Heat-Diffusion: Pareto Optimal Dynamic Routing for Time-Varying Wireless Networks

Reza Banirazi, Edmond Jonckheere, Bhaskar Krishnamachari
Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089
E-mail: {banirazi, jonckhee, bkrishna}@usc.edu

***Abstract*—A new routing policy, named Heat-Diffusion (HD), is developed for multihop wireless networks subject to stochastic arrivals, time-varying topology, and inter-channel interference, using only current queue congestion and current channel states, without requiring the knowledge of topology and arrivals. Besides throughput optimality, HD minimizes a quadratic routing cost defined by endowing each channel with a cost-factor. It also minimizes average total queue congestion, and so average network delay, within the class of routing policies that base decision only on current queue lengths and current channel states. Further, within this class, HD provides a Pareto optimal tradeoff between average delay and average routing cost, meaning that no policy can improve either one without detriment to the other. Finally, HD fluid limit follows graph combinatorial heat equation, which opens a new way to study wireless networks using heat calculus, a very active area of pure mathematics.**

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 that 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, possibly time-varying, 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 that use only current queue occupancies and current channel states, possibly together with the knowledge of arrival statistics and 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-mentioned 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 these two objectives.

Similar to BP, also HD requires a centralized scheduling whose complexity is prohibitive for practice. However, much progress has recently been made to ease this difficulty by deriving decentralized schedulers with the performance of arbitrarily close to the centralized version [8], [9].

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, max-weight 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-proven conjecture that heavy traffic delay is minimized when $\alpha \to 0$, where a discussion of this was given in [10] 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 "patchwork" of localized piecewise quadratic Lyapunov functions. In wireless networks, shadow queues enabled BP to handle multicast sessions with reduced number of actual queues that need to be maintained [11]. Replacing queue-length by packet-age, [12] introduced a delay-based BP policy. To improve BP delay performance, [13] proposed place-holders with Last-In-First-Out (LIFO) forwarding. Adaptive redundancy was used in [14] 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 made to adopt the original framework for handling finite buffers [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].

Contributions—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 takes the form of heat flow on a suitably-weighted directed graph, 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 paves a way to take advantage of powerful tools from heat calculus in the analysis and optimization of stochastic, packet-based, time-slotted queuing networks constrained by link interference.

Second, the new policy minimizes average network delay in the class of all algorithms that 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].

Third, the new policy reduces the Dirichlet routing cost to its minimum feasible value. 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.

Fourth, in the above-mentioned class of routing algorithms, through changing a Lagrange control parameter, HD provides a Pareto optimal performance with respect to average delay and the 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 criteria.

Last but not least, HD enjoys the same algorithmic structure, complexity, and overhead as BP, giving them the same wide-reaching impact. This also provides an easy way to leverage all advanced improvements to BP to further enhance HD quality. At the same time, it simplifies the way to practice via a smooth software transition from BP to HD.

Note: The page limit does not let us include the proofs, but all the proofs are available in [27]. A very infant idea of HD algorithm was introduced in [28].

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 randomly arrive into different nodes, requiring a multihop routing. 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.

For a link ij, the *capacity* $\mu_{ij}(n)$, which is frequently called *transmission rate* in literature, counts the maximum number of packets the link can transmit at the slot n. The *actual-transmission* $f_{ij}(n)$, on the other hand, counts the number of 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

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

where $\mathbb{E}$ denotes expectation. Throughout this paper, an over-bar notation will denote the $\limsup$ expected time average as defined in (1). A queuing network is *stable* if all its queues are stable. 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 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 [21]–[24] that 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:*

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

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 exponentially 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:*

$$\begin{aligned} \text{Minimize:} \quad & \overline{Q} := \sum\nolimits_{i\in\mathcal{V}} \overline{q_i} \\ \text{Subject to:} \quad & \text{Network constraints.} \end{aligned} \tag{3}$$

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 [25]. 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 any of the above-mentioned difficulties.

- *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 trade average delay for average routing cost, which also plays the role of Lagrange multiplier. 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

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

1) BP weighing: For every link ij find queue differential $q_{ij}(n) := q_i(n) - q_j(n)$ and give a weight to the link as

$$w_{ij}(n) := \mu_{ij}(n)\, q_{ij}(n)^+ \tag{5}$$

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

2) BP scheduling: Find the scheduling vector such that

$$\boldsymbol{\pi}(n) = \arg\max_{\boldsymbol{\pi}\in\Pi} \sum\nolimits_{ij\in\mathcal{E}} \pi_{ij} w_{ij}(n) \tag{6}$$

where ties are broken randomly.

3) BP forwarding: Over each activated link with $w_{ij}(n) > 0$ transmit packets at full capacity $\mu_{ij}(n)$. If there is no enough packets 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 (5) with

$$w_{ij}(n) := \mu_{ij}(n)\big(\, q_{ij}(n) - V\rho_{ij}(n)\, \mu_{ij}(n)\,\big)^+ \tag{7}$$

where $V \in [0, \infty)$ determines the importance of 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]. Therefore, 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

At every timeslot n, HD control policy for the network layer observes link queue-differentials $q_{ij}(n) := q_i(n) - q_j(n)$ and estimates channel capacities $\mu_{ij}(n)$ and channel cost-factors $\rho_{ij}(n)$ to make a routing decision as follows.

