of

$$\boldsymbol{f}^{\star}(t) = \boldsymbol{\Sigma} \max\{\mathbf{0},\, \boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}^{\star}(t)\} \tag{56}$$

$$\begin{aligned} \dot{\boldsymbol{q}}_{\circ}^{\star}(t) &= -\vec{\boldsymbol{L}}_{\circ} \boldsymbol{q}_{\circ}^{\star}(t) + \overline{\boldsymbol{a}_{\circ}} \\ \vec{\boldsymbol{L}}_{\circ} &:= \boldsymbol{B}_{\circ} \boldsymbol{\Sigma} \boldsymbol{B}_{\circ}^{\top} \operatorname{diag}\big(\mathbb{I}_{\boldsymbol{B}_{\circ}^{\top} \boldsymbol{q}_{\circ}^{\star}(t) \succ \mathbf{0}}\big) \end{aligned} \tag{57}$$

where $\boldsymbol{q}_{\circ}^{\star}(t)$, $\overline{\boldsymbol{a}_{\circ}}$ and $\boldsymbol{f}^{\star}(t)$ are conformably structured as (15), (16) and (17), respectively, and $\boldsymbol{\Sigma} := \boldsymbol{I}_{|\mathcal{K}|} \otimes \operatorname{diag}(\boldsymbol{\sigma})$ with $\boldsymbol{\sigma}$ the vector of edge weights.[2] Then we claim that the HD fluid model (47)–(55) follows the target model (56)–(57).

Note that $\sigma_{ij} = \overline{\phi_{ij}}$ depends not only on the link cost-factor ρ_{ij}, but also on the control parameter β, meaning that varying β leads to different edge weights, and so different graph topologies. Also note that in our target model, the sum of multiclass heat flows on each link is assumed to be less than the expected average of link capacity $\overline{\mu_{ij}}$. This is imposed since the graph heat equations (45)–(46) are derived under the assumption that heat flow on each link follows the Fick's law of diffusion, not disturbed by the link capacity.

To gain a better insight, let us first consider the case of a uniclass network (single destination). Let packets be routed under the HD policy (microscopic flow). Thus at every timeslot, the policy activates a particular set of links to transmit a specific number of packets over them. Obviously, each link transmits packets at some slots and is switched off at some other slots. Then the claim says that if we look at the average flow in limit (macroscopic flow), it mimics the heat flow on the underlying directed graph, where the limit flow of each link is determined by counting the total number of packets transmitted over the link during a sufficiently large period of time divided by the time duration. But the more interesting part appears when a multiclass network is considered. Now if we look at the limit flow of each individual class, it still

[2] As an extension of nonlinear Dirichlet Laplacian $\vec{\boldsymbol{L}}_{\circ}^{(d)}$ with the size of $(|\mathcal{V}|-1) \times (|\mathcal{V}|-1)$ in (46), one may call $\vec{\boldsymbol{L}}_{\circ}$ the *generalized* nonlinear Dirichlet Laplacian which is of the size $(|\mathcal{V}|-1)|\mathcal{K}| \times (|\mathcal{V}|-1)|\mathcal{K}|$ and acts on a multiclass directed graph with $|\mathcal{K}|$ different classes.

mimics the heat flow on the underlying directed graph, but of course with respect to its corresponding destination node.

Theorem 5: Consider the generalized graph heat equations as defined in (56)–(57), and suppose that the multiclass heat flow (56) satisfies the link capacity constraints (23). Then the HD fluid model (47)–(55) with $\beta \in [0,1]$ asymptotically converges to the generalized graph heat equations (56)–(57) with $\sigma_{ij} = \overline{\phi_{ij}}$. In particular, the HD fluid model with $\beta = 0$ converges to the generalized heat equations on an unweighted directed graph, and the HD fluid model with $\beta = 1$ converges to the generalized heat equations on a directed graph with the edge weights $\sigma_{ij} = 1/\overline{\rho_{ij}}$.

Proof: We show that $\boldsymbol{q}_\circ(t)$ and $\boldsymbol{f}(t)$ respectively converge to $\boldsymbol{q}^\star_\circ(t)$ and $\boldsymbol{f}^\star(t)$ in limit. Note in (56)–(57) that $\boldsymbol{q}^\star_\circ(t)$, and accordingly $\boldsymbol{f}^\star(t)$, reach their steady-state conditions exponentially fast. Thus it is sufficient to show that $\boldsymbol{q}_\circ(t)$ and $\boldsymbol{f}(t)$ respectively converge to the stationary values of $\boldsymbol{q}^\star_\circ(t)$ and $\boldsymbol{f}^\star(t)$ where $\dot{\boldsymbol{q}}^\star_\circ(t) = 0$. In other words, we prove that the two systems have the same behavior in steady-state condition, while they can possibly be different in transients.

Consider the continuous-time Lyapunov function

$$Y(t) := \big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big)^\top \, \overline{\boldsymbol{M}}_\circ \big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big) \tag{58}$$

where $\overline{\boldsymbol{M}}_\circ = \big(\boldsymbol{B}_\circ\boldsymbol{B}_\circ^\top\big)^{-1}\boldsymbol{B}_\circ\overline{\boldsymbol{\Phi}}\boldsymbol{B}_\circ^\top$ represents the expected time average of non-symmetric matrix $\boldsymbol{M}_\circ(n)$ as defined in Lem. 1, and where $\boldsymbol{q}^\star_\circ$ represents the time-independent stationary value of $\boldsymbol{q}^\star_\circ(t)$. Recall that all eigenvalues of $\boldsymbol{M}_\circ(t)$ are strictly positive by Lem. 1 and so $Y(t)$ is a positive-definite and radially unbounded function with respect to $\big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big)$. Taking time derivative from (58), we obtain

$$\dot{Y}(t) = \dot{\boldsymbol{q}}_\circ(t)^\top \big(\overline{\boldsymbol{M}}_\circ^\top + \overline{\boldsymbol{M}}_\circ\big)\big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big) \tag{59}$$

Exploiting Lem. 2 in the latter leads to

$$\dot{Y}(t) \leqslant \eta\, \dot{\boldsymbol{q}}_\circ(t)^\top \, \overline{\boldsymbol{M}}_\circ \big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big) \tag{60}$$

where η takes the value either 1 or 3 depending on if the functional $(\boldsymbol{a}_\circ - \boldsymbol{B}_\circ\boldsymbol{f})^\top \boldsymbol{M}_\circ \boldsymbol{q}_\circ$ is either negative or positive. Since η is a positive coefficient, it has no impact on the Lyapunov argument and can simply be omitted, but for the sake of consistency we prefer to keep it in the statements.

Plugging (55) in (47) and taking time derivative lead to

$$\dot{\boldsymbol{q}}_\circ(t) = \overline{\boldsymbol{a}_\circ} - \boldsymbol{B}_\circ \dot{\boldsymbol{f}}^{\mathrm{tot}}(t). \tag{61}$$

Note in (51) that the entry of $\widehat{\boldsymbol{f}(t)}$ corresponding to link ij and class d specifies the number of d-classes the link will send *per unit time* if it is activated at time t. Then $\boldsymbol{f}(t)$ identifies the vector of rate of actual-transmissions realized at time t. Now assume that the entry of $\boldsymbol{f}(t)$ corresponding to link ij and class d at time t is equal to $x \geqslant 0$, meaning that at time t the link transmits x number of d-classes per unit time. Then it should be obvious that the same entry of $\dot{\boldsymbol{f}}^{\mathrm{tot}}(t)$ at time t must be also equal to x. This can be seen more formally from

$$\dot{\boldsymbol{f}}^{\mathrm{tot}}(t) = \lim_{\delta\to 0} \frac{\boldsymbol{f}^{\mathrm{tot}}(t+\delta) - \boldsymbol{f}^{\mathrm{tot}}(t)}{\delta} = \boldsymbol{f}(t) \tag{62}$$

in light of $\lim_{\delta\to 0} \boldsymbol{f}^{\mathrm{tot}}(t+\delta) = \boldsymbol{f}^{\mathrm{tot}}(t) + \delta \boldsymbol{f}(t)$.

On the fictitious graph, when the system in steady-state condition, i.e. when $\dot{\boldsymbol{q}}^\star_\circ(t) = 0$, we obtain $\overline{\boldsymbol{a}_\circ} = \vec{\boldsymbol{L}}_\circ \boldsymbol{q}^\star_\circ = \boldsymbol{B}_\circ \boldsymbol{f}^\star$ in which $\boldsymbol{f}^\star$ represents the time-independent stationary value of $\boldsymbol{f}^\star(t)$. Plugging this and (62) in (61) yields

$$\dot{\boldsymbol{q}}_\circ(t) = \boldsymbol{B}_\circ \boldsymbol{f}^\star - \boldsymbol{B}_\circ \boldsymbol{f}(t). \tag{63}$$

Substituting (63) for $\dot{\boldsymbol{q}}_\circ(t)$ in (60) yields

$$\eta^{-1}\,\dot{Y}(t) \leqslant \big(\boldsymbol{B}_\circ \boldsymbol{f}^\star - \boldsymbol{B}_\circ \boldsymbol{f}(t)\big)^\top \, \overline{\boldsymbol{M}}_\circ \big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big).$$

Then exploiting the equality (25) of Lem. 1 leads to

$$\eta^{-1}\,\dot{Y}(t) \leqslant \big(\boldsymbol{f}^\star - \boldsymbol{f}(t)\big)^\top \, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \big(\boldsymbol{q}_\circ(t) - \boldsymbol{q}^\star_\circ\big).$$

