

Boltzmann Meets Nash: Energy-Efficient Routing in Optical Networks under Uncertainty

Panayotis Mertikopoulos, *Member, IEEE*, Aris L. Moustakas, *Senior Member, IEEE*, and Anna Tzanakaki, *Senior Member, IEEE*

Abstract

Motivated by the massive deployment of power-hungry data centers for IT service provisioning, we examine the problem of routing in optical networks with the aim of minimizing traffic-driven power consumption. To tackle this issue, routing must take into account energy efficiency as well as capacity considerations; moreover, in rapidly-varying network environments, this must be accomplished in a real-time, distributed manner that remains robust in the presence of random disturbances and noise. In view of this, we derive a pricing scheme whose Nash equilibria coincide with the network's socially optimum states, and we propose a distributed learning method based on the Boltzmann distribution of statistical mechanics. Using tools from stochastic calculus, we show that the resulting Boltzmann routing scheme exhibits remarkable convergence properties under uncertainty: specifically, the long-term average of the network's power consumption converges within ε of its minimum value in time which is at most $\tilde{O}(1/\varepsilon^2)$, irrespective of the fluctuations' magnitude; additionally, if the network admits a strict, non-mixing optimum state, the algorithm converges to it – again, no matter the noise level. Our analysis is supplemented by extensive numerical simulations which show that Boltzmann routing can lead to a significant decrease in power consumption over basic, shortest-path routing schemes in realistic network conditions.

Index Terms

Boltzmann routing; energy efficiency; fluctuations; optical networks; uncertainty; Nash equilibrium.

I. INTRODUCTION

In recent years, the Internet has become the cornerstone of global telecommunications, to the point that it now plays an integral part in sustaining the growth of the world economy. Specifically, as more and more users are using the Internet for backing up data, video on demand (VoD), device synchronization and other IT services, data-hungry applications are experiencing a wildfire growth that affects almost all aspects of modern-day, always-connected societies. As a result, optical networks must grow fast enough so as to avoid experiencing a “capacity crunch” that will cripple their ability to handle the requested traffic in a reliable and efficient manner [3].

P. Mertikopoulos is with the French National Center for Scientific Research (CNRS) and the Laboratoire d'Informatique de Grenoble (LIG), F-38000, Grenoble, France. A. L. Moustakas and A. Tzanakaki are with the Physics Department of the National & Kapodistrian University of Athens, Athens, Greece. Part of this work appeared in the conference paper [1] and the report [2].

At the same time however, energy consumption in data networks must also be kept in check: in total, power consumption from information and communications technology (ICT) applications has been estimated to 4% of the global energy consumption [4], with a steeply rising trend as we progress toward the Internet of Things (IoT) era. In particular, just in the United States, Internet equipment consumes a staggering 8% of the total energy with a predicted growth of 50% within a decade [5]. Therefore, in addition to the always-present requirements of minimizing latencies and maximizing throughput, there is a key developing need to operate optical networks in an energy-aware way that minimizes power expenditure to the bare minimum. This last goal becomes especially prominent when taking into account the power needs of large switching and data centers (DCs) which are rapidly becoming a dominant term in the power consumption equation due the huge amounts of traffic that they service [6, 7].

In this paper, motivated by the massive deployment of power-hungry data centers described above, we examine the problem of routing traffic with the aim of minimizing the network’s traffic-driven power consumption. To achieve this, centralized optimization is not an option, because it is neither scalable nor flexible. Instead, one needs distributed load management as close as possible to the link level while ensuring that routing choices strike an optimal balance between “greener”, more efficient – but possibly congested – network facilities and “hungrier”, legacy resources that may be under-utilized. To tackle this issue, we derive a game-theoretic pricing scheme [8] which reconciles the selfish solution concept of (nonatomic) Nash equilibrium with social optimality (i.e. states that minimize global power consumption in the network). Building on this equivalence, we further propose a distributed learning scheme based on the Boltzmann distribution of statistical mechanics, which bears strong ties to replicator-based methods [9–12] and the class of exponential learning schemes used in finite (atomic) games [13–15].

Exploiting these links, we are able to show that the proposed routing scheme converges rapidly to a socially efficient state – in practice, within a few iterations. However, a key limiting factor in these considerations is that data traffic may exhibit significant spatiotemporal fluctuations because different applications have highly variable characteristic time scales. In addition, obtaining perfect information on the energy consumption of each switch/data center is a very challenging task in large networks due to real-time propagation delays, hierarchical topology aggregation and update sparsity (so as to keep the signaling overhead in check) [16–18]. This uncertainty is exacerbated further in highly dynamic networks where state information becomes obsolete very fast; accordingly, it is not clear whether a routing scheme designed for static, deterministic networks can maintain its convergence properties under uncertainty and noise.

To account for this crucial challenge, we focus on the performance of the proposed Boltzmann routing method in the presence of unpredictable stochastic disturbances. Using tools from stochastic analysis and Itô calculus [19, 20], we show that Boltzmann routing is exceptionally robust in this regard: first, by finetuning the choice of the method’s so-called *inverse temperature*, we show that the network’s long-term average consumption converges within ε of its minimum value in at most $\tilde{O}(1/\varepsilon^2)$ time, irrespective of the fluctuations’ magnitude. Moreover, if the network admits a strict, non-mixing optimum state, the method converges to it almost surely (a.s.), no matter the degree of volatility. Finally, if the network instead admits an interior, fully-mixing optimum state, the proposed scheme remains arbitrarily close to said state with probability arbitrarily close to 1, provided that the scheme’s inverse temperature parameter is

taken sufficiently small.