1) HD weighing: Every link first calculates the number of packets it would transmit if it were activated:

$$\widehat{f_{ij}}(n) := \min\{\phi_{ij}(n) q_{ij}(n)^+,\ \mu_{ij}(n)\}$$
$$\phi_{ij}(n) := (1-\beta) + \beta/\rho_{ij}(n) \tag{8}$$

TABLE I
Contrasting HD policy with V-parameter BP policy.

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	$\mu_{ij}(n)\big(q_{ij}(n) - V\rho_{ij}(n)\,\mu_{ij}(n)\big)^+$
		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			$f_{ij}(n) = \begin{cases} \lceil\widehat{f_{ij}}(n)\rceil & \text{if } \pi_{ij}(n)=1 \\ 0 & \text{otherwise}\end{cases}$

where the Lagrange control parameter β is as defined in (4) to make a tradeoff between queue occupancy and routing penalty, and the hat notation denotes a predicted value which would not necessarily be realized. Then the link weight

$$w_{ij}(n) := 2\,\phi_{ij}(n) q_{ij}(n)\widehat{f_{ij}}(n) - \widehat{f_{ij}}(n)^2. \tag{9}$$

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

$$\boldsymbol{\pi}(n) = \arg\max_{\boldsymbol{\pi}\in\Pi} \sum\nolimits_{ij\in\mathcal{E}} \pi_{ij} w_{ij}(n) \tag{10}$$

where ties are broken randomly.

3) HD forwarding: Over each activated link transmit $\widehat{f_{ij}}(n)$ number of packets, meaning that

$$f_{ij}(n) = \begin{cases} \lceil \widehat{f_{ij}}(n) \rceil & \text{if } \pi_{ij}(n) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

where $f_{ij}(n)$ represents the number of packets genuinely sent over the link ij at the slot n, and the ceiling function $\lceil x \rceil$ maps a real number x to the smallest following integer.

It is critical to discriminate among actual transmissions $f_{ij}(n)$, transmission predictions $\widehat{f_{ij}}(n)$, and link capacities $\mu_{ij}(n)$. Table 1 compares HD and V-parameter BP algorithms, emphasizing the same structure and complexity.

Remark 1: 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.

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 indirectly it also takes into account the link capacity through (8). Thus HD allocates resources based only on genuinely transmittable packets, without counting on null packets as is practiced in BP schemes.

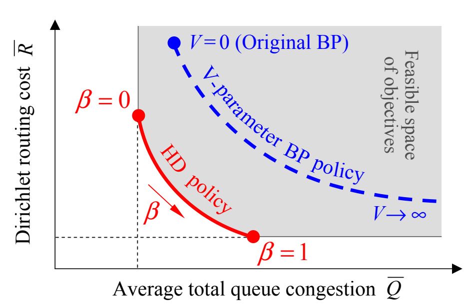

Fig. 1. Graphical description of HD Pareto optimality with respect to average queue congestion and Dirichlet routing cost, contrasted with BP performance.

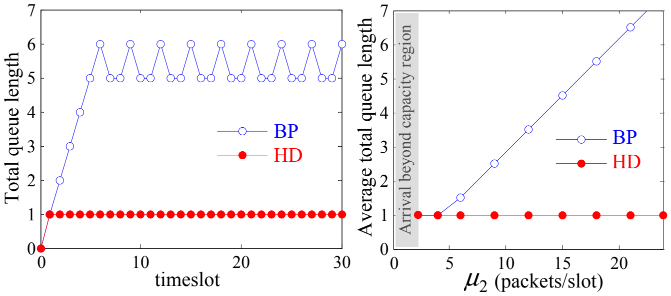

Fig. 2. Performance of HD versus BP in the two-queue downlink.

H2: The link weight (9), 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 $w_{ij}(n) = \phi_{ij}(n)^2 q_{ij}(n)^2$. 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) which is fundamental to other HD qualities.

H3: Varying the Lagrange multiplier 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 that rely only on current queue backlogs and current channel states (Th. 3).
- Raising β adds to 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 criteria (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 that act based only on current queue congestion and current channel states.

H4: Unlike BP that forwards the highest possible 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 routing penalty. Forwarding a portion of queue differentials rather than filling up the link capacities also complies with resembling *heat flow* on the underlying directed graph (Th. 5) that in effect minimizes time average routing cost via Dirichlet Principle (Th. 8).

C. 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.

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 $\rho_{ij} = \mathrm{ETX}_{ij}$. For zero initial conditions, Fig. 3 compares the performance of HD with BP. While HD easily stabilizes the total queued packets at 1 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 (10) rather than an *addend* in the V-parameter BP formula (7).

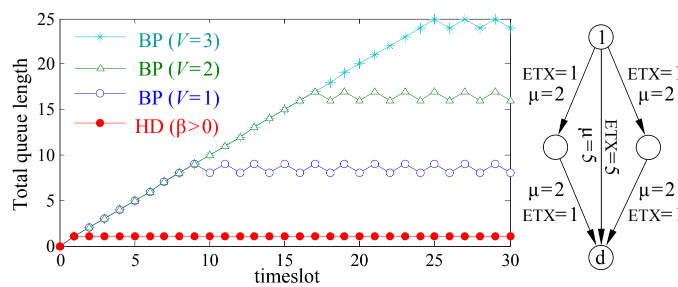

Fig. 3. Performance of HD versus BP in the lossy link network.