Multiply two sides by 2, add and subtract $\boldsymbol{f}(t)^\top \boldsymbol{f}(t) + \boldsymbol{f}^{\star\top}\boldsymbol{f}^\star$ to the right-hand side, and expand the parentheses to get

$$\begin{aligned}
2\,\eta^{-1}\,\dot{Y}(t) \leqslant &-\big(2\,\boldsymbol{f}(t)^\top\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t) - \boldsymbol{f}(t)^\top \boldsymbol{f}(t)\big) && \text{(64a)}\\
&+\big(2\,\boldsymbol{f}^{\star\top}\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t) - \boldsymbol{f}^{\star\top} \boldsymbol{f}^\star\big) && \text{(64b)}\\
&-\big(2\,\boldsymbol{f}^{\star\top}\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ - \boldsymbol{f}^{\star\top} \boldsymbol{f}^\star\big) && \text{(64c)}\\
&+\big(2\,\boldsymbol{f}(t)^\top\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ - \boldsymbol{f}(t)^\top \boldsymbol{f}(t)\big) && \text{(64d)}
\end{aligned}$$

On the right-hand side of (64), first observe that (64a) and (64b) are characterized on the wireless network given queue backlogs $\boldsymbol{q}_\circ(t)$, while (64c) and (64d) are characterized on the fictitious graph given temperatures $\boldsymbol{q}^\star_\circ$. Recalling the D functional of (54) evaluated at $\boldsymbol{\rho}(t) = \overline{\boldsymbol{\rho}}$ of (55), then observe that (64a) equals $-D(\boldsymbol{f}, \beta, t)$ yielded by the HD forwarding $\boldsymbol{f}(t)$ and (64b) equals $D(\boldsymbol{f}^\star, \beta, t)$ yielded by $\boldsymbol{f}^\star$. On the wireless network, Cor. 2 asserts that given current queue backlogs $\boldsymbol{q}_\circ(t)$, the D functional yielded by HD forwarding is larger than that yielded by any alternative forwarding which respects link capacity constraints. We assumed, on the other hand, that under $\boldsymbol{f}^\star$ the sum of multiclass heat flows remain lower than average link capacities, and so $\boldsymbol{f}^\star$ respects link capacity constraints. This implies that $-D(\boldsymbol{f}, \beta, t) + D(\boldsymbol{f}^\star, \beta, t) \leqslant 0$, meaning that on the right-hand side of (64) the summation of (64a) and (64b) is nonpositive. Therefore,

$$\begin{aligned}
2\,\eta^{-1}\,\dot{Y}(t) \leqslant &-\big(2\,\boldsymbol{f}^{\star\top}\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ - \boldsymbol{f}^{\star\top} \boldsymbol{f}^\star\big) && \text{(65a)}\\
&+\big(2\,\boldsymbol{f}(t)^\top\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ - \boldsymbol{f}(t)^\top \boldsymbol{f}(t)\big) && \text{(65b)}
\end{aligned}$$

Now we show that on the fictitious graph, the functional $H(\boldsymbol{f}) := 2\,\boldsymbol{f}^\top\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ - \boldsymbol{f}^\top \boldsymbol{f}$ is maximized at $\boldsymbol{f} = \boldsymbol{f}^\star$ subject to link directionality. The statement easily follows from

$$H(\boldsymbol{f}) = \sum\nolimits_{ij\in\mathcal{E}} \sum\nolimits_{d\in\mathcal{K}} 2\,\overline{\phi_{ij}}\, q^{\star(d)}_{ij} f^{(d)}_{ij} - \big(f^{(d)}_{ij}\big)^2$$

where link directionality entails $f^{(d)}_{ij} \geqslant 0$ for every link ij and every class d. Then to maximize $H(\boldsymbol{f})$ one needs to assign $f^{(d)}_{ij} = 0$ whenever $q^{\star(d)}_{ij} \leqslant 0$, and $f^{(d)}_{ij} = \overline{\phi_{ij}}\, q^{\star(d)}_{ij}$ otherwise. Translating this into matrix form, the maximizing $\boldsymbol{f}$ is obtained as $\boldsymbol{f}^{\mathrm{opt}} = \overline{\boldsymbol{\Phi}} \max\big\{\boldsymbol{0}, \boldsymbol{B}_\circ^\top \boldsymbol{q}^\star_\circ\big\}$. Comparing the latter with (56), it is seen that $\boldsymbol{f}^\star = \boldsymbol{f}^{\mathrm{opt}}$ for $\boldsymbol{\Sigma} = \overline{\boldsymbol{\Phi}}$.

On the right-hand side of (65), observe that (65a) equals $-H(\boldsymbol{f}^\star)$ and (65b) equals $H(\boldsymbol{f}(t))$. By the argument above, given temperatures $\boldsymbol{q}^\star_\circ$ on the fictitious graph, $H(\boldsymbol{f}^\star)$ yielded by $\boldsymbol{f}^\star$ is larger than that yielded by any other alternative including $\boldsymbol{f}(t)$. Therefore, $-H(\boldsymbol{f}^\star) + H(\boldsymbol{f}(t)) \leqslant 0$, meaning that the summation of (65a) and (65b) is nonpositive, which leads to $2\,\eta^{-1}\,\dot{Y}(t) \leqslant 0$. As η takes the value either 1 or 3, the latter is equivalent to $\dot{Y}(t) \leqslant 0$.

Consider a joint solution $\big(\boldsymbol{q}_\circ(t), \boldsymbol{q}^\star_\circ\big)$ to the HD fluid model (47)–(55) with $\beta \in [0,1]$ and to the generalized graph heat equations (56)–(57) with $\sigma_{ij} = \overline{\phi_{ij}}$. Let Ω be the largest invariant set in the set of all points for which $\dot{Y}(t) = 0$. Since $Y(t)$ is a radially unbounded, positive definite function with $\dot{Y}(t) \leqslant 0$, LaSalle's Invariance Principle implies that every joint solution $\big(\boldsymbol{q}_\circ(t), \boldsymbol{q}^\star_\circ\big)$ asymptotically converges to Ω. Now we show that Ω contains only the point $\big(\boldsymbol{q}^\star_\circ, \boldsymbol{q}^\star_\circ\big)$.

If $\dot{Y}(t) = 0$, then as $\overline{\boldsymbol{M}}_\circ{}^\top + \overline{\boldsymbol{M}}_\circ$ is a symmetric positive-definite matrix, (59) entails $\boldsymbol{q}_\circ(t) = \boldsymbol{q}^\star_\circ$ and/or $\dot{\boldsymbol{q}}_\circ(t) = 0$. Recall that $\boldsymbol{q}^\star_\circ$ is the stationary value of $\boldsymbol{q}^\star_\circ(t)$, which means $\dot{\boldsymbol{q}}^\star_\circ = 0$. Thus also $\boldsymbol{q}_\circ(t) = \boldsymbol{q}^\star_\circ$ leads to $\dot{\boldsymbol{q}}_\circ(t) = 0$. Now if $\dot{\boldsymbol{q}}_\circ(t) = 0$, then (63) entails $\boldsymbol{B}_\circ \boldsymbol{f}(t) = \boldsymbol{B}_\circ \boldsymbol{f}^\star$, which implies $\boldsymbol{B}_\circ \boldsymbol{f}(t) \to \boldsymbol{B}_\circ \boldsymbol{f}^\star$ as $t \to \infty$. Noting that $\boldsymbol{f}^\star$ is not in the null space of $\boldsymbol{B}_\circ$ as $\boldsymbol{B}_\circ \boldsymbol{f}^\star = \overline{\boldsymbol{a}_\circ} \neq 0$, it equally implies that

$$\lim\nolimits_{t\to\infty} \boldsymbol{f}(t) \to \boldsymbol{f}^\star. \tag{66}$$

It also implies that $\boldsymbol{B}_\circ \boldsymbol{f}^{\rm tot}(t) \to \boldsymbol{B}_\circ \boldsymbol{f}^{\star\rm tot}$ as $t \to \infty$. Now let us take a limit of $t \to \infty$ from (47) to obtain

$$\lim_{t\to\infty} \boldsymbol{q}_\circ(t) = \boldsymbol{q}_\circ(0) + \lim_{t\to\infty} \big(\boldsymbol{a}^{\rm tot}_\circ(t) - \boldsymbol{B}_\circ \boldsymbol{f}^{\rm tot}(t)\big).$$

Observe that $\boldsymbol{a}^{\rm tot}_\circ(t)$ is the same for both of the HD fluid model on wireless network and the generalized heat equation on fictitious graph. Then the latter equation is equivalent to

$$\begin{aligned}\lim_{t\to\infty} \boldsymbol{q}_\circ(t) &= \boldsymbol{q}_\circ(0) + \lim_{t\to\infty} \big(\boldsymbol{q}^\star_\circ(t) - \boldsymbol{q}^\star_\circ(0)\big)\\ &= \lim_{t\to\infty} \boldsymbol{q}^\star_\circ(t) + \big(\boldsymbol{q}_\circ(0) - \boldsymbol{q}^\star_\circ(0)\big)\\ &= \boldsymbol{q}^\star_\circ + \big(\boldsymbol{q}_\circ(0) - \boldsymbol{q}^\star_\circ(0)\big)\end{aligned}$$

Note that the system (56)–(57) reaches its steady-state condition independent of its initial condition. Thus as we concern only about steady-state value $\boldsymbol{q}^\star_\circ$, the initial value $\boldsymbol{q}^\star_\circ(0)$ can be taken arbitrarily. Then taking $\boldsymbol{q}^\star_\circ(0) = \boldsymbol{q}_\circ(0)$ leads to

$$\lim\nolimits_{t\to\infty} \boldsymbol{q}_\circ(t) \to \boldsymbol{q}^\star_\circ\,. \tag{67}$$

Therefore, if $\dot{Y}(t) = 0$, then (67). This entails that the invariant set Ω has only one element $\big(\boldsymbol{q}_\circ(t) = \boldsymbol{q}^\star_\circ\,,\, \boldsymbol{q}^\star_\circ\big)$, which ensures that the limits (66) and (67) are globally valid. ■

Remark 4: The assumption that the multiclass heat flow (56) satisfies the link capacity constraints (23) indeed guarantees that the sum of multiclass heat flows on each link remains lower than the expected average of link capacity $\overline{\mu_{ij}}$. Thus it guarantees that for a given arrival rate vector $\overline{\boldsymbol{a}_\circ}$, it is possible at all to stabilize the network such that the fluid limit asymptotically follows the generalized graph heat equations. In essence, the assumption imposes a mild constraint in the sense that it can be satisfied by most practical networks where the networks are not in heavy traffic condition. At the same time, notice that in heavy traffic condition, there is indeed no room for minimizing a routing cost because technically, every existing path has to carry a data flow near to its maximum capacity. Also worth to observe that the assumption intrinsically deals with network topology, arrival rates, link capacities, and link cost-factors, meaning that despite being mild in practice, it bears a theoretically deep and comprehensive concept.