The performance of Boltzmann routing is validated by means of extensive numerical simulations modeling an optical data network deployed over the largest metropolitan centers of the continental US. Our results show that Boltzmann routing represents a scalable and flexible method that reduces traffic-driven energy consumption by a factor of 40% over simple shortest-path/closest-destination routing modes, even under high degrees of volatility and uncertainty.

Paper outline

In Section II, we present our network energy consumption model and a Pigouvian pricing scheme based on the theory of nonatomic congestion games. Subsequently, in Section III, we introduce the proposed Boltzmann routing method, first in a deterministic and then in a fully stochastic environment. Our main theoretical results (outlined above) are presented in Section IV, in both a deterministic and a stochastic context. Finally, Section V contains our numerical simulations in practical optical networks with diverse link, switching, and data center specifications. To streamline our presentation, technical proofs have been relegated to a series of appendices at the end.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we describe our model for power consumption in optical networks where traffic demands are routed towards a set of service provisioning infrastructures – the network’s *data centers* (DCs). After a few preliminaries in Section II-A, we discuss the specifics of power consumption in Section II-B and we present our game-theoretic model and pricing scheme in Section II-C.

A. Preliminaries

Following the standard traffic model of [21–25], we begin with a finite undirected graph with vertex set \mathcal{V} and edge set \mathcal{E} , describing respectively the network’s *nodes* and *links*. We assume that traffic in the network is generated in continuous streams by a set of *source nodes* $\mathcal{S} \subseteq \mathcal{V}$, each node $s \in \mathcal{S}$ generating traffic demands at rate $\varrho_s > 0$ towards a set of *destination nodes* $\mathcal{D}_s \subseteq \mathcal{V}$. In our setting, the traffic generated at each node represents the aggregation of the demands of a very large number of infinitesimal, nonatomic users – for instance, a city’s smartphone users accessing a cloud-based service. By the same token, the destination set \mathcal{D}_s represents the network’s data centers (or a subset thereof), so each traffic element generated at source s is to be delivered in an *anycast* fashion to any one of the destination nodes in \mathcal{D}_s (for a schematic representation, see Fig. 1).¹

Based on the routing choices of each individual traffic element, this traffic is split over a small, relevant set \mathcal{A}_s of *paths* (or *routes*) that have been chosen offline so as to join source $s \in \mathcal{S}$ to the associated destination set \mathcal{D}_s (for

¹This generic setup can easily account for a variety of different delivery semantics, such as unicast or multicast routing. For instance, in the unicast case, each infinitesimal traffic element generated at s would simply be addressed to a *specific* destination node in \mathcal{D}_s instead of any $d \in \mathcal{D}_s$ [23, 24]. Essentially, switching to a different delivery scheme is just a matter of notation; given our motivation, we focus on anycast routing throughout, but our results also apply to these different routing schemes.

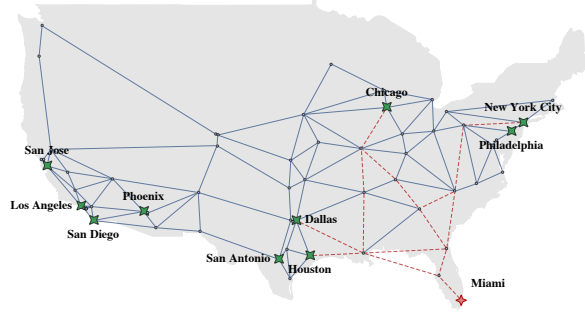


Fig. 1. A typical network topology with source nodes and data centers deployed over the continental US. For illustration purposes, we have highlighted a source node $s \in \mathcal{S}$ (red diamond) and its routing choices $\alpha \in \mathcal{A}_s$ (dashed lines).

instance, all paths with the lowest hop-count). As such, if x_α denotes the amount of traffic routed from source $s \in \mathcal{S}$ via the path $\alpha \in \mathcal{A}_s$, the corresponding *flow profile* $x_s = (x_\alpha)_{\alpha \in \mathcal{A}_s}$ will be an element of the scaled simplex

$$\mathcal{X}_s = \{x_s \in \mathbb{R}^{\mathcal{A}_s} : \sum_{\alpha \in \mathcal{A}_s} x_\alpha = \varrho_s \text{ and } x_\alpha \geq 0 \text{ for all } \alpha \in \mathcal{A}_s\}. \quad (1)$$

Coalescing all these individual flow distributions into a single state variable, the aggregate profile $x = (x_s)_{s \in \mathcal{S}} \in \prod_s \mathbb{R}^{\mathcal{A}_s}$ will be referred to as the network's *state* and the set $\mathcal{X} \equiv \prod_s \mathcal{X}_s$ of all such states will be called the *state space* of the network. For concision, it will often be convenient to write $x = (x_\alpha)_{\alpha \in \mathcal{A}}$ where $\mathcal{A} = \bigcup_s \mathcal{A}_s$ denotes the set of all utilized paths in the network; since a path $\alpha \in \mathcal{A}$ uniquely determines its source and destination nodes, there is no risk of confusion.