Power minimization: Consider the sensor network of Fig. 4 subject to 1-hop interference model, i.e. two links with a common node cannot transmit at the same time. Each link ij has a noise intensity $N_{ij} \in [1, 5]$ which is randomly assigned at first and keeps constant during the simulation. We adopt Shannon capacity $\mu_{ij} = \Omega_{ij} \log_2(1 + P_{ij}/N_{ij})$ with P_{ij} the power transmission and Ω_{ij} the bandwidth. Each timeslot two packets arrive at nodes 1, 2, 3, and 4 destined for the node d.

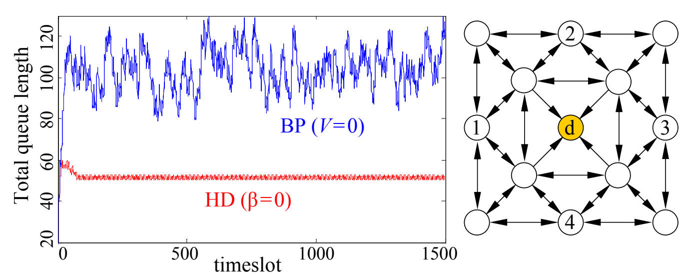

Fig. 4. Total queue backlog of HD versus BP in the power minimization.

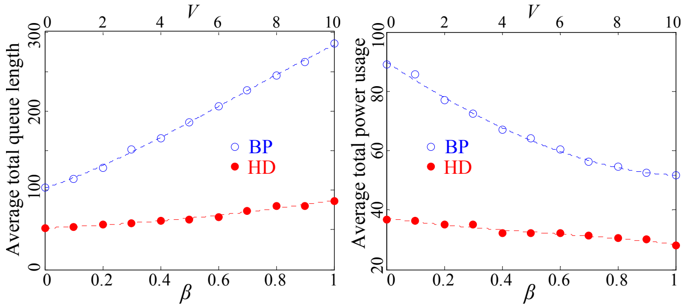

Fig. 5. Trading queue congestion for power consumption in HD as a function of β, and in BP as a function of V. Dashed lines represent interpolation.

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 $\Omega_{ij} = 5$ 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, i.e. $V = 0$. 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 V-parameter BP with $V = 10$. Smaller oscillations in HD endorses both H1 and H4, showing the defect of capacity-based scheduling and maximum packet forwarding in BP.

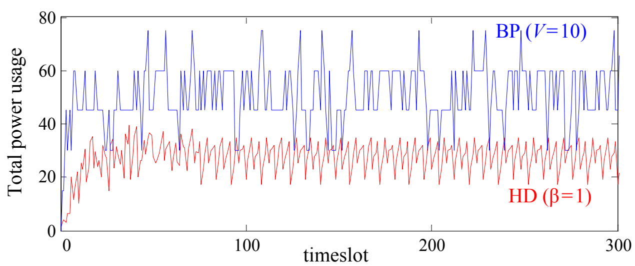

Fig. 6. Timeslot evolution of total power consumption in HD versus BP.

IV. Heat-Diffusion Throughput Optimality

Consider a general queuing network with $q_i(n)$ being the integer number of packets in the node i at the slot n as before. The state variables of the system are represented by the vector

$$\boldsymbol{q}_\circ(n) := \big[\, q_1(n), \dots, q_{d-1}(n), q_{d+1}(n), \dots, q_{|\mathcal{V}|}(n) \,\big]$$

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

Notation 1: 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 $a_i(n)$ represent the integer number of exogenous packets arriving into the node i at the slot n. Discarding $a_d(n) \equiv 0$, the vector of node arrivals

$$\boldsymbol{a}_\circ(n) := \big[\, a_1(n), \dots, a_{d-1}(n), a_{d+1}(n), \dots, a_{|\mathcal{V}|}(n) \,\big].$$

Likewise, the vector of link actual-transmissions

$$\boldsymbol{f}(n) := \big[\, f_1(n), \dots, f_{|\mathcal{E}|}(n) \,\big]$$

where as before, $f_{ij}(n)$ is the integer number of packets actually sent over the link ij at the slot 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. Then $\boldsymbol{B}_\circ$ denotes a reduction of $\boldsymbol{B}$ through discarding the row related to the destination node d. One can verify that $\boldsymbol{B}_\circ \boldsymbol{f}(n)$ is a node vector with the entry corresponding to node i as

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

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 queuing network is captured by

$$\boldsymbol{q}_\circ(n+1) = \boldsymbol{q}_\circ(n) + \boldsymbol{a}_\circ(n) - \boldsymbol{B}_\circ \boldsymbol{f}(n). \tag{12}$$

A. The Key Property of Heat-Diffusion Policy

Having (12), 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 maximizes the $\boldsymbol{f}$-controlled functional

$$D(\boldsymbol{f}, \beta, n) := 2\, \boldsymbol{f}(n)^\top \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) - \boldsymbol{f}(n)^\top \boldsymbol{f}(n) \tag{13}$$

subject to network constraints, where $\boldsymbol{\Phi}(n) = \mathrm{diag}(\boldsymbol{\phi}(n))$ denotes the diagonal matrix expansion of vector $\boldsymbol{\phi}(n)$ with $\phi_{ij}(n) = (1-\beta) + \beta/\rho_{ij}(n)$ as defined in (8).