VII. Cost Minimizing Routing Policy

To establish the second pillar of HD Pareto optimality, we show in this section, via Dirichlet Principle, that the Dirichlet routing cost $\overline{R}$ in (2) reaches its minimum feasible value under HD with $\beta = 1$. In fact, we prove a more general result showing that HD with any $\beta \in [0, 1]$ solves the following β-dependent optimization problem:

$$\begin{aligned}&\text{Minimize:} && \sum_{ij\in\mathcal{E}} \sum_{d\in\mathcal{K}} \overline{\big(f^{(d)}_{ij}\big)^2 / \phi_{ij}}\\ &\text{Subject to:} && \text{Throughput optimality}\end{aligned} \tag{68}$$

where for $\beta = 1$ we get $\phi_{ij} = 1/\rho_{ij}$ that recovers the problem of Dirichlet routing cost minimization in (2).

A. Classic Dirichlet Principle

In steady-state conduction, the amount of heat entering any region of an object is equal to the amount of heat leaving out the region. Thus while partial derivatives of temperature with respect to space may either be zero or have nonzero values, all time derivatives of temperature at any point are uniformly zero. This leads to the classic Poisson equation

$$\operatorname{div}\big(\sigma(\boldsymbol{x})\, \nabla Q(\boldsymbol{x})\big) + A(\boldsymbol{x}) = 0$$

which formulates the stationary heat transfer by substituting zero for the time derivative of temperature in (39). Then Dirichlet Principle states that the Poisson equation has a unique solution which minimizes the Dirichlet energy

$$E\big(Q(\boldsymbol{x})\big) := \int_{\mathcal{M}} \Big(\frac{1}{2}\, \sigma \|\nabla Q(\boldsymbol{x})\|^2 - Q(\boldsymbol{x}) A(\boldsymbol{x})\Big) \mathrm{d}\boldsymbol{x}$$

among all twice differentiable functions $Q(\boldsymbol{x})$ that respect the boundary conditions on $\partial\mathcal{M}$.

B. Dirichlet Principle on Undirected Graphs

To derive the combinatorial analogue of Poisson equation on *undirected* graphs, one identifies classic div with the boundary operator $\boldsymbol{B}$, and classic gradient ∇ with the minus[3] of coboundary operator $-\boldsymbol{B}^\top$. Fixing $q_d(t) = 0$, we obtain

$$-\boldsymbol{L}^{(d)}_\circ \boldsymbol{q}^{(d)}_\circ + \boldsymbol{a}^{(d)}_\circ = \boldsymbol{0} \tag{69}$$

which correctly realizes (44) in steady-state condition. Like the classic case, this equation has a unique solution which minimizes the combinatorial Dirichlet energy

$$E(\boldsymbol{q}^{(d)}_\circ) := \frac{1}{2}\, \boldsymbol{q}^{(d)\top}_\circ \boldsymbol{L}^{(d)}_\circ \boldsymbol{q}^{(d)}_\circ - \boldsymbol{q}^{(d)\top}_\circ \boldsymbol{a}^{(d)}_\circ. \tag{70}$$

The proof of Dirichlet Principle is much simpler in the combinatorial case. In fact, since $\boldsymbol{L}^{(d)}_\circ$ is positive definite, $E(\boldsymbol{q}^{(d)}_\circ)$ is convex and thus has a minimum at the critical point, where its first order variation vanishes. This readily leads to the combinatorial Poisson equation (69).

C. Dirichlet Principle on Directed Graphs

Essentially, the Poisson equation on a *directed* graph should capture the steady-state behavior of combinatorial heat diffusion process established as (46). This leads to

$$-\vec{\boldsymbol{L}}^{(d)}_\circ \boldsymbol{q}^{(d)}_\circ + \boldsymbol{a}^{(d)}_\circ = \boldsymbol{0}. \tag{71}$$

The difficulty arises from the fact that contrary to the linear Laplacian $\boldsymbol{L}^{(d)}_\circ$ on undirected graphs, the $\vec{\boldsymbol{L}}^{(d)}_\circ$ here is a

[3] In vector calculus, the gradient of a scalar field is positive in the direction of increase of the field. On a graph, on the other hand, the gradient of a node variable is taken positive in the direction of decrease of the variable. By the same reason, the classic Laplace operator is a negative semi-definite operator, while the graph Laplacian is a positive semi-definite matrix.

nonlinear operator. Thus the easy way of proving Dirichlet Principle on undirected graphs is ceased to exist on here. Observe that in (70), the directional derivative of $\frac{1}{2}\,\boldsymbol{q}_\circ^{(d)\top}\boldsymbol{L}_\circ^{(d)}\boldsymbol{q}_\circ^{(d)}$ along $\boldsymbol{q}_\circ^{(d)}$ is simply $\boldsymbol{L}_\circ^{(d)}\boldsymbol{q}_\circ^{(d)}$, as appeared in (69), by the reason that $\boldsymbol{L}_\circ^{(d)}$ is a symmetric positive definite matrix. On a directed graph, on the other hand, $\vec{\boldsymbol{L}}_\circ^{(d)}$ is an operand-dependent nonlinear operator which retains neither linearity nor symmetricity properties. Therefore, we can not claim that $\vec{\boldsymbol{L}}_\circ^{(d)}\boldsymbol{q}_\circ^{(d)}$ in (71) is the directional derivative of $\frac{1}{2}\,\boldsymbol{q}_\circ^{(d)\top}\vec{\boldsymbol{L}}_\circ^{(d)}\boldsymbol{q}_\circ^{(d)}$. Nevertheless, by the next theorem we extend the concept of Dirichlet Principle to directed graphs with nonlinear Laplacian.

Theorem 6: On a directed graph, the Poisson equation (71) has a unique solution which minimizes combinatorial Dirichlet-like energy functional

$$\vec{E}(\boldsymbol{q}_\circ^{(d)}) := \frac{1}{2}\,\boldsymbol{q}_\circ^{(d)\top}\vec{\boldsymbol{L}}_\circ^{(d)}\boldsymbol{q}_\circ^{(d)} - \boldsymbol{q}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)} \tag{72}$$

where $\vec{\boldsymbol{L}}_\circ^{(d)}$ is the nonlinear Dirichlet Laplacian as in (46).

Proof: Using (46) one can expand $\vec{\boldsymbol{L}}_\circ^{(d)}$ to obtain

$$\vec{E}(\boldsymbol{q}_\circ^{(d)}) = \frac{1}{2}\left(\boldsymbol{q}_\circ^{(d)\top}\boldsymbol{B}_\circ^{(d)}\right)^{+}\operatorname{diag}(\boldsymbol{\sigma})\big(\boldsymbol{B}_\circ^{(d)\top}\boldsymbol{q}_\circ^{(d)}\big)^{+} - \boldsymbol{q}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)}.$$

Note that each entry of $\boldsymbol{B}_\circ^{(d)\top}\boldsymbol{q}_\circ^{(d)} \in \mathbb{R}^{|\mathcal{E}|}$ represents the temperature-difference along the corresponding edge.

Let $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ be the $\vec{E}$ minimizing solution. We partition $\boldsymbol{B}_\circ^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ according to the sign of its entries. Accordingly, the incidence matrix $\boldsymbol{B}_\circ^{(d)}$ gets partitioned as

$$\boldsymbol{B}_\circ^{(d)} = \left(\boldsymbol{B}_{\circ+}^{(d)} \mid \boldsymbol{B}_{\circ\varnothing}^{(d)} \mid \boldsymbol{B}_{\circ-}^{(d)}\right) \tag{73}$$

where $\boldsymbol{B}_{\circ+}^{(d)}$, $\boldsymbol{B}_{\circ\varnothing}^{(d)}$, and $\boldsymbol{B}_{\circ-}^{(d)}$ respectively denote the partitions of $\boldsymbol{B}_\circ^{(d)}$ which contain the incidence information of edges with positive, zero, and negative temperature-difference. Likewise, the edge weights get partitioned into $\boldsymbol{\sigma} = \big(\boldsymbol{\sigma}_+ \,|\, \boldsymbol{\sigma}_\varnothing \,|\, \boldsymbol{\sigma}_-\big)$. Then $\vec{E}(\boldsymbol{q}_\circ^{(d)})$ at $\boldsymbol{q}_\circ^{(d)} = \overset{*}{\boldsymbol{q}}_\circ^{(d)}$ can be shown as

$$\begin{aligned}\vec{E}(\overset{*}{\boldsymbol{q}}_\circ^{(d)}) = &-\overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)} + \frac{1}{2}\,\overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{B}_{\circ+}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_+)\,\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\\ &+\frac{1}{2}\big(\overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{B}_{\circ\varnothing}^{(d)}\big)^{+}\operatorname{diag}(\boldsymbol{\sigma}_\varnothing)\big(\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} &&\text{(74a)}\\ &+\frac{1}{2}\big(\overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{B}_{\circ-}^{(d)}\big)^{+}\operatorname{diag}(\boldsymbol{\sigma}_-)\big(\boldsymbol{B}_{\circ-}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} &&\text{(74b)}\end{aligned}$$

where we used $\big(\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}$. Observe that (74b) is strongly zero due to the $(\cdot)^+$ operation. Also (74a) vanishes because $\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)} = \boldsymbol{0}$. Therefore,

$$\vec{E}(\overset{*}{\boldsymbol{q}}_\circ^{(d)}) = \frac{1}{2}\,\overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{B}_{\circ+}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_+)\,\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)} - \overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)}.$$