Now, if the network is at state $x \in \mathcal{X}$, the *load* w_e on edge $e \in \mathcal{E}$ is defined as the sum of all flows through e , namely

$$w_e = \sum_{\substack{\alpha \in \mathcal{A} \\ \alpha \ni e}} x_\alpha = \sum_{\alpha \in \mathcal{A}} \Pi_{e\alpha} x_\alpha, \quad (2)$$

where $\Pi = (\Pi_{e\alpha})_{e \in \mathcal{E}, \alpha \in \mathcal{A}}$ denotes the *edge-path incidence matrix* of the network, viz.

$$\Pi_{e\alpha} = \begin{cases} 1 & \text{if } e \in \alpha, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

With this in mind, we associate to each edge $e \in \mathcal{E}$ a *consumption* (or *cost*) *function* $c_e: \mathbb{R}_+ \rightarrow \mathbb{R}$ reflecting the negative externalities induced on edge e due to its traffic load w_e . In realistic network scenarios, the externalities incurred on a given edge $e \in \mathcal{E}$ increase (or, at best, do not decrease) with the underlying load w_e ; furthermore, these externalities accrue at an increased rate as the load increases. On that account, our only assumption will be:

Assumption 1. *The network's consumption functions $c_e: \mathbb{R}_+ \rightarrow \mathbb{R}$ are convex and nondecreasing.*

Remark 1. The implications of Assumption 1 in the context of power consumption in optical networks are detailed in the following section; we only state here that this hypothesis is sufficiently broad so as to accommodate a wide variety of congestion-limited frameworks, ranging from latency minimization [21, 23, 25] to urban traffic management

[22, 24], etc.² For this reason, all our analysis will be carried out under the general umbrella of Assumption 1.

With all this at hand, we obtain the global *consumption minimization problem*

$$\begin{aligned} \text{minimize} \quad & C(x) = \sum_{e \in \mathcal{E}} c_e(w_e) \\ \text{subject to} \quad & w = \Pi x, \quad x \in \mathcal{X}. \end{aligned} \tag{CM}$$

In words, the network states x^* that solve (CM) minimize the aggregate consumption $C(x)$ over the entire network; as such, the solutions of (CM) will be called the network's *socially optimum states*. As we show below, the set of such states has a particularly simple structure under Assumption 1:

Proposition 1. *The solution set \mathcal{X}^* of (CM) is nonempty, convex and compact. Moreover, if the network's consumption functions are strictly increasing and convex, every solution of (CM) induces the same load profile $w^* = (w_e^*)_{e \in \mathcal{E}}$; finally, if the utilized paths $\alpha \in \mathcal{A}$ are also linearly independent (in the sense that the network's path-edge incidence matrix Π is invertible on \mathcal{X}), the global minimization problem (CM) admits a unique solution $x^* \in \mathcal{X}$.*

Proof: The first part of our claim is a classic result that follows immediately from the compactness of \mathcal{X} and the convexity of each c_e , $e \in \mathcal{E}$ – see e.g. [23, 26]. For our second assertion, simply note that if the network's cost functions are strictly convex and increasing, C is strictly convex as a function of w , so it admits a unique solution w^* . Hence, if Π is invertible on \mathcal{X} , the equation $w^* = \Pi x^*$ admits a unique solution $x^* \in \mathcal{X}$, as claimed. ■

The above result shows that the network's set of socially optimum states has some fairly desirable attributes; however, it does not provide a way for the network's users to compute (or otherwise converge to) such a state in a distributed way. This question underlies much of this paper's motivation, so we address it in detail in Sections III and IV.

B. Energy consumption in optical data networks

As we mentioned in the introduction, traffic-driven power consumption in optical data networks occurs at two basic levels: (i) at the *link level* (including all wavelength-related and transmission line elements); and (ii) at the *data center level* (including all costs to service, process and/or store demands at the network's data centers, cooling, etc.) [27, 28].

We will thus consider two types of power consumption functions:

- 1) *Link-based consumption.* At the transmission line, power consumption typically comprises two parts: a) a *traffic-independent* component (due to amplification and other factors); and b) a *traffic-dependent* component, proportional to the number of active wavelengths per link (including the cost of switching, electrical-to-optical and optical-to-electrical conversion, etc.). Assuming that wavelength granularity is sufficiently fine, the number of wavelengths needed to carry a traffic load w is itself proportional to w , leading to the link-based model [28, 29]

$$c_{\text{link}}(w) = A_{\text{link}} + B_{\text{link}} w, \tag{4a}$$

²For instance, assuming M/M/1 service queues in a latency minimization framework, a standard choice for c_e is $c_e(w_e) = w_e/(\mu_e - w_e)$, where μ_e denotes the capacity of edge $e \in \mathcal{E}$ [21, 23].

where the values of A_{link} and B_{link} depend on the power specifications of each link.