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 the uncontrollable factors that affect link capacities and cost-factors. We assume that $\boldsymbol{S}(n)$ evolves according to an ergodic stationary process and takes values in a finite set $\mathcal{S}$. Considering a connectivity graph $(\mathcal{V}, \mathcal{E})$ together with a channel state process $\boldsymbol{S}(n)$, an arrival *rate* vector $\overline{\boldsymbol{a}}$ is in the network capacity region $\mathcal{C}$ iff there exists a set of link actual-transmissions such that

$$\overline{a_i} = \sum\nolimits_{b \in \mathrm{out}(i)} \overline{f_{ib}} - \sum\nolimits_{a \in \mathrm{in}(i)} \overline{f_{ai}} \tag{14}$$

$$\overline{f_{ij}} \leqslant \sum\nolimits_{\boldsymbol{S} \in \mathcal{S}} s\, \mathbb{E}\{\mu_{ij}(n) \big| \boldsymbol{S}(n) = \boldsymbol{S}\}. \tag{15}$$

C. HD Throughput-Optimality for All β

To analyze the HD stability, we use a Lyapunov argument which, unlike most of previous results in literature that use the sum of squares of queues, involves a Lyapunov candidate with nonsymmetric weighting matrix and no trivial interpretation of a specific energy in the system. Let the functional

$$W(n) := \boldsymbol{q}_\circ(n)^\top \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n)$$

where $\boldsymbol{M}_\circ(n) := \big(\boldsymbol{B}_\circ \boldsymbol{B}_\circ^\top\big)^{-1} \boldsymbol{B}_\circ \boldsymbol{\Phi}(n) \boldsymbol{B}_\circ^\top$ is a nonsymmetric matrix. Note that for $\beta = 0$ where $\boldsymbol{\Phi}(n) = \boldsymbol{I}$ with $\boldsymbol{I}$ the identity matrix, or for the case that all links are of the same cost-factor where $\boldsymbol{\Phi}(n) = \alpha \boldsymbol{I}$ with $\alpha > 0$ a scalar, the Lyapunov function $W(n)$ reduces to the sum of squares of queue lengths.

Lemma 1: On a connected-topology network, $\boldsymbol{M}_\circ(n)$ 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,

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

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 $\boldsymbol{x}^\top \boldsymbol{M}_\circ(n)\,\boldsymbol{y} > 0$.

Letting the Lyapunov drift $\Delta W(n) := W(n+1) - W(n)$, exploiting Lem. 2, replacing $\boldsymbol{f}^\top \boldsymbol{B}_\circ^\top \boldsymbol{M}_\circ \boldsymbol{q}_\circ$ with $\boldsymbol{f}^\top \boldsymbol{\Phi} \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ$ in light of Lem. 1, and using the $D(\boldsymbol{f}, \beta, n)$ expression of (13),

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

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.

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 (14), there exists a packet flow $\boldsymbol{f'}(n)$ such that $\overline{\boldsymbol{B}_\circ \boldsymbol{f'}} = \overline{\boldsymbol{a}_\circ} + \boldsymbol{\epsilon}$. At the same time, Th. 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 $\Delta W(n)$ is bounded for all $\beta \in [0, 1]$.

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.

V. Congestion Minimizing Policy at $\beta = 0$

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 section we establish the first pillar. Prior to stating the main result in Th. 3, we propose two lemmas that are used in the proof.

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

$$2\,\boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(n) - \boldsymbol{f}(n)^\top \boldsymbol{B}_\circ^\top \boldsymbol{B}_\circ \boldsymbol{f}(n) \tag{16}$$

at every timeslot n subject to network constraints.

Considering Lem. 3 and Th. 1 together, it turns out that both the functional (13) at $\beta = 0$, where $\boldsymbol{\Phi}(n)$ becomes the identity matrix, and the functional (16) 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

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

where for two random variables $\boldsymbol{X}$ and $\boldsymbol{Y}$, $\mathbb{C}\mathrm{ov}\{\boldsymbol{X}, \boldsymbol{Y}\} := \mathbb{E}\{\boldsymbol{X}^\top \boldsymbol{Y}\} - \mathbb{E}\{\boldsymbol{X}\}^\top \mathbb{E}\{\boldsymbol{Y}\}$ and $\mathbb{V}\mathrm{ar}\{\boldsymbol{X}\} := \mathbb{C}\mathrm{ov}\{\boldsymbol{X}, \boldsymbol{X}\}$.

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 that do not incorporate the Markov structure of topology process into their decisions, among them is BP [1] and most its derivations [4]–[17]. It also includes all stationary randomized algorithms that 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 that 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 (3).

VI. Classical versus Combinatorial Heat Process

To formulate the graph heat diffusion, we use the theory of *combinatorial geometry* where the notion of *chains-cochains* provides a genuine counterpart for *differential forms* in geometry. The details are found in [26] and references therein.