Since the network is stabilizable, for every heat source there must exist at least one directed path to the destination d. Then $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ must be entrywise finite because it minimizes functional (72). This implies that under the optimum heat distribution $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$, the set of edges of positive temperature-difference must build a connected graph with node d. Therefore, $\boldsymbol{B}_{\circ+}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_+)\,\boldsymbol{B}_{\circ+}^{(d)\top}$ is a symmetric positive-definite matrix (see the proof of Lem. 1). It turns out that $\vec{E}(\overset{*}{\boldsymbol{q}}_\circ^{(d)})$ is convex and thus has a minimum at the critical point, where its first order variation vanishes. Thus $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ must satisfy

$$\boldsymbol{a}_\circ^{(d)} = \boldsymbol{B}_{\circ+}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_+)\,\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}. \tag{75}$$

In view of $\big(\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ and $\big(\boldsymbol{B}_{\circ-}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \big(\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \boldsymbol{0}$, (75) can be stated as

$$\begin{aligned}\boldsymbol{a}_\circ^{(d)} = &\,\boldsymbol{B}_{\circ+}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_+)\big(\boldsymbol{B}_{\circ+}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+}\\ &+\boldsymbol{B}_{\circ\varnothing}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_\varnothing)\big(\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+}\\ &+\boldsymbol{B}_{\circ-}^{(d)}\operatorname{diag}(\boldsymbol{\sigma}_-)\big(\boldsymbol{B}_{\circ-}^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+}.\end{aligned}$$

Using matrix form (73), the latter is equivalent to

$$\boldsymbol{a}_\circ^{(d)} = \boldsymbol{B}_\circ^{(d)}\operatorname{diag}(\boldsymbol{\sigma})\big(\boldsymbol{B}_\circ^{(d)\top}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \vec{\boldsymbol{L}}_\circ^{(d)}\overset{*}{\boldsymbol{q}}_\circ^{(d)}$$

which recovers the Poisson equation (71).

Next, take a variation along the links with zero temperature-difference and let $\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$ be the new heat distribution. By sufficiently small variation, the links which have either positive or negative temperature-difference under $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ distribution keep the same sign of temperature-difference under $\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$ distribution too, but those with zero temperature-difference under $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ may go into different conditions under $\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$. Therefore, we still have the equalities $\big(\boldsymbol{B}_{\circ+}^{(d)\top}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \boldsymbol{B}_{\circ+}^{(d)\top}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$ and $\big(\boldsymbol{B}_{\circ-}^{(d)\top}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}\big)^{+} = \boldsymbol{0}$, but the entries of $\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$ may not be zero anymore. Let us define a positive scalar

$$\begin{aligned}\varepsilon &:= \limsup_{\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}\to\overset{*}{\boldsymbol{q}}_\circ^{(d)}} \frac{1}{2}\left(\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)\top}\vec{\boldsymbol{L}}_\circ^{(d)}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)} - \overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\vec{\boldsymbol{L}}_\circ^{(d)}\overset{*}{\boldsymbol{q}}_\circ^{(d)}\right)\\ &= \limsup_{\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}\to\overset{*}{\boldsymbol{q}}_\circ^{(d)}} \frac{1}{2}\big(\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{B}_{\circ\varnothing}^{(d)}\big)^{+}\operatorname{diag}(\boldsymbol{\sigma}_\varnothing)\big(\boldsymbol{B}_{\circ\varnothing}^{(d)\top}\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}\big)^{+}\end{aligned}$$

where sup is taken over all directions of $\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}$. Then defining $\delta := \big|\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)} - \overset{*}{\boldsymbol{q}}_\circ^{(d)\top}\boldsymbol{a}_\circ^{(d)}\big|$ leads to

$$\vec{E}(\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}) - \vec{E}(\overset{*}{\boldsymbol{q}}_\circ^{(d)}) > \varepsilon - \delta.$$

Thus for every $\varepsilon > 0$, picking $|\,\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)} - \overset{*}{\boldsymbol{q}}_\circ^{(d)}|$ sufficiently small such that $\delta < \varepsilon$ guarantees that $\vec{E}(\overset{\diamond}{\boldsymbol{q}}_\circ^{(d)}) > \vec{E}(\overset{*}{\boldsymbol{q}}_\circ^{(d)})$. Therefore, the optimum heat distribution is robust against variation, which together with the convexity of $\vec{E}$ around its critical point entail that $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ is the unique solution that minimizes $\vec{E}(\boldsymbol{q}_\circ^{(d)})$. As a consequence, $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ must be the unique root for the first order variation of $\vec{E}$ as well. This entails the minimizing $\overset{*}{\boldsymbol{q}}_\circ^{(d)}$ to be the unique solution of the Poisson equation (71). ■

Remark 5: Though the Dirichlet Principle on undirected graphs has been known for a long time, its extension to directed graphs is completely new to literature. As a model of heat flow on a directed graph, conceptualize an electrical network with placing a diode on each edge. The current, which mimics the combinatorial heat flux, moves along the negative gradient of voltage, but only under the condition of respecting the diode direction. Another example is a piping network of liquid/gas with a check valve on each line. Again the liquid/gas flows along the negative gradient of pressure, while each check valve allows the flow in only one direction.

D. Dirichlet Routing Cost Minimization

Extending Th. 6 to the multiclass heat diffusion process on directed graphs, the steady-state solution to (57) must minimize the multiclass version of functional (72). More precisely, the generalized Poisson equation

$$-\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star + \overline{\boldsymbol{a}_\circ} = \boldsymbol{0} \tag{76}$$

has a unique solution which minimizes the generalized Dirichlet-like energy functional

$$\vec{E}(\boldsymbol{q}_\circ^\star) = \frac{1}{2}\,\boldsymbol{q}_\circ^{\star\top}\vec{\boldsymbol{L}}_\circ\,\boldsymbol{q}_\circ^\star - \boldsymbol{q}_\circ^{\star\top}\overline{\boldsymbol{a}_\circ} \tag{77}$$

where $\boldsymbol{q}_\circ^\star \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|}$ and $\vec{\boldsymbol{L}}_\circ \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|\times(|\mathcal{V}|-1)|\mathcal{K}|}$ are as defined in (57). Yet the framework of this result is not aligned with what we need for the optimization problem (68). The next theorem resolves this incongruity by showing that minimizing (77) is indeed the dual of minimizing energy dissipation on the directed graph with zero duality gap.

Theorem 7: On a directed graph, minimizing the generalized Dirichlet-like energy functional (77) is equivalent to solving the constrained optimization problem

$$\begin{aligned} \text{Minimize:} \quad & \vec{E}_R(\boldsymbol{f}^\star) := \boldsymbol{f}^{\star\top}\boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star \\ \text{Subject to:} \quad & \boldsymbol{B}_\circ\boldsymbol{f}^\star = \overline{\boldsymbol{a}_\circ} \end{aligned} \tag{78}$$

which represents total energy dissipated on the graph.

Proof: By the principle of duality, the primal optimization problem (78) can be viewed as the Lagrangian dual problem

$$\underset{\boldsymbol{\lambda}\succcurlyeq\boldsymbol{0}}{\text{Maximize:}}\ \inf_{\boldsymbol{f}^\star}\Big(J := \boldsymbol{f}^\star\,\boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star + 2\,\boldsymbol{\lambda}^\top\big(\overline{\boldsymbol{a}_\circ} - \boldsymbol{B}_\circ\boldsymbol{f}^\star\big)\Big) \tag{79}$$

where $\boldsymbol{\lambda}$ is the vector of Lagrange multipliers. Since $\boldsymbol{f}^\star$ is continuously differentiable, so is J, and thus the infimum occurs where the gradient is equal to zero, leading to

$$\nabla_{\boldsymbol{f}^\star} J = 0 \quad \Rightarrow \quad \boldsymbol{B}_\circ^\top\boldsymbol{\lambda} = \boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star.$$

Because $\boldsymbol{f}^\star$ is entrywise nonnegative, the latter is equivalent to $\big(\boldsymbol{B}_\circ^\top\boldsymbol{\lambda}\big)^+ = \boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star$. Multiplying both sides by $\boldsymbol{B}_\circ\boldsymbol{\Sigma}$ and using the definition of $\vec{\boldsymbol{L}}_\circ$ in (57) yield $\vec{\boldsymbol{L}}_\circ\boldsymbol{\lambda} = \boldsymbol{B}_\circ\boldsymbol{f}^\star$ where $\boldsymbol{f}^\star$ represents the steady-state value of $\boldsymbol{f}^\star(t)$. Further, in (57) at the steady-state condition, i.e. $\dot{\boldsymbol{q}}_\circ^\star(t) = \boldsymbol{0}$, we have $\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star = \overline{\boldsymbol{a}_\circ}$ in which $\boldsymbol{q}_\circ^\star$ represents the steady-state value of $\boldsymbol{q}_\circ^\star(t)$. Also from (56) and (57) together we have $\boldsymbol{B}_\circ\boldsymbol{f}^\star = \vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star$, which with $\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star = \overline{\boldsymbol{a}_\circ}$ implies $\boldsymbol{B}_\circ\boldsymbol{f}^\star = \overline{\boldsymbol{a}_\circ}$. Considering this with the previous result of $\vec{\boldsymbol{L}}_\circ\boldsymbol{\lambda} = \boldsymbol{B}_\circ\boldsymbol{f}^\star$, we obtain $\vec{\boldsymbol{L}}_\circ\boldsymbol{\lambda} = \overline{\boldsymbol{a}_\circ}$. In light of Th. 6, on the other hand, the Poisson equation $\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star = \overline{\boldsymbol{a}_\circ}$ has a unique solution. Comparing this with $\vec{\boldsymbol{L}}_\circ\boldsymbol{\lambda} = \overline{\boldsymbol{a}_\circ}$ results in $\boldsymbol{\lambda} = \boldsymbol{q}_\circ^\star \succcurlyeq \boldsymbol{0}$.