- 2) *Destination-based consumption.* At the data center level, power consumption again comprises two parts: *a*) a *traffic-independent* element (due to cooling, infrastructure maintenance, etc.); and *b*) a *traffic-dependent* part, proportional to the induced CPU load – and hence, proportional to the traffic load at each data center. We thus obtain the consumption model [27]:

$$c_{\text{DC}}(w) = A_{\text{DC}} + B_{\text{DC}} w. \quad (4b)$$

Of course, the consumption model (4b) is node-based, so the edge-based formulation of the previous section does not immediately apply. However, since data centers are *terminal* destination nodes, this can be remedied as follows: First, adjoin to each data center node d a virtual node d' and a virtual edge $e = (d, d')$; assume further that all traffic reaching d is rerouted to d' via e . Then, the load at d can be represented by the load on $e = (d, d')$ and the cost function of this fictitious edge reflects the energy consumption at the data center.

In addition to the above, in practical optical networks, each link/data center has a finite capacity, leading to the state constraint

$$w_e \leq L_e \quad \text{for all } e \in \mathcal{E}, \quad (5)$$

where L_e represents the *maximum load* supported on edge e . We thus obtain the *capacity-constrained consumption minimization problem*

$$\begin{aligned} & \text{minimize} && C(x), \\ & \text{subject to} && x \in \mathcal{X}_0, \end{aligned} \quad (\text{CM}_0)$$

where

$$\mathcal{X}_0 = \{x \in \mathcal{X} : w_e = (\Pi x)_e \leq L_e \text{ for all } e \in \mathcal{E}\} \quad (6)$$

denotes the network's *capacity-constrained feasible region*. Clearly, if the traffic generation rates ϱ_s are very high, the set \mathcal{X}_0 may be empty, in which case (CM₀) does not admit a feasible solution. It will therefore be useful to introduce the following ε -relaxation of (CM₀):

$$\begin{aligned} & \text{minimize} && C_\varepsilon(x) = \sum_{e \in \mathcal{E}} c_e^\varepsilon(w_e) \\ & \text{subject to} && w = \Pi x, \quad x \in \mathcal{X}, \end{aligned} \quad (\text{CM}_\varepsilon)$$

where the ε -adjusted consumption functions c_e^ε are given by

$$c_e^\varepsilon(w) = \begin{cases} c_e(w) & \text{if } w \leq L_e, \\ c_e(L_e) + (w - L_e)/\varepsilon & \text{if } w \geq L_e. \end{cases} \quad (7)$$

In words, the consumption model (7) of (CM_ε) coincides with that of (CM₀) up to the edge's maximum load L_e and then increases sharply with slope $1/\varepsilon$. Since c_e^ε is strictly increasing, any feasible solution of (CM₀) will also be a solution of (CM_ε). On the other hand, given that \mathcal{X} is compact and nonempty, (CM_ε) always admits a solution, even when (CM₀) does not; in this case, the solutions of (CM_ε) can be considered as “approximate” solutions to (CM₀) in the limit $\varepsilon \rightarrow 0$. We formalize this in the following proposition:

Proposition 2. Let $C_\varepsilon^* = \min\{C_\varepsilon(x) : x \in \mathcal{X}\}$ denote the minimum value of the global minimization problem (CM_ε) , and let C_0^* denote the corresponding quantity for the capacity-constrained problem (CM_0) . Then, $\lim_{\varepsilon \rightarrow 0} C_\varepsilon^* = C_0^*$ (with the standard convention $\min \emptyset = \infty$).

Proof: Write $\mathcal{X}_\varepsilon = \mathcal{X}$ for the feasible set of (CM_ε) , $\varepsilon > 0$. By construction, the set-valued correspondence $\varepsilon \mapsto \mathcal{X}_\varepsilon$, $\varepsilon \geq 0$, is upper hemicontinuous [30] and the objective function of (CM_ε) coincides with that of (CM) on \mathcal{X}_0 . Thus, by a precursor to Berge's maximum theorem [30, Lemma 16.30], it follows that the function $\varepsilon \mapsto C_\varepsilon^*$, $\varepsilon \geq 0$, is lower semicontinuous. Since $\lim_{\varepsilon \rightarrow 0} c_\varepsilon^\varepsilon(w) = \infty$ for all $w > L_\varepsilon$, our claim is immediate. ■

In what follows, we will assume for simplicity that the capacity-constrained feasible region \mathcal{X}_0 of (CM_0) is nonempty (after all, there is nothing to optimize if \mathcal{X}_0 is empty). In this case, using the ε -relaxation (CM_ε) allows us to recover the solutions of (CM_0) in the limit $\varepsilon \rightarrow 0$. Given that the ε -adjusted consumption functions $c_\varepsilon^\varepsilon$ satisfy Assumption 1, this will allow us to apply the general analysis for (CM) , even in the presence of capacity constraints of the form (5).

C. Pigouvian pricing and Nash equilibrium

In a fully distributed environment, the system designer would like to ensure that the network's nonatomic (infinitesimal) users make socially optimal routing choices while unilaterally minimizing the price they pay for accessing the network and utilizing its resources. Following Pigou's theory of pricing [31], our approach to achieve this will be to charge agents for the externalities that they induce at a socially optimum state; in so doing, one can then ensure that such states constitute an equilibrium of the agents' selfish interactions [8, 25].