A. Continuous Heat Diffusion on Manifolds

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

$$\frac{\partial Q(\boldsymbol{z}, t)}{\partial t} = -\mathrm{div} \boldsymbol{F}(\boldsymbol{z}, t) + A(\boldsymbol{z}, t). \tag{17}$$

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

$$\boldsymbol{F}(\boldsymbol{z}, t) = -\sigma(\boldsymbol{z})\, \nabla Q(\boldsymbol{z}, t) \tag{18}$$

where $\sigma(\boldsymbol{z})$ is *thermal diffusivity* quantifying how fast heat moves through the material. Putting (17) and (18) together,

$$\frac{\partial Q(\boldsymbol{z}, t)}{\partial t} = \mathrm{div}\big(\sigma(\boldsymbol{z})\, \nabla Q(\boldsymbol{z}, t)\big) + A(\boldsymbol{z}, t). \tag{19}$$

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 representing the discrete domain, pointwise functions $Q(\boldsymbol{z}, t)$ and $A(\boldsymbol{z}, t)$ are respectively replaced by 0-*cochain* vectors $\boldsymbol{q}(t)$ and $\boldsymbol{a}(t)$ (node variables), line integral $\boldsymbol{F}(\boldsymbol{z}, t)$ is replaced by a 1-*cochain* vector $\boldsymbol{f}(t)$ (edge variable), and thermal diffusivity σ is replaced by an edge weight vector $\boldsymbol{\sigma}$.

As a 1-complex, the structure of a graph is fully described by the *node-edge* incidence matrix $\boldsymbol{B}$ of Sec. IV, except that we substitute edge direction for algebraic topological edge orientation. Then on an *undirected* graph with node d as the heat sink, the combinatorial analogue of the classic heat equations (17)–(19) are given by

$$\dot{\boldsymbol{q}}(t) = -\boldsymbol{B}\, \boldsymbol{f}(t) + \boldsymbol{a}(t)\,,\;\; q_d(t) = 0 \tag{20}$$

$$\boldsymbol{f}(t) = \mathrm{diag}(\boldsymbol{\sigma})\, \boldsymbol{B}^\top \boldsymbol{q}(t) \tag{21}$$

$$\dot{\boldsymbol{q}}(t) = -\boldsymbol{B}\, \mathrm{diag}(\boldsymbol{\sigma})\, \boldsymbol{B}^\top \boldsymbol{q}(t) + \boldsymbol{a}(t)\,,\;\; q_d(t) = 0 \tag{22}$$

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, which absorbs all the generated heat.

Enforcing the boundary condition $q_d(t) = 0$, one can eliminate the sink d from (20)–(22). This yields the reduced set of *continuous-time* graph heat equations

$$\boldsymbol{f}(t) = \operatorname{diag}(\boldsymbol{\sigma})\, \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t) \tag{23}$$

$$\dot{\boldsymbol{q}}_\circ(t) = -\boldsymbol{L}_\circ\, \boldsymbol{q}_\circ(t) + \boldsymbol{a}_\circ(t)\ ,\ \boldsymbol{L}_\circ := \boldsymbol{B}_\circ \operatorname{diag}(\boldsymbol{\sigma})\, \boldsymbol{B}_\circ^\top \tag{24}$$

where the $\circ$ subscript denotes a reduced quantity as before. The linear operator $\boldsymbol{L}_\circ$ 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 (20) remains unchanged, but the Fick's law (21) 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

$$\boldsymbol{f}(t) = \operatorname{diag}(\boldsymbol{\sigma}) \max\big\{\boldsymbol{0},\ \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t)\big\} \tag{25}$$

$$\begin{aligned} \dot{\boldsymbol{q}}_\circ(t) &= -\vec{\boldsymbol{L}}_\circ\, \boldsymbol{q}_\circ(t) + \boldsymbol{a}_\circ(t) \\ \vec{\boldsymbol{L}}_\circ &:= \boldsymbol{B}_\circ \operatorname{diag}(\boldsymbol{\sigma})\, \boldsymbol{B}_\circ^\top \operatorname{diag}\big(\mathbb{I}_{\boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t) \succ \boldsymbol{0}}\big) \end{aligned} \tag{26}$$

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. We call $\vec{\boldsymbol{L}}_\circ$ the *nonlinear Dirichlet Laplacian* acting on a directed graph.

VII. Wireless Network Thermodynamics

Though defined on a directed graph, (25)–(26) still represent a deterministic, continuous-time process with no link interference. The latter, particularly, makes the wireless problem quite intractable. Nevertheless, this section advocates a genuine diffusion process on stochastic, time-slotted, interference networks by showing that the HD *fluid limit* follows combinatorial 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 [19], [20] 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.

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 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 $\widehat{f_{ij}}(n)$ be the predicted number of packets that the link ij would transmit if it were activated in the slot n, and form the vector $\widehat{\boldsymbol{f}}(n)$. 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

$$\boldsymbol{f}^{\rm tot}(n) = \sum\nolimits_{\boldsymbol{\pi}\in\Pi} \sum\nolimits_{k=1}^{n} \big(T_{\boldsymbol{\pi}}(k) - T_{\boldsymbol{\pi}}(k-1)\big)\big(\boldsymbol{\pi} \odot \widehat{\boldsymbol{f}}(k)\big)$$

where $\odot$ denotes the entrywise product.