In (79), the first term of J is equal to $\boldsymbol{f}^\star\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star$, while $\boldsymbol{\Sigma}^{-1}\boldsymbol{f}^\star = \big(\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ^\star\big)^+$ from (56). Thus the first term of J can be replaced by $\big(\boldsymbol{q}_\circ^{\star\top}\boldsymbol{B}_\circ\big)^+\boldsymbol{\Sigma}\big(\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ^\star\big)^+ = \boldsymbol{q}_\circ^{\star\top}\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star$. Regarding the second term of J, we already showed that $\boldsymbol{B}_\circ\boldsymbol{f}^\star = \vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star$. Thus the second term can be replaced by $2\,\boldsymbol{\lambda}^\top\big(\overline{\boldsymbol{a}_\circ} - \vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star\big)$. Then using the result of $\boldsymbol{\lambda} = \boldsymbol{q}_\circ^\star$ leads to

$$J = -\boldsymbol{q}_\circ^{\star\top}\vec{\boldsymbol{L}}_\circ\boldsymbol{q}_\circ^\star + 2\,\boldsymbol{q}_\circ^{\star\top}\overline{\boldsymbol{a}_\circ}.$$

Substituting the latter for J in the dual problem (79) leads to the minimization problem (77). Furthermore, since the objective function in (79) is obviously convex, the duality gap is zero and thus the dual problem (77) results in the same optimal values that the primal problem (78) does. ■

It is worth to compare Th. 7 with the famous law of minimum dissipative energy on electrical networks [30]. In essence, Th. 7 extends this law to directed graphs, or to resistive-diode networks for that matter. Then the upshot is due to the connection between heat diffusion process on a directed graph and HD fluid limit on a wireless network.

Theorem 8: Consider the generalized graph heat equations (56)–(57), and suppose that the multiclass heat flow (56) satisfies the link capacity constraints (23). Then the HD policy with $\beta \in [0,1]$ solves the β-dependent optimization problem (68). In particular, if the generalized heat flow satisfies the link capacity constraints for $\sigma_{ij} = 1/\overline{\rho_{ij}}$, then the HD policy with $\beta = 1$ solves the Dirichlet routing cost minimization problem as defined in (2).

Proof: In Th. 8, it is shown that under the multiclass heat equations (56)–(57), the stationary value of generalized Dirichlet-like energy $\vec{E}$ is strictly minimized. In Th. 8, on the other hand, it is shown that minimizing $\vec{E}$ is exactly the same as minimizing the stationary value of graph energy dissipation $\vec{E}_R$. Then the proof immediately follows from Th. 5 which states that the limiting dynamics of HD with $\beta \in [0,1]$ comply with the deterministic, multiclass heat equations (56)–(57) with $\sigma_{ij} = \overline{\phi_{ij}}$. Note that in light of Th. 5, every expected time average value on the stochastic wireless network governed by HD policy follows the corresponding stationary value obtained from multiclass heat equations on the suitably weighted directed graph. In particular, the objective function in (68) complies with the graph energy dissipation $\vec{E}_R$ in (78) for $\sigma_{ij} = \overline{\phi_{ij}}$. By the same token, the Dirichlet routing cost $\overline{R}$ in (2) follows $\vec{E}_R$ for $\sigma_{ij} = 1/\overline{\rho_{ij}}$. ■

Remark 6: This is the first time a network-layer routing policy asserts the strict minimization of a routing penalty subject to network stability. As an unpractical solution, it can be shown that there exists a stationary randomized algorithm which solves this minimization problem, but it requires a full knowledge of arrival statistics and channel state probabilities, and also involves solving an intractable dynamic programming [2], [3]. Thus far, the only practical approach to this problem has been the V-parameter BP policy of Sec. II that can get close to the optimal solution, but at the expense of unbounded increase in the average delay, which causes network instability. We remark, however, that the V-parameter approach holds for more general cost functions, and is not restricted to the particular structure of Dirichlet routing cost.

VIII. Pareto Optimal Performance

Minimizing average delay and minimizing average routing cost are often conflicting objectives, meaning that as one decreases the other increases. This leads to a natural multi-objective optimization framework. Then the ideal operating points are on the *Pareto boundary* since they correspond to equilibria from which any deviation will lead to the performance degradation in at least one objective.

A. Strong Pareto Optimality for Nonuniform Link Costs

We have shown that HD with $\beta = 0$ minimizes the average network delay $\overline{Q}$ within the class of all routing algorithms which depend only on current queue congestion and current channel states (optimization problem (3)). We have also shown that under a mild assumption on link capacities, HD with $\beta = 1$ strictly minimizes the Dirichlet routing cost $\overline{R}$ among all possible routing algorithms and subject to network stability (optimization problem (2)). Now consider the region built on the joint variables $(\overline{Q}, \overline{R})$ in which $\overline{Q}$ is achievable by the aforementioned class of routing algorithms. Assume that this region has a convex Pareto boundary. Then the next theorem claims that the HD with $\beta \in [0,1]$ operates on this Pareto boundary, meaning that it solves the multi-objective optimization problem (4).

Theorem 9: Define $\overline{Q}$ as in (3) and $\overline{R}$ as in (2). Consider the region of all joint variables $(\overline{Q}\,,\overline{R})$ with $\overline{Q}$ obtained by a routing algorithm which acts based only on current queue backlogs and current channel states, and suppose that this region has a convex Pareto boundary. Also consider the generalized graph heat equations (56)–(57) with $\sigma_{ij} = 1/\overline{\rho_{ij}}$, and suppose that the multiclass heat flow (56) satisfies the link capacity constraints (23). Then the HD policy with $\beta \in [0,1]$ operates on the Pareto boundary of $(\overline{Q}\,,\overline{R})$ region.

Proof: Since the region $(\overline{Q}\,,\overline{R})$ has a convex Pareto boundary, it is known that the entire boundary can be reached by the weighted-sum method [31], in which the Pareto front is obtained by changing a weight between the two objective functions. Then the proof easily follows with observing that the HD policy minimizes $\overline{Q}$ at $\beta=0$, minimizes $\overline{R}$ at $\beta=1$, and varies the relative importance between these two objectives by changing the convexity factor β between 0 and 1. ■

Remark 7: This is the first time a network-layer routing policy asserts Pareto optimal performance with respect to average delay and an average routing cost subject to throughput optimality. To the best of our knowledge, addressing such a multi-objective optimization problem at the network layer, even without throughput optimality constraint, is completely new. Also worth to note that in the case of non-convex Pareto boundary, HD with $\beta \in [0,1]$ still covers the operating points on convex parts of the boundary, although some Pareto optimal operating points lie on non-convex parts [31].

B. Weak Pareto Optimality for Unit Link Costs

When all links are of *unit* cost-factor, we get $\phi_{ij} = 1$ for all β, and thus HD policy has the same performance independent of β. Considering this observation together with Th. 3 implies that the average delay $\overline{Q}$ must be minimized for all $\beta \in [0,1]$. Considering it together with Th. 8, on the other hand, implies that the Dirichlet routing cost $\overline{R}$ must also be minimized for all $\beta \in [0,1]$. Holding these two requirements at the same time entails that $\overline{Q}$ and $\overline{R}$ must be minimized together, which equivalently means that the Pareto boundary of $(\overline{Q}\,,\overline{R})$ region must shrink into one point. Such an operating point is called *weakly* Pareto optimal in the sense that no tradeoff is allowed as it is impossible to strictly improve at least one objective. The upshot is formalized by the next corollary.

Corollary 3: Consider $(\overline{Q}\,,\overline{R})$ region as defined in Th. 9. Suppose that the generalized heat flow (56) on an unweighted multiclass graph satisfies the link capacity constraints (23). Then under unit cost-factors for all links, the Pareto boundary of $(\overline{Q}\,,\overline{R})$ region reduces to one point, and the HD policy operates at this point for all $\beta \in [0,1]$. ◇

Figure 7 depicts the feasible region on the joint variables $(\overline{Q}\,,\overline{R})$ under unit cost-factors for all links. In contrast with Fig. 1, it also emphasizes the HD operation at weakly Pareto optimal point for all $\beta \in [0,1]$ in comparison with V-parameter BP performance for $V \in [0,\infty)$.

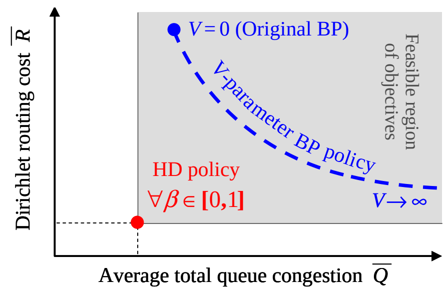

Fig. 7. Graphical description of weak Pareto boundary under unit cost-factors for all links, depicting the HD operating point compared with BP performance.

IX. Conclusion

To summarize, we have introduced the new Heat-Diffusion (HD) routing policy for network layer. Within the class of all routing algorithms which perform based only on current queue congestion and current channel states, HD obtains a Pareto optimal tradeoff between average delay and Dirichlet routing cost. In particular, in the above-mentioned class, HD achieves the minimum average delay. Further, under a mild assumption on link capacities, HD strictly minimizes the Dirichlet routing cost among all possible routing algorithms (not restricted to the above-mentioned class). Besides these, HD is strongly connected to the classical world of Thermodynamics that we believe opens the door to a rich array of theoretical techniques that could be applied in future work to analyze and optimize communication in stochastic networks constrained with channel interference.

Acknowledgment

The authors would like to thank Professor Michael Neely at USC, for his helpful feedback and discussions on this work.