To make this precise, suppose that a user that sends an infinitesimal amount of traffic dx through edge $e \in \mathcal{E}$ is charged $\pi_e(w_e) dx$ where $\pi_e(w_e)$ is the *price per unit of traffic* on edge e at load w_e . The total amount charged over path $\alpha \in \mathcal{A}$ will then be

$$\pi_\alpha(x) dx = \sum_{e \in \alpha} \pi_e(w_e) dx, \quad (8)$$

so a (nonatomic) user will be satisfied with his routing choice if the charged price $\pi_\alpha(x)$ is the lowest among all other available paths/destinations [32]. Formally:

Definition 1. A state $x^* \in \mathcal{X}$ is at *Nash equilibrium* if, for all paths $\alpha \in \mathcal{A}$ with $x_\alpha^* > 0$, we have

$$\pi_\alpha(x^*) = \min_{\beta \sim \alpha} \pi_\beta(x^*), \quad (\text{NE})$$

where the minimum is taken over all paths $\beta \sim \alpha$ originating at the same source as α . In words, x^* is a *Nash equilibrium* when every nonatomic user chooses the least expensive path.

We will also say that x^* is an *interior* (or *fully-mixing*) Nash equilibrium if $x^* \in \text{int}(\mathcal{X})$, i.e. if it employs all paths $\alpha \in \mathcal{A}$; by contrast, x^* will be called *strict* (or *non-mixing*) if $\arg \min_{\alpha \in \mathcal{A}_s} \pi_\alpha(x^*)$ is a singleton for every source node $s \in \mathcal{S}$.

A standard result in the theory of nonatomic congestion games is that Nash equilibria can be characterized as the

solutions of a certain convex program. Specifically, following [22–24, 26, 33], consider the *potential function*

$$U(x) = \sum_{e \in \mathcal{E}} U_e(w_e) \quad (9)$$

where $w = \Pi x$ and

$$U_e(x) = \int_0^{w_e} \pi_e(w) dw. \quad (10)$$

Then, as was shown in [26, 33], x^* is a Nash equilibrium if and only if it solves the *potential minimization problem*

$$x^* \in \arg \min_{x \in \mathcal{X}} U(x). \quad (\text{PM})$$

Hence, by comparing (PM) and (CM), we obtain the following Pigouvian pricing scheme:

Proposition 3. *If the charged price per unit of traffic on edge e is $\pi_e(w_e) = \partial_{-}c_e(w_e)$, Nash equilibria coincide with the solutions of the global consumption minimization problem (CM).*

Remark 2. In the above, $\partial_{-}c_e(w_e)$ denotes the left derivative of c_e at w_e . This derivative always exists because c_e is assumed convex [34]; if c_e is smooth, we can simply take $\pi_e = c'_e$.

Proposition 3 describes the pricing scheme that the system designer should use so that selfish, cost-minimizing users end up minimizing global power consumption at equilibrium. As such, in the rest of this paper, we will assume that prices are set in accordance to Proposition 3, so that socially optimum states for (CM) coincide with the network’s Nash equilibria under π .

III. BOLTZMANN ROUTING

In this section, we present a distributed routing scheme to attain a socially efficient state based on a logit-type learning rule inspired from statistical mechanics. This scheme consists of two basic steps: First, each source node $s \in \mathcal{S}$ keeps a running “score” for each path starting at s by aggregating the associated end-to-end price over time. Then, every nonatomic user based at s chooses a path (and hence, a destination node $d \in \mathcal{D}_s$) with probability inversely proportional to the exponential of this performance score – in analogy to the Boltzmann distribution of statistical mechanics.

For generality, throughout this section, we focus on the global minimization problem (CM) with consumption functions satisfying Assumption 1; the specifics of energy efficiency in optical data networks are discussed in Section V. To streamline our presentation, we first develop our routing scheme in a deterministic setting with perfect information. Subsequently, we extend our model to a fully stochastic setting where the network is subject to random perturbations (arising e.g. from fluctuations in background traffic, link quality, estimation errors, etc.). We then present our convergence analysis and theoretical results in Section IV.

A. Boltzmann routing

The first step in setting up our routing scheme is to introduce a path’s “score”, interpreted here as a cumulative measure of the path’s price over time. More precisely, if $x(t)$ denotes the state of the system at time t , we define the

score of path $\alpha \in \mathcal{A}$ as

$$y_\alpha(t) = y_\alpha(0) + \int_0^t \pi_\alpha(x(s)) ds, \quad (11)$$

i.e. as the cumulative price of path α over the time interval $[0, t]$ (the value $y_\alpha(0)$ is an arbitrary constant that represents an initial assessment of the path's price).

Clearly, high scores indicate commensurately high prices, so such paths should be selected with low probability. We thus posit that a nonatomic user starting at s selects path $\alpha \in \mathcal{A}_s$ with probability given by the *Boltzmann distribution*

$$P_\alpha(t) \propto \exp(-\eta(t)y_\alpha(t)), \quad (B)$$

where $\eta(t) > 0$ is an *inverse temperature* parameter whose role is explained below. Thus, writing (11) in differential form and invoking the law of large numbers to obtain the flow induced by (B), we obtain the *Boltzmann routing scheme*

$$\begin{aligned} \dot{y}_\alpha &= \pi_\alpha(x), \\ x_\alpha &= \frac{\varrho_s \exp(-\eta y_\alpha)}{\sum_{\beta \in \mathcal{A}_s} \exp(-\eta y_\beta)}, \end{aligned} \quad (BR)$$

where $s \equiv s(\alpha)$ denotes the starting point of α and ϱ_s is the corresponding traffic generation rate. These routing dynamics will be at the core of our paper, so a few remarks are in order:

a) *Basic properties*: The Boltzmann scheme (BR) enjoys the following desirable properties:

(P1) It is *consistent*: $x_\alpha(t) \geq 0$ and $\sum_{\alpha \in \mathcal{A}_s} x_\alpha(t) = \varrho_s$ for all $t \geq 0$, so $x(t)$ is a valid state variable.