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

$$\boldsymbol{q}_\circ(t) = \boldsymbol{q}_\circ(0) + \boldsymbol{a}_\circ^{\rm tot}(t) - \boldsymbol{B}_\circ \boldsymbol{f}^{\rm tot}(t) \tag{27}$$

$$\dot{\boldsymbol{f}}^{\rm tot}(t) = \sum\nolimits_{\boldsymbol{\pi}\in\Pi} \dot{T}_{\boldsymbol{\pi}}(t)\big(\boldsymbol{\pi} \odot \widehat{\boldsymbol{f}}(t)\big) \tag{28}$$

$$\underset{\boldsymbol{\pi}\in\Pi}{\dot{T}_{\boldsymbol{\pi}}(t)} = \begin{cases} 1 & \text{if } \boldsymbol{\pi} \text{ is chosen at time } t \\ 0 & \text{otherwise} \end{cases} \tag{29}$$

$$\sum\nolimits_{\boldsymbol{\pi}\in\Pi} T_{\boldsymbol{\pi}}(t) = t \ \text{ with } T_{\boldsymbol{\pi}}(t) \text{ nondecreasing.} \tag{30}$$

The equality (28) entails the existence of a $\delta > 0$ such that

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

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

To get a deterministic first-order approximation, all stochastic input variables and all time-varying system parameters are replaced by their expected time average values. Therefore,

$$\boldsymbol{a}_\circ^{\rm tot}(t) = \overline{\boldsymbol{a}}_\circ\, t. \tag{31}$$

Likewise, the cost factor $\rho_{ij}(n)$ is replaced by $\overline{\rho_{ij}}$ and the capacity $\mu_{ij}(n)$ is replaced by $\overline{\mu_{ij}}$ for each link $ij \in \mathcal{E}$.

Particular equations: While (27)–(31) hold for any stable network operating under an arbitrary non-idling control policy, each policy determines $\widehat{\boldsymbol{f}}(t)$ and $T_{\boldsymbol{\pi}}(t)$ in its own particular way. Specifically, referring to (8)–(11), HD enforces

$$\widehat{\boldsymbol{f}}(t) \overset{\rm HD}{=} \min\Big\{\, \overline{\boldsymbol{\Phi}}\big(\boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t)\big)^+,\ \overline{\boldsymbol{\mu}}\, \Big\} \tag{32}$$

$$\boldsymbol{w}(t) \overset{\rm HD}{=} \widehat{\boldsymbol{f}}(t) \odot \big(\, 2\, \overline{\boldsymbol{\Phi}}\, \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ(t) - \widehat{\boldsymbol{f}}(t)\, \big) \tag{33}$$

$$\boldsymbol{\pi}(t) = \arg\max\nolimits_{\boldsymbol{\pi}\in\Pi}\ \boldsymbol{\pi}^\top \boldsymbol{w}(t) \tag{34}$$

where $\min$ is taken entrywise and $\boldsymbol{w}(t)$ is the vector of weights assigned by HD policy to each link at time t.

Theorem 4: Define the *HD fluid model* as the collection of deterministic continuous-time equations (27)–(34). Then on a network controlled by the HD policy, every fluid limit $\tilde{\boldsymbol{X}}(t) = \big(\tilde{\boldsymbol{q}}_\circ(t), \tilde{\boldsymbol{f}}^{\rm tot}(t), \tilde{T}_{\boldsymbol{\pi}}(t)\big)$ satisfies the HD fluid model.

With the transmission prediction (32) and the link weighing (33), the immediate result is that the HD key property of Th. 1 in the packet-level is valid in the rate-level too.

Corollary 1: In the fluid limit, the HD control policy maximizes the continuous-time $\boldsymbol{f}$-controlled functional

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

subject to network constraints.

B. Thermodynamic-Like Packet Routing

Consider a wireless network with packets being routed under HD policy with $\beta \in [0,1]$ (microscopic flow). Thus at every timeslot, the policy activates a particular set of links to send a specific number of packets over them. Obviously, each link transmits packets at some slots and is switched off at some other slots. We claim that looking at the average flow in limit (macroscopic flow), it takes the form of the heat flow on the underlying directed graph with suitably-weighted edges.

To formalize this claim, consider a directed graph 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 $a_i(n)$ on the wireless network a corresponding heat source with intensity $\overline{a_i}$ on the graph, and set the node corresponding to the destination node on the wireless network as the heat sink. The flow of heat on the graph is governed by (25)–(26), providing us with the following *reference heat model:*

$$\boldsymbol{f}^\star(t) = \overline{\boldsymbol{\Phi}}\,\max\{\boldsymbol{0},\ \boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ^\star(t)\} \tag{35}$$

$$\begin{aligned} \dot{\boldsymbol{q}}_\circ^\star(t) &= -\vec{\boldsymbol{L}}_\circ^\star\,\boldsymbol{q}_\circ^\star(t) + \overline{\boldsymbol{a}_\circ} \\ \vec{\boldsymbol{L}}_\circ^\star &:= \boldsymbol{B}_\circ \overline{\boldsymbol{\Phi}}\,\boldsymbol{B}_\circ^\top\,\mathrm{diag}\big(\mathbb{I}_{\boldsymbol{B}_\circ^\top \boldsymbol{q}_\circ^\star(t) \succ \boldsymbol{0}}\big). \end{aligned} \tag{36}$$

Then the next theorem shows that the HD fluid model complies with the above reference heat model. Note that $\overline{\phi_{ij}}$ depends not only on the average link cost-factor $\overline{\rho_{ij}}$, but also on the Lagrange control parameter β, i.e. varying β leads to different edge weights on the graph, and so to different graph topology.