Appendix

A. Proof of Theorem 1

Under the assumption that each link may transmit packets from only one class at each timeslot, the maximization problem (14) can be solved in two steps: First, for each link find

$$[d_{ij}^{\text{opt}}(n), f_{ij}^{\text{opt}}(n)] = \arg\max\nolimits_{d,f}\ 2\,\phi_{ij}(n)\,q_{ij}^{(d)}(n) f - f^2 \quad (80)$$

$$w_{ij}^{\text{opt}}(n) := 2\,\phi_{ij}(n) q_{ij}^{(d_{ij}^{\text{opt}})}(n) f_{ij}^{\text{opt}}(n) - f_{ij}^{\text{opt}}(n)^2. \quad (81)$$

Second, use the link weights $w_{ij}^{\text{opt}}(n)$ to solve

$$\boldsymbol{\pi}^{\text{opt}}(n) = \arg\max_{\boldsymbol{\pi}\in\Pi} \sum\nolimits_{ij\in\mathcal{E}} \pi_{ij}\ w_{ij}^{\text{opt}}(n). \quad (82)$$

Comparing (81) with the HD weighing (11), and (82) with the HD scheduling (12), the proof is fulfilled by showing that the solution to (82) follows $d_{ij}^*(n)$ in (9) and $\widehat{f_{ij}}(n)$ in (10). Using the first derivative test on $2\,\phi_{ij}(n)\,q_{ij}^{(d)}(n) f - f^2$ with respect to $q_{ij}^{(d)}(n)$, it is seen that $d_{ij}^{\text{opt}}(n) = \arg\max_d q_{ij}^{(d)}(n)$ which equals $d_{ij}^*(n)$ in (9). By taking the first derivative with respect to f, on the other hand, we obtain $f_{ij}^{\text{opt}}(n) = \phi_{ij}(n)\,q_{ij}^{(d^*)}(n)$. Adding to this the link constraints that $f_{ij}(n)$ must be nonnegative and at most equal to the link capacity, results in $f_{ij}^{\text{opt}}(n) = \min\{\phi_{ij}(n)\,q_{ij}^{(d^*)}(n)^+,\ \mu_{ij}(n)\}$ which follows $\widehat{f_{ij}}(n)$ in (10) and concludes the proof. ■

B. Proof of Lemma 1

First we show that the matrix $\boldsymbol{B}_\circ \boldsymbol{\Sigma} \boldsymbol{B}_\circ^\top$ is symmetric positive definite for any diagonal matrix $\boldsymbol{\Sigma}$ with positive diagonal entries. Since the wireless network has a connected topology, it is known that any sub-matrix obtained from its incidence matrix by removing one or more rows has full row rank [32]. This means that $\boldsymbol{B}_\circ^{(d)}$ has full row rank for any $d \in \mathcal{K}$. Accordingly, the $\boldsymbol{B}_\circ$ matrix of (18) has full row rank, implying that the symmetric positive semi-definite matrix $\boldsymbol{B}_\circ \boldsymbol{\Sigma} \boldsymbol{B}_\circ^\top$ has full rank, and so is symmetric positive definite.

As $\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top$ is symmetric positive definite, so is its inverse. Also $\boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ is symmetric positive definite since the $\boldsymbol{\Phi}(n)$ matrix of (24) is diagonal with positive diagonal entries for all n. Thus $\boldsymbol{M}_\circ(n)$ is a product of two symmetric positive definite matrices, and therefore it is similar to $\big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{1/2} \boldsymbol{M}_\circ(n) \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1/2} = \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1/2} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1/2}$. The latter is congruent to $\boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ and so symmetric positive definite. Thus $\boldsymbol{M}_\circ(n)$ has positive eigenvalues for all n, meaning that $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} > 0$ for an arbitrary vector $\boldsymbol{x} \neq \boldsymbol{0}$.

Here we show that $\boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} = \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$ for an arbitrary vector $\boldsymbol{x}$. Since $\boldsymbol{M}_\circ(n) = \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$, and since $\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top$ is a symmetric positive definite matrix, we have $\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} = \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$. Now assume $\boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} \neq \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$. Then since $\boldsymbol{x}$ is an arbitrary vector, the equality $\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} = \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$ implies that both $\boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x}$ and $\boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$ must belong to the null space of $\boldsymbol{B}_\circ$, which results in $\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} = \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x} = \boldsymbol{0}$. But this is in contradiction with the positive definiteness of $\boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$, and therefore the assumption that $\boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{x} \neq \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{x}$ is not correct. ■

C. Proof of Lemma 2

Replacing $\boldsymbol{M}_\circ + \boldsymbol{M}_\circ^\top$ with $2\boldsymbol{M}_\circ + (\boldsymbol{M}_\circ^\top - \boldsymbol{M}_\circ)$, we need to prove that for arbitrary vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{(|\mathcal{V}|-1)|\mathcal{K}|}$,

$$2\, \boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y} + \boldsymbol{x}^\top \big(\boldsymbol{M}_\circ(n)^\top - \boldsymbol{M}_\circ(n)\big)\, \boldsymbol{y} \leqslant \eta\, \boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}$$

for η being either 1 or 3 depending on if $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}$ is either negative or positive. This is equivalent to prove that

$$\boldsymbol{x}^\top \big(\boldsymbol{M}_\circ(n)^\top - \boldsymbol{M}_\circ(n)\big)\, \boldsymbol{y} \leqslant (\eta - 2)\, \boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}.$$

Considering that η switches between 1 and 3 corresponding to the sign of $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}$, it is easy to see that to prove the above inequality, it is sufficient to prove that

$$\big|\boldsymbol{x}^\top \big(\boldsymbol{M}_\circ(n)^\top + \boldsymbol{M}_\circ(n)\big)\, \boldsymbol{y}\big| \leqslant \big|\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\, \boldsymbol{y}\big|.$$

Dropping (n) for brevity, the latter is equivalent to

$$\boldsymbol{x}^\top (\boldsymbol{M}_\circ^\top - \boldsymbol{M}_\circ)\, \boldsymbol{y}\boldsymbol{y}^\top (\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top)\, \boldsymbol{x} \leqslant \boldsymbol{x}^\top \boldsymbol{M}_\circ\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{M}_\circ^\top \boldsymbol{x},$$

which by some matrix algebra can be transformed to

$$\boldsymbol{x}^\top (2\, \boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top)\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{M}_\circ\, \boldsymbol{x} \geqslant 0. \tag{83}$$

First observe that as $\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top$ is a skew-symmetric matrix, $\boldsymbol{z}^\top (\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top)\, \boldsymbol{z} = 0$ for an arbitrary vector $\boldsymbol{z}$. Therefore, $\boldsymbol{z}^\top (2\, \boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top)\, \boldsymbol{z} = \boldsymbol{z}^\top \boldsymbol{M}_\circ\, \boldsymbol{z}$ which shows that $2\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top$ is a quasi-positive matrix in the sense of Lem. 1.

Now we show that for any quasi-positive matrix $\boldsymbol{A}$, there exist two symmetric positive definite matrices $\boldsymbol{L}$ and $\boldsymbol{R}$ such that $\boldsymbol{A} = \boldsymbol{L}\boldsymbol{R}$. As a square matrix, $\boldsymbol{A}$ has a Jordan decomposition as $\boldsymbol{A} = \boldsymbol{V}\boldsymbol{J}\boldsymbol{V}^{-1}$ where $\boldsymbol{J}$ is a Jordan form with positive diagonal entries, as $\boldsymbol{A}$ has all eigenvalues positive. Then choose $\boldsymbol{L} = \boldsymbol{V}\boldsymbol{E}\boldsymbol{V}^\top$ and $\boldsymbol{R} = (\boldsymbol{V}^{-1})^\top \boldsymbol{E}^{-1} \boldsymbol{J} \boldsymbol{V}^{-1}$ where $\boldsymbol{E}$ is a symmetric positive definite matrix commuting with $\boldsymbol{J}$. Since $\boldsymbol{L}$ is congruent to $\boldsymbol{E}$, it has positive eigenvalues, and so it is a symmetric positive definite matrix. Likewise, since $\boldsymbol{R}$ is congruent to $\boldsymbol{E}^{-1}\boldsymbol{J}$, it has positive eigenvalues, and so it is a symmetric positive definite matrix. To see why $\boldsymbol{E}^{-1}\boldsymbol{J}$ has positive eigenvalues, observe that it is similar to $\boldsymbol{E}^{1/2}\boldsymbol{E}^{-1}\boldsymbol{J}\boldsymbol{E}^{-1/2} = \boldsymbol{E}^{-1/2}\boldsymbol{J}\boldsymbol{E}^{-1/2}$ and so congruent to $\boldsymbol{J}$ which has positive eigenvalues.

Turing back to (83), let $\boldsymbol{D} := (2\, \boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top)\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{M}_\circ$. Since both $2\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top$ and $\boldsymbol{M}_\circ$ are quasi-positive, they can be decomposed into the product of two symmetric positive definite matrices as $2\boldsymbol{M}_\circ - \boldsymbol{M}_\circ^\top = \boldsymbol{L}_1\boldsymbol{R}_1$ and $\boldsymbol{M}_\circ = \boldsymbol{L}_2\boldsymbol{R}_2$ that leads to $\boldsymbol{D} = \boldsymbol{L}_1\boldsymbol{R}_1\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{L}_2\boldsymbol{R}_2$. Then $\boldsymbol{D}$ is similar to $\boldsymbol{R}_2^{1/2}\boldsymbol{L}_1\boldsymbol{R}_1\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{L}_2\boldsymbol{R}_2^{1/2}$ and so congruent to $\boldsymbol{L}_1\boldsymbol{R}_1\, \boldsymbol{y}\boldsymbol{y}^\top \boldsymbol{L}_2$. Applying this rule iteratively, it is seen that $\boldsymbol{D}$ is congruent to $\boldsymbol{y}\boldsymbol{y}^\top$, and therefore both of them has the same number of negative eigenvalues. But $\boldsymbol{y}\boldsymbol{y}^\top$ is a symmetric positive semi-definite matrix for any vector $\boldsymbol{y}$ and so has no negative eigenvalue, meaning that $\boldsymbol{D}$ also has no negative eigenvalue. Then it is obvious that $\boldsymbol{x}^\top \boldsymbol{D}\, \boldsymbol{x} \geqslant 0$ which concludes (83). ■

D. Proof of Lemma 3

The time variable (n) is often dropped for brevity. We want to show that at every slot n, HD with $\beta = 0$ maximizes

$$D_P(\boldsymbol{f}) := 2\, \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) - \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{B}_\circ \boldsymbol{f}(n) \tag{84}$$

subject to network constraints. At first we *temporarily* ignore the network constraints including link interference, link capacity, and link directionality. Afterwards, we will add the effect of these constraints into the picture.