(P2) It is *reinforcing*: traffic elements tend to be routed along paths with lower prices.

(P3) It is *stateless*: routing choices do not require knowledge of the network's state.

(P4) It is *distributed*: each node only needs to monitor the price of the paths that start at said node.

We should also note that (BR) has fairly low requirements in terms of computational complexity. Specifically, each node $s \in \mathcal{S}$ only needs to run a cheap update on $|\mathcal{A}_s|$ variables; since the set of paths being utilized by each source node is typically small (for instance, those with the minimum hop count), every node $s \in \mathcal{S}$ only needs to process at most $\mathcal{O}(\mathcal{D}_s) \leq \mathcal{O}(\mathcal{V})$ variables. By comparison, the application of distributed flow-deviation methods [21] would require each node to keep track of $\mathcal{O}(\mathcal{E})$ link variables and then run at each step an $\mathcal{O}(\mathcal{V}^2)$ shortest path algorithm, leading to updates of significantly higher complexity.

b) *Relation to the replicator dynamics*: If we differentiate the network's state variable x in (BR), a straightforward calculation leads to the evolutionary dynamics

$$\dot{x}_\alpha = -\eta x_\alpha \left[\pi_\alpha(x) - \sum_{\beta \sim \alpha} x_\beta \pi_\beta(x) \right] + \frac{\dot{\eta}}{\eta} x_\alpha \left[\log x_\alpha - \sum_{\beta \sim \alpha} x_\beta \log x_\beta \right], \quad (12)$$

where the summation is taken over all paths $\beta \sim \alpha$ with the same source as α , and we have taken $\varrho_s = 1$ for simplicity.

Thus, in the baseline case $\eta = 1$, we obtain the dynamical system

$$\dot{x}_\alpha = -x_\alpha \left[\pi_\alpha(x) - \sum_{\beta \sim \alpha} x_\beta \pi_\beta(x) \right], \quad (RD)$$

which is the classical (multi-population) *replicator equation* of evolutionary game theory [24, 35]. In this way, (BR) can be seen as an extension of the replicator-based routing schemes of [9–12], the key difference being the extra,

price-independent term of the dynamics (12). As we shall see, this additional term will be crucial for the convergence properties of (BR) under uncertainty – an open problem posed by [11].

c) *On the inverse temperature η* : From an algorithmic viewpoint, the role of the inverse temperature parameter $\eta(t)$ in (BR) is to act as an extrinsic weight that “normalizes” the paths’ scores $y(t)$. In particular, for low η , the Boltzmann distribution (B) tends to select paths uniformly; by contrast, for high η , the induced choice probabilities “freeze” down to a hard best-response scheme which routes all traffic along the path with the lowest score. Thus, to counterbalance the $\mathcal{O}(t)$ growth of the variables $y_\alpha(t)$, we will assume throughout that $\eta(t)$ is *nonincreasing* with decay rate slower than $1/t$:

Assumption 2. $\eta(t)$ is C^1 -smooth, nonincreasing, and $\lim_{t \rightarrow \infty} \eta(t)t = \infty$.

In thermodynamic terms, Assumption 2 means that the system is being *heated* over time (instead of being cooled). This comes in stark contrast with simulated annealing and log-linear learning [36] where the system begins at a high sampling temperature ($\eta \approx 0$) and subsequently freezes to very low temperatures ($\eta \rightarrow \infty$) to approach a state of least energy. The reason for taking a heating schedule in (BR) is that the energy levels $y_\alpha(t)$ of the system at hand are not fixed (as in simulated annealing) but, instead, they grow over time. As a result, if the paths’ prices (whose aggregation determines the system’s energy levels) are subject to randomness, freezing the system may lead it to quench prematurely to a suboptimal state; we explore this issue in detail in Section IV.

B. Routing under uncertainty

A key assumption underlying the (deterministic) Boltzmann routing scheme (BR) is that prices are assumed immune to exogenous stochastic fluctuations. In practice however, this assumption often fails: background traffic fluctuations can be quite substantial due to burst-like user demands, open market energy costs are highly volatile (for instance, at certain locations power may be supplied by renewable energy sources whose output depends on the weather), whereas load measurements are typically subject to errors due to delay [16], inaccurate routing information and noise [17], etc.