Theorem 5: Consider a wireless network with the reference heat model (35)–(36), and suppose that the heat flow (35) meets the link capacity constraint (15). Then the HD fluid model (27)–(34) with $\beta \in [0,1]$ asymptotically converges to the nonlinear graph heat equations (35)–(36). In particular, the HD fluid model with $\beta = 0$ converges to the heat equations on an unweighted directed graph, and the HD fluid model with $\beta = 1$ converges to the heat equations on a weighted directed graph with the edge weights $1/\overline{\rho_{ij}}$.

Remark 2: The assumption that the reference heat flow (35) satisfies the link capacity constraint (15) indeed guarantees that the flow of heat on each link remains lower than the expected time average of link capacity $\overline{\mu_{ij}}$. Thus it guarantees that for a given arrival rate vector $\overline{\boldsymbol{a}_\circ}$, it is possible in principle to stabilize the network such that the fluid limit asymptotically follows the directed graph heat equations. Also worth to observe that the assumption intrinsically deals with network topology, arrival rates, link capacities, and link cost-factors, bearing a theoretically deep and comprehensive concept.

VIII. Routing Cost Minimizing Policy at $\beta = 1$

To establish the second pillar of HD Pareto optimality, this section shows, via Dirichlet Principle, that the routing cost $\overline{R}$ 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 β-dependent optimization problem

$$\begin{aligned} \text{Minimize:} &\quad \sum\nolimits_{ij\in\mathcal{E}} \overline{(f_{ij})^2/\phi_{ij}} \\ \text{Subject to:} &\quad \text{Throughput optimality} \end{aligned} \tag{37}$$

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

A. Classic Dirichlet Principle

Consider the classical heat diffusion equations (17)–(19) subject to constant heat sources $A(\boldsymbol{z})$. In steady-state thermal conduction, the amount of heat entering any region of the manifold 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 will remain uniformly zero. This leads to the classic Poisson equation

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

which formulates the stationary heat transfer by substituting zero for the time derivative of temperature in (19). Then Dirichlet Principle states that the Poisson equation has a unique solution that minimizes the Dirichlet energy

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

among all twice differentiable functions $Q(\boldsymbol{z})$ 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 of coboundary operator $-\boldsymbol{B}^\top$. Fixing $q_d(t) = 0$ yields

$$-\boldsymbol{L}_\circ\,\boldsymbol{q}_\circ + \boldsymbol{a}_\circ = \boldsymbol{0}$$

which correctly realizes (24) in steady-state condition. Like the classic case, this equation has a unique solution that minimizes the combinatorial Dirichlet energy

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

C. Dirichlet Principle on Directed Graphs

Essentially, the Poisson equation on a *directed* graph should capture the steady-state behavior of nonlinear diffusion process (26) subject to constant heat sources $\boldsymbol{a}_\circ$, leading to

$$-\vec{\boldsymbol{L}}_\circ\,\boldsymbol{q}_\circ + \boldsymbol{a}_\circ = \boldsymbol{0}. \tag{38}$$

The difficulty arises from the fact that contrary to the linear Laplacian $\boldsymbol{L}_\circ$ on undirected graphs, the $\vec{\boldsymbol{L}}_\circ$ is an operand-dependent operator that retains neither linearity nor symmetricity. Thus the easy way of proving Dirichlet Principle on undirected graphs is ceased to exist here as we can not claim that $\vec{\boldsymbol{L}}_\circ \boldsymbol{q}_\circ$ in (38) is the directional derivative of $\frac{1}{2}\,\boldsymbol{q}_\circ^\top \vec{\boldsymbol{L}}_\circ \boldsymbol{q}_\circ$. Nevertheless, the next theorem extends the concept of Dirichlet Principle to directed graphs with nonlinear Laplacian.

Theorem 6: On a directed graph with the nonlinear Dirichlet Laplacian $\vec{\boldsymbol{L}}_\circ$, the Poisson equation (38) has a unique solution that minimizes the Dirichlet-like energy

$$\vec{E}(\boldsymbol{q}_\circ) := \frac{1}{2}\,\boldsymbol{q}_\circ^\top \vec{\boldsymbol{L}}_\circ\,\boldsymbol{q}_\circ - \boldsymbol{q}_\circ^\top \boldsymbol{a}_\circ. \tag{39}$$

D. Dirichlet Routing Cost Minimization

The framework of Th. 6 is not yet aligned with what we need for the optimization problem (37). The next theorem resolves this incongruity by showing that minimizing the functional (39) is indeed the dual of minimizing energy dissipation on the directed graph with a zero duality gap.

Theorem 7: Considering the Poisson equation (38), minimization of the Dirichlet-like energy (39) is equivalent to solving the constrained optimization problem

$$\begin{array}{rl} \text{Minimize:} & \vec{E}_R(\boldsymbol{f}) := \boldsymbol{f}^\top \mathrm{diag}(\boldsymbol{\sigma})^{-1} \boldsymbol{f} \\ \text{Subject to:} & \boldsymbol{B}_\circ \boldsymbol{f} = \boldsymbol{a}_\circ \end{array}$$

which represents total energy dissipated on the graph.