Ignoring link interference, the vector space $\mathcal{F}$ formed by the collection of admissible link actual-transmissions $\boldsymbol{f}(n)$ becomes a convex set. To see this, observe that in the absence of link interference, for any two admissible link actual-transmissions $\boldsymbol{f}_1(n), \boldsymbol{f}_2(n) \in \mathcal{F}$ and for any scalar $c \in [0,1]$, the link actual-transmission $c\boldsymbol{f}_1(n) + (1-c)\boldsymbol{f}_2(n)$ is also admissible and so is in $\mathcal{F}$. Then we show that $D_P(\boldsymbol{f})$ is strictly concave on $\mathcal{F}$. Considering any $\boldsymbol{f}_1(n), \boldsymbol{f}_2(n) \in \mathcal{F}$ and any scalar $c \in [0,1]$, we need to show that

$$D_P(c\boldsymbol{f}_1 + (1-c)\boldsymbol{f}_2) > c\, D_P(\boldsymbol{f}_1) + (1-c) D_P(\boldsymbol{f}_2).$$

Expanding the D_P function as defined in (84), and after some elementary matrix algebra, we need to show that

$$c\,(1-c)\big(\boldsymbol{B}_\circ \boldsymbol{f}_1 - \boldsymbol{B}_\circ \boldsymbol{f}_2\big)^\top \big(\boldsymbol{B}_\circ \boldsymbol{f}_1 - \boldsymbol{B}_\circ \boldsymbol{f}_2\big) > 0$$

which is correct for any $c \in [0,1]$ and any $\boldsymbol{f}_1(n)$ and $\boldsymbol{f}_2(n)$ not in the null space of $\boldsymbol{B}_\circ$. Assuming that the destination node d can only receive d-classes, it is a simple network flow argument to show that $\boldsymbol{B}_\circ \boldsymbol{f} \neq \boldsymbol{0}$ for any admissible $\boldsymbol{f}(n) \in \mathcal{F}$, meaning that $\boldsymbol{f}(n)$ can not be in the null space of $\boldsymbol{B}_\circ$.

Having shown that $D_P(\boldsymbol{f})$ is strictly concave, the $D_P(\boldsymbol{f})$ maximization problem has a unique solution which solves $\nabla_{\boldsymbol{f}} D_P(\boldsymbol{f}) = \boldsymbol{0}$ leading to $\boldsymbol{B}_\circ^\top \boldsymbol{B}_\circ \boldsymbol{f} = \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ$. Let

$$g_{ij}^{(d)}(n) := \big(\boldsymbol{B}_\circ^\top \boldsymbol{B}_\circ \boldsymbol{f}(n)\big)_{ij}^{(d)}$$

which represents the entry of hyper-vector $\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f}$ corresponding to the edge ij and class d at the slot n. Then the maximizing $\boldsymbol{f}(n)$ is the one which respects $g_{ij}^{(d)}(n) = q_{ij}^{(d)}(n)$ for any link ij and any class d.

Next assume that a routing policy takes $\alpha f_{ij}^{(d)} = q_{ij}^{(d)}$ for a scalar $\alpha > 0$ which will be determined later. By the argument above, such a policy can maximize $D_P(\boldsymbol{f})$ every timeslot, if it complies with $g_{ij}^{(d)} = \alpha f_{ij}^{(d)}$. Plugging this in (84), the maximum of $D_P(\boldsymbol{f})$ is found as

$$D_P^{\max} = \sum\nolimits_{ij\in\mathcal{E}} \sum\nolimits_{d\in\mathcal{K}} f_{ij}^{(d)}\big(2\,q_{ij}^{(d)} - \alpha f_{ij}^{(d)}\big). \tag{85}$$

Substituting $\alpha f_{ij}^{(d)} = q_{ij}^{(d)}$ yields

$$D_P^{\max} = \frac{1}{\alpha}\sum\nolimits_{ij\in\mathcal{E}} \sum\nolimits_{d\in\mathcal{K}} \big(q_{ij}^{(d)}\big)^2 \tag{86}$$

which represents the maximum of $D_P(\boldsymbol{f})$ for a specific value of $\alpha > 0$ and in the absence of all network constraints.

Now let us add the network constraints into the picture. Obviously, the $D_P^{\max}$ as (86) is no longer attainable. However, going one step back to the $D_P^{\max}$ in (85), under $\alpha f_{ij}^{(d)} = q_{ij}^{(d)}$ forwarding regime, one can find the maximizing $\boldsymbol{f}$ by solving the following optimization problem at every slot n:

$$\begin{aligned} &\text{Maximize:} && \sum_{ij\in\mathcal{E}}\sum_{d\in\mathcal{K}} 2\,q_{ij}^{(d)}(n) f_{ij}^{(d)}(n) - f_{ij}^{(d)}(n)^2 \\ &\text{Subject to:} && \text{Network constraints.}\end{aligned} \tag{87}$$

Let us first determine α^* which maximizes (87). Obviously, the maximizing α is the smallest one. However, the following expression shows that α can not go lower than one:

$$\boldsymbol{f}^\top\boldsymbol{f} \overset{(a)}{\leqslant} \big(\boldsymbol{B}_\circ^+\boldsymbol{f}\big)^\top\big(\boldsymbol{B}_\circ^+\boldsymbol{f}\big) \overset{(b)}{\leqslant} \boldsymbol{q}_\circ^\top\boldsymbol{q}_\circ \overset{(c)}{=} \boldsymbol{f}^\top\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f} \overset{(d)}{=} \alpha\boldsymbol{f}^\top\boldsymbol{f}$$

which leads to $\alpha^* = 1$. The inequality (a) reads the fact that

$$\sum_{i\in\mathcal{V}}\sum_{b\in\text{out}(i)}\sum_{d\in\mathcal{K}} \big(f_{ib}^{(d)}\big)^2 \leqslant \Big(\sum_{i\in\mathcal{V}}\sum_{b\in\text{out}(i)}\sum_{d\in\mathcal{K}} f_{ib}^{(d)}\Big)^2.$$

Note that $\boldsymbol{B}_\circ^+$ represents the incidence matrix in (18) where all -1 entries are replaced by 0. Thus $\boldsymbol{B}_\circ^+\boldsymbol{f}$ is a hyper-vector in which the entry corresponding to node i and class d shows the outgoing flow of class d from the node i. The inequality (b) says that in each timeslot the number of d-classes leaving a node i is at most equal to the queue backlog of class d at the node i. The equality (c) comes from the D_P maximizing condition $\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f} = \boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ$. To see this, multiply both sides by $\boldsymbol{B}_\circ$ to get $(\boldsymbol{B}_\circ\boldsymbol{B}_\circ^\top)\,\boldsymbol{B}_\circ\boldsymbol{f} = (\boldsymbol{B}_\circ\boldsymbol{B}_\circ^\top)\,\boldsymbol{q}_\circ$ and then the positive definiteness of $\boldsymbol{B}_\circ\boldsymbol{B}_\circ^\top$ gives the result. Finally, the equality (d) reads $\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f} = \alpha\boldsymbol{f}$, which follows from the D_P maximizing condition $\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f} = \boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ$ together with the forwarding regime $\alpha f_{ij}^{(d)} = q_{ij}^{(d)}$, or equivalently $\alpha\boldsymbol{f} = \boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ$.

Fixing $\alpha = 1$, the optimization problem (87) is identical to the one in (14) for $\phi_{ij}(n) = 1$, or equivalently $\beta = 0$. We have shown in Th. 1 that HD with $\beta \in [0, 1]$ solves (14) at each slot n. Taking $\beta = 0$, it implies that HD with $\beta = 0$ solves (87) at each slot n, which concludes the proof. ∎

E. Proof of Lemma 4

The time variable (n) is often dropped for brevity. Consider the basic quadratic Lyapunov function $W(n) := \boldsymbol{q}_\circ(n)^\top\boldsymbol{q}_\circ(n)$ and find the Lyapunov drift $\Delta W(n) = W(n+1) - W(n)$ by substituting for $\boldsymbol{q}_\circ(n+1)$ from (19). This leads to

$$\Delta W = 2\boldsymbol{a}_\circ^\top\boldsymbol{q}_\circ - 2\boldsymbol{a}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f} - 2\boldsymbol{f}^\top\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ + \boldsymbol{a}_\circ^\top\boldsymbol{a}_\circ + \boldsymbol{f}^\top\boldsymbol{B}_\circ^\top\boldsymbol{B}_\circ\boldsymbol{f}.$$

Taking expectation from ΔW above yields

$$\begin{aligned}\mathbb{E}\{\Delta W\} = \;& 2\,\mathbb{E}\{\boldsymbol{a}_\circ\}^\top\mathbb{E}\{\boldsymbol{q}_\circ\} - 2\,\mathbb{E}\{\boldsymbol{a}_\circ\}^\top\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\} \\ & - 2\,\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\}^\top\mathbb{E}\{\boldsymbol{q}_\circ\} - 2\,\mathbb{C}\text{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\} \\ & + \mathbb{E}\{\boldsymbol{a}_\circ\}^\top\mathbb{E}\{\boldsymbol{a}_\circ\} + \mathbb{V}\text{ar}\{\boldsymbol{a}_\circ\} \\ & + \mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\}^\top\mathbb{E}\{\boldsymbol{B}_\circ\boldsymbol{f}\} + \mathbb{V}\text{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\}.\end{aligned}$$

where the expectation is with respect to the randomness of arrivals, channel states, and routing decision—in case of a randomized routing algorithm. The assumption in the above equality is that the network layer controller has no impact on arrivals. This means that $\boldsymbol{a}_\circ(n)$ is an independent variable in the sense that it is not influenced by the controller. Thus $\boldsymbol{a}_\circ(n)$ is statistically uncorrelated with the random variables $\boldsymbol{q}_\circ(n)$ and $\boldsymbol{f}(n)$, meaning that $\mathbb{C}\text{ov}\{\boldsymbol{a}_\circ, \boldsymbol{q}_\circ\} = 0$ and $\mathbb{C}\text{ov}\{\boldsymbol{a}_\circ, \boldsymbol{B}_\circ\boldsymbol{f}\} = 0$.