Starting from (8), we will model such disturbances via the random perturbation model

$$\tilde{\pi}_\alpha(x; t) = \sum_{e \in \alpha} (\pi_e(w_e) + \xi_e(t)), \quad (13)$$

where ξ_e is a zero-mean stochastic process.³ The score y_α of path α will then follow the dynamics

$$\dot{y}_\alpha = \tilde{\pi}_\alpha(x) = \pi_\alpha(x) + \xi_\alpha, \quad (14)$$

where $\xi_\alpha = \sum_{e \in \alpha} \xi_e = \sum_{e \in \mathcal{E}} \Pi_{e\alpha} \xi_e$ denotes the aggregate perturbation over path α . Hence, writing (14) as an Itô (non-anticipative) stochastic differential equation [19, 20, 37], we obtain the *stochastic Boltzmann routing* dynamics

$$\begin{aligned} dY_\alpha &= \pi_\alpha(X) dt + dZ_\alpha, \\ X_\alpha &= \frac{\rho_s \exp(-\eta Y_\alpha)}{\sum_{\beta \in \mathcal{A}_s} \exp(-\eta Y_\beta)}. \end{aligned} \quad (\text{SBR})$$

³At the most basic level, ξ_e can be assumed to be a simple additive white Gaussian noise (AWGN) process.

where, to be consistent with (14), the *path-noise process* Z_α is of the form

$$dZ_\alpha = \sum_{e \in \alpha} \sigma_e dW_e, \quad (15)$$

where $W = (W_e)_{e \in \mathcal{E}}$ is a standard Wiener process (Brownian motion) in $\mathbb{R}^{\mathcal{E}}$ and $\sigma_e \equiv \sigma_e(x, t)$ is the (possibly state- and time-dependent) *volatility coefficient* of the fluctuations on edge $e \in \mathcal{E}$.

The dynamical system (SBR) will be our core stochastic model so we proceed with some explanatory remarks:

d) Assumptions on the noise: The noise model (15) implies that fluctuations are independent across *links*, but not across *paths*. Indeed, a simple calculation reveals the correlation structure

$$dZ_\alpha \cdot dZ_\beta = \sum_{e, e' \in \mathcal{E}} \sigma_e \sigma_{e'} dW_e \cdot dW_{e'} = \sum_{e, e' \in \mathcal{E}} \sigma_e \sigma_{e'} \delta_{ee'} dt = \sum_{e \in \alpha \cap \beta} \sigma_e^2 dt, \quad (16)$$

i.e. *fluctuations along two paths are correlated along their common edges*.⁴ With this in mind, it will be convenient to introduce the *volatility matrix* $\Sigma = (\sigma_{\alpha\beta}^2)_{\alpha, \beta \in \mathcal{A}}$ defined as

$$\sigma_{\alpha\beta}^2 \equiv \sum_{e \in \alpha \cap \beta} \sigma_e^2 = \sum_{e \in \mathcal{E}} \Pi_{e\alpha} \Pi_{e\beta} \sigma_e^2. \quad (17)$$

By construction, the matrix $\Sigma \equiv \Sigma(x, t)$ describes the *quadratic covariation* of the noise process Z in the sense that $dZ_\alpha \cdot dZ_\beta = \Sigma_{\alpha\beta} dt$ [19, 20]. On that account, our only assumption will be:

Assumption 3. *The fluctuations' volatility matrix Σ is bounded: $\sup_{x,t} \|\Sigma(x, t)\| \equiv \sigma_*^2 < \infty$.*

Remark 3. We should clarify here that Assumption 3 means that fluctuations are only bounded in *mean square*; at any given time $t \geq 0$, the disturbances $dZ_e(t)$ could be arbitrarily large.

e) The stochastic replicator dynamics: Even though $X(t)$ is fully determined under (SBR) for all $t \geq 0$, this is an indirect description of the governing dynamics of $X(t)$. To obtain an explicit description of these (stochastic) dynamics, we can follow the same approach as in (12); however, because of the noise, we now have to employ the rules of (Itô) stochastic calculus [19, 20, 37]. This derivation is quite convoluted, so we only present here the end result (proven in Appendix A):

Proposition 4. *The solutions of (SBR) satisfy the stochastic differential equation:*

$$dX_\alpha = -\eta X_\alpha \left[\pi_\alpha(X) - \sum_\beta X_\beta \pi_\beta(X) \right] dt \quad (18a)$$

$$- \eta X_\alpha \left[dZ_\alpha - \sum_\beta X_\beta dZ_\beta \right] \quad (18b)$$

$$+ \frac{\dot{\eta}}{\eta} X_\alpha \left[\log X_\alpha - \sum_\beta X_\beta \log X_\beta \right] dt \quad (18c)$$

$$+ \frac{\eta^2}{2} X_\alpha \left[\sum_{\beta, \gamma} (\delta_{\alpha\gamma} - X_\gamma) (\delta_{\beta\gamma} - 2X_\beta) \sigma_{\beta\gamma}^2 \right] dt, \quad (18d)$$

where all sums are taken over the paths $\beta, \gamma \sim \alpha$ with the same starting point as α , and all traffic generation rates have been taken equal to $\varrho_s = 1$ for simplicity.

⁴More generally, the edge processes W_e could be themselves correlated along different edges; we do not consider this for simplicity.

Despite the complex appearance of (18), each of the constituent terms admits a relatively straightforward interpretation:

- a) The term (18a) drives the process in the baseline, deterministic case $\sigma = 0$, $\eta = \text{constant}$; as such, it coincides with the deterministic replicator dynamics (RD).
- b) The martingale term (18b) reflects the direct impact of the noise on the evolution of $X(t)$.
- c) The term (18c) is due to the temporal variation of the inverse temperature parameter $\eta(t)$ and plays the same role as in (12).
- d) Finally, the term (18d) is the Itô correction induced by the non-anticipative nature of the Itô integral [19, 20]. This term is *price-independent* and does not depend on $\dot{\eta}$, so it does not vanish for constant η ; also, the correlation structure of the noise process Z appears explicitly in (18d) via the volatility matrix Σ (cf. the remarks preceding Assumption 3).

Apart from a vaguely similar structure, there is no overlap between (18) and the stochastic replicator dynamics with “aggregate shocks” that have been studied extensively in evolutionary biology [38–42]. The exponential learning approach of [13, 43] is much closer in spirit but it does not account for the nonlinear nature of the cost functions π_e and the correlation structure of the noise processes Z along paths; we explore these connections in more detail below.

IV. CONVERGENCE ANALYSIS

Our aim in this section will be to analyze the long-term convergence properties of the proposed Boltzmann routing scheme in the presence of noise and uncertainty. To establish a baseline, we begin with the scheme’s convergence properties in a noiseless, deterministic setting:

Theorem 1. *Under (BR), $x(t)$ converges to a solution of the global minimization problem (CM).*

Given the connection of (BR) to the replicator dynamics, Theorem 1 (proven in Appendix B) represents a strengthening of known convergence results for potential [44] and stable games [45]. Specifically, [44] showed that the replicator dynamics (RD) converge to the set of rest points in potential games; however, since the rest points of (RD) are not necessarily solutions of (CM), we cannot use the analysis of [24, 44] to deduce the convergence of (BR) to a globally efficient state. More recently, [12] established the convergence of (RD) to the game’s equilibrium set, but under the implicit assumption that the network’s routes are linearly independent (an assumption that fails automatically if $|\mathcal{A}| > |\mathcal{E}|$ because Π fails to be invertible on \mathcal{X} in this case). Instead, Theorem 1 dispenses with all such assumptions and provides an unconditional convergence result for (BR).

This behavior of (BR) is fairly encouraging, but it is contingent on the absence of fluctuations and uncertainty. If the network is constantly subject to stochastic disturbances, there is no reason to expect that this convergence still holds beyond the “small noise” regime. Nevertheless, as we show below, the proposed Boltzmann routing scheme remains exceptionally robust in the presence of noise: under (SBR), the long-term average of the total network consumption is minimized, *irrespective of the level of uncertainty*. Formally:

Theorem 2. Assume that (SBR) is run with a variable parameter $\eta(t)$ satisfying Assumption 2. Then, the long-term average $\bar{C}(t) = t^{-1} \int_0^t C(X(s)) ds$ of the network's total consumption enjoys the bound:

$$\bar{C}(t) \leq C^* + \frac{\sum_s \log A_s}{\eta(t)t} + \frac{\sigma_*^2}{2t} \int_0^t \eta(s) ds + 2\sigma_*^2 \sqrt{\frac{\log \log t}{t}} + \mathcal{O}(1/t) \quad (a.s.) \quad (19)$$

where C^* is the minimum value of (CM) and $A_s = |\mathcal{A}_s|$ is the number of paths utilized by source $s \in \mathcal{S}$. In particular, if $\eta(t) \rightarrow 0$, we have $\lim_{t \rightarrow \infty} \bar{C}(t) = C^*$ (a.s.).

The proof of Theorem 2 (presented in detail in Appendix C) hinges on the so-called *Kullback–Leibler* (KL) divergence [46–48], defined here as

$$D_{\text{KL}}(x^*, x) = \sum_{s \in \mathcal{S}} \sum_{\alpha \in \mathcal{A}_s} x_\alpha^* \log(x_\alpha^*/x_\alpha), \quad x^*, x \in \mathcal{X}. \quad (20)$$

The KL divergence provides an asymmetric measure of the information-theoretic distance between x^* and x . Using this, we express the average global consumption $\bar{C}(t)$ in terms of this distance, and we then use Itô's lemma [19, 20, 37] to bound it by a vanishing function of t . In so doing, we obtain the three main components of the convergence rate estimate (19): the first represents the convergence rate of the noiseless process (BR), the second is due to the Itô correction (18d), while the third one stems from the law of the iterated logarithm [37].

Importantly, the second term of (19) does not vanish for constant η , explaining the requirement $\lim_{t \rightarrow \infty} \eta(t) = 0$. On the other hand, this requirement can be dropped if there is no noise: for $\sigma_* = 0$ and $\eta = \text{constant}$, the RHS of (19) reduces to $\mathcal{O}(1/t)$. In addition, the form of (19) also highlights a trade-off between more aggressive (slowly decaying) schedules for η and the underlying uncertainty. Specifically, the first (deterministic) term of (19) is decreasing in η while the second (stochastic) one increases with η . Hence, in the absence of noise, it is better to use a large, constant η instead of letting $\eta(t) \rightarrow 0$; however, a constant η may be detrimental under uncertainty.

These considerations can be illustrated by choosing a specific heating schedule of the form $\eta(t) \sim 1/t^a$ for some $a \in (0, 1)$; in this case, we obtain the explicit convergence rate:

Corollary 1. Assume that (SBR) is run with $\eta(t) \sim 1/t^a$ for some $a \in (0, 1)$. Then:

$$\bar{C}(t) = C^* + \begin{cases} \mathcal{O}(1/t^a) & \text{if } 0 < a < \frac{1}{2}, \\