It is worth to compare Th. 7 with the celebrated law of minimum dissipative energy on electrical networks. 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, which brings circuit theory and wireless networking together.

Theorem 8: Consider a wireless network with the reference heat model (35)–(36), and suppose that the heat flow (35) meets the link capacity constraint (15). Then the HD policy with $\beta \in [0,1]$ solves the β-dependent constrained optimization problem (37). In particular, supposing that the heat flow (35) with $\overline{\phi_{ij}} = 1/\overline{\rho_{ij}}$ meets the link capacity constraint (15), the HD policy with $\beta = 1$ solves the Dirichlet routing cost minimization problem (2).

IX. 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.

We have shown that HD with $\beta = 0$ minimizes the average network delay $\overline{Q}$ within the class of all routing algorithms that depend only on current queue congestion and current channel states (optimization problem (3)). We have also shown that under the assumption of Th. 8 on link capacities, HD with $\beta = 1$ strictly minimizes the Dirichlet routing cost $\overline{R}$ among all stabilizing routing algorithms (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 and 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.

Theorem 9: Defining $\overline{Q}$ 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 routing algorithms that act based only on current queue backlogs and current channel states, and suppose that this region has a convex Pareto boundary. Also consider the reference heat model (35)–(36) with $\overline{\phi_{ij}} = 1/\overline{\rho_{ij}}$, and suppose that the heat flow (35) meets the link capacity constraint (15). Then the HD policy with $\beta \in [0,1]$ operates on the Pareto boundary of $(\overline{Q}\,, \overline{R})$ region, solving the Pareto optimization problem (4).

X. Conclusion

We have introduced the new Heat-Diffusion (HD) routing policy that can provide a Pareto optimal tradeoff between average network delay and Dirichlet routing cost. In particular, HD reduces the Dirichlet routing cost to its minimum feasible value, and within an important class of routing algorithms, achieves the minimum average delay. 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 to analyze and optimize wireless networking.

Acknowledgment

The authors would like to thank Professor Michael Neely at USC for his helpful feedback. This work has been supported in part by the NSF grants CNS-1017881 and CNS-1049541.

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] L. Bui, S. Sanghavi, R. Srikant. Distributed link scheduling with constant overhead. *Tr. Net.,* 17:1467–1480, 2009.
[9] L. Jiang, J. Walrand. Approaching throughput-optimality in distributed CSMA scheduling algorithms with collisions. *Tr. Net.,* 19:816–829, 2011.
[10] T. Ji, E. Athanasopoulou, R. Srikant. Optimal scheduling policies in small generalized switches. *INFOCOM,* 2009.
[11] L. Bui, R. Srikant, A. Stolyar. A novel architecture for delay reduction in back-pressure scheduling algorithm. *Tr. Net.,* 19:1597–1609, 2011.
[12] B. Ji, C. Joo, N. Shroff. Delay-based back-pressure scheduling in multi-hop wireless networks. *INFOCOM,* 2011.
[13] L. Huang, S. Moeller, M. Neely, B. Krishnamachari. LIFO-backpressure achieves near optimal utility-delay tradeoff. *WiOpt,* 2011.
[14] M. Alresaini, M. Sathiamoorthy, B. Krishnamachari, M. Neely. Back-pressure with adaptive redundancy (BWAR). *INFOCOM,* 2012.
[15] D. Xue, R. Murawski, E. Ekici. Distributed utility-optimal scheduling with finite buffers. *WiOpt,* 2012.
[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] J. Dai. Stability of fluid and stochastic processing networks. *MaPhySto Misc. Pub.,* no. 9, 1999.
[20] M. Bramson. Stability of queueing networks. *Prob. Surv.,* 169-345, 2008.
[21] X. Lin, N. Shroff. Joint rate control and scheduling in multihop wireless networks. *CDC,* 2004.
[22] M. Neely, E. Modiano, C. Li. Fairness and optimal stochastic control for heterogeneous networks. *INFOCOM,* 2005.
[23] A. Eryilmaz, R. Srikant. Fair resource allocation wireless networks using queue-length-based scheduling and congestion control. *INFOCOM,* 2005.
[24] A. Stolyar. Maximizing queueing network utility subject to stability: greedy primal-dual algorithm. *Que. Sys.,* 50:401-457, 2005.
[25] R. Berry, R. Gallager. Communication over fading channels with delay constraints. *Tr. Info. Theo.,* 48:1135–1149, 2002.
[26] D. Arnold, R. Falk, R. Winther. Finite element exterior calculus: from Hodge theory to numerical stability. *Amer. Math. Soc.,* 47:281–354, 2010.
[27] R. Banirazi, E. Jonckheere, B. Krishnamachari. Heat-diffusion: Pareto optimal dynamic routing for time-varying wireless networks. *USC Tech. Report 2013,* http://anrg.usc.edu/www/papers/HDrouting.pdf.
[28] R. Banirazi, E. Jonckheere, B. Krishnamachari. Heat Diffusion Algorithm for Resource Allocation and Routing in Multihop Wireless Networks. *Globecom,* 2012.