By some elementary matrix algebra, the drift expectation $\mathbb{E}\{\Delta W\}$ above can be recast as

$$\begin{aligned}\mathbb{E}\{\Delta W\} = \;& \mathbb{E}\{\boldsymbol{a}_\circ - \boldsymbol{B}_\circ\boldsymbol{f}\}^\top\mathbb{E}\{\boldsymbol{a}_\circ - \boldsymbol{B}_\circ\boldsymbol{f} + 2\,\boldsymbol{q}_\circ\} \\ & - 2\,\mathbb{C}\text{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\} + \mathbb{V}\text{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\} + \mathbb{V}\text{ar}\{\boldsymbol{a}_\circ\}\end{aligned}$$

where the equality holds for every timeslot. Summing over timeslots 0 until $\tau - 1$, dividing the sum by τ, and taking a $\limsup$ of $\tau \to \infty$ result in

$$\begin{aligned}&\limsup_{\tau\to\infty} 1/\tau \sum\nolimits_{n=0}^{\tau-1} \mathbb{E}\{\boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ\boldsymbol{f}(n)\}^\top\mathbb{E}\{\boldsymbol{h}(n)\} = \\ &\qquad + 2\,\overline{\mathbb{C}\text{ov}\{\boldsymbol{B}_\circ^\top\boldsymbol{q}_\circ, \boldsymbol{f}\}} - \overline{\mathbb{V}\text{ar}\{\boldsymbol{B}_\circ\boldsymbol{f}\}} - \overline{\mathbb{V}\text{ar}\{\boldsymbol{a}_\circ\}}\end{aligned}$$

where $\boldsymbol{h}(n) := \boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ\boldsymbol{f}(n) + 2\,\boldsymbol{q}_\circ(n)$, and where we have used $\limsup_{\tau\to\infty}(W(\tau) - W(0))/\tau = 0$ as the routing policy stabilizes $\overline{\boldsymbol{a}_\circ}$ and so keeps $W(n) = \boldsymbol{q}(n)^\top\boldsymbol{q}(n)$ finite at each timeslot. Next we show that the left hand side of the equality above is zero which yields the result.

Observe that $\boldsymbol{h}(n)$ is entrywise nonnegative because $\boldsymbol{a}_\circ(n)$ is entrywise nonnegative and $\boldsymbol{B}_\circ\boldsymbol{f}(n) \preccurlyeq \boldsymbol{q}_\circ(n)$, saying that in each timeslot the number of packets leaving a node is at most equal to the node queue backlog, where curly-inequality denotes an entrywise comparison. Also notice that $\boldsymbol{h}(n)$ is finite because (i) $\boldsymbol{a}_\circ(n)$ has finite mean and variance, (ii) $\boldsymbol{f}(n)$ is finite since forwarding on each link is bounded with the link capacity which is finite, and (iii) $\mathbb{E}\{\boldsymbol{q}_\circ(n)\}$ is finite due to the queue stability. Thus there exist constant vector bounds $\boldsymbol{h}_{\min}$ and $\boldsymbol{h}_{\max}$ such that $\boldsymbol{0} \preccurlyeq \boldsymbol{g}_{\min} \preccurlyeq \mathbb{E}\{\boldsymbol{g}(n)\} \preccurlyeq \boldsymbol{g}_{\max} \prec \boldsymbol{\infty}$ for each slot n. This implies that

$$\begin{aligned}&(\,\overline{\boldsymbol{a}_\circ} - \overline{\boldsymbol{B}_\circ\boldsymbol{f}}\,)^\top\boldsymbol{g}_{\min} \leqslant \\ &\qquad \lim_{\tau\to\infty} 1/\tau \sum\nolimits_{n=0}^{\tau-1} \mathbb{E}\{\boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ\boldsymbol{f}(n)\}^\top\mathbb{E}\{\boldsymbol{g}(n)\} \leqslant \\ &\qquad\qquad (\,\overline{\boldsymbol{a}_\circ} - \overline{\boldsymbol{B}_\circ\boldsymbol{f}}\,)^\top\boldsymbol{g}_{\max}\end{aligned}$$

Since $\overline{\boldsymbol{a}_\circ}$ is stabilized by the routing policy, the feasibility condition (21) compels $\overline{\boldsymbol{a}_\circ} = \overline{\boldsymbol{B}_\circ\boldsymbol{f}}$ leading to

$$\lim_{\tau\to\infty} 1/\tau \sum\nolimits_{n=0}^{\tau-1} \mathbb{E}\{\boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ\boldsymbol{f}(n)\}^\top\mathbb{E}\{\boldsymbol{g}(n)\} = 0$$

which concludes the proof. ∎

REFERENCES

[1] L. Tassiulas, A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *Tr. Autom. Cont.,* 37:1936–1949, 1992.

[2] L. Georgiadis, M. Neely, R. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Found. Tren. Net.,* 1:1–144, 2006.

[3] M. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems.* Morgan & Claypool, 2010.

[4] J. Dai, W. Lin. Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Ann. Appl. Prob.,* 18:2239–2299, 2008.

[5] D. Shah, D. Wischik. Optimal scheduling algorithms for input-queued switches. *INFOCOM,* 2006.

[6] K. Ross, N. Bambos. Projective cone scheduling (PCS) algorithms for packet switches of maximal throughput. *Tr. Net.,* 17:976–989, 2009.

[7] M. Naghshvar, H. Zhuang, T. Javidi. A general class of throughput-optimal routing policies in multi-hop wireless networks. *Tr. Info. Theo.,* 58:2175–2193, 2012.

[8] T. Ji, E. Athanasopoulou, R. Srikant. Optimal scheduling policies in small generalized switches. *INFOCOM,* 2009.

[9] L. Bui, R. Srikant, A. Stolyar. A novel architecture for delay reduction in back-pressure scheduling algorithm. *Tr. Net.,* 19:1597–1609, 2011.

[10] B. Ji, C. Joo, N. Shroff. Delay-based back-pressure scheduling in multi-hop wireless networks. *INFOCOM,* 2011.

[11] L. Huang, S. Moeller, M. Neely, B. Krishnamachari. LIFO-backpressure achieves near optimal utility-delay tradeoff. *WiOpt,* 2011.

[12] M. Alresaini, M. Sathiamoorthy, B. Krishnamachari, M. Neely. Back-pressure with adaptive redundancy (BWAR). *INFOCOM,* 2012.

[13] D. Xue, R. Murawski, E. Ekici. Distributed utility-optimal scheduling with finite buffers. *WiOpt,* 2012.

[14] L. Bui, S. Sanghavi, R. Srikant. Distributed link scheduling with constant overhead. *Tr. Net.,* 17:1467–1480, 2009.

[15] L. Jiang, J. Walrand. Approaching throughput-optimality in distributed CSMA scheduling algorithms with collisions. *Tr. Net.,* 19:816–829, 2011.

[16] S. Moeller, A. Sridharan, B. Krishnamachari, O. Gnawali. Routing without routes: the backpressure collection protocol. *IPSN,* 2010.

[17] J. Martnez, J. Bafalluy. Design, implementation, tracing of dynamic backpressure routing for NS-3. *SIMUTools,* 2011.

[18] E. Stai, J. Baras, S. Papavassiliou. Throughput-delay tradeoff in wireless multi-hop networks via greedy hyperbolic embedding. *MTNS,* 2012.

[19] G. Sharma, R. Mazumdar, N. Shroff. On the complexity of scheduling in wireless networks. *MobiCom,* 2006.

[20] J. Dai. Stability of fluid and stochastic processing networks. *MaPhySto Misc. Pub.,* no. 9, 1999.

[21] M. Bramson. Stability of queueing networks. *Prob. Surv.,* 169-345, 2008.

[22] X. Lin, N. Shroff. Joint rate control and scheduling in multihop wireless networks. *CDC,* 2004.

[23] M. Neely, E. Modiano, C. Li. Fairness and optimal stochastic control for heterogeneous networks. *INFOCOM,* 2005.

[24] A. Eryilmaz, R. Srikant. Fair resource allocation wireless networks using queue-length-based scheduling and congestion control. *INFOCOM,* 2005.

[25] A. Stolyar. Maximizing queueing network utility subject to stability: greedy primal-dual algorithm. *Que. Sys.,* 50:401-457, 2005.

[26] R. Berry, R. Gallager. Communication over fading channels with delay constraints. *Tr. Info. Theo.,* 48:1135–1149, 2002.

[27] D. Arnold, R. Falk, R. Winther. Finite element exterior calculus: from Hodge theory to numerical stability. *Amer. Math. Soc.,* 47:281–354, 2010.

[28] S. Ohta, K. Sturm. Heat flow on Finsler manifolds. *Comm. Pure Appl. Math.,* 62:13861433, 2009.

[29] D. Boley, G. Ranjan, Z. Zhang. Commute times for a directed graph using an asymmetric Laplacian. *Lin. Algeb. Applic.,* 435:224–242, 2011.

[30] P. Doyle, J. Snell. Random walks and electrical networks. *Math. Assoc. America,* 1984.

[31] I. Kim, O. Weck. Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Struct. Mult. Opt.* 29:149–158, 2005.

[32] W. Chen. *Applied Graph Theory: Graphs and Electrical Networks.* North Holland Publishing, 1976.

[33] E. Leonardi, M. Melia, F. Neri, M. Marson. Bounds on average delays and queue size averages and variances in input-queued cell-based switches. *INFOCOM,* 2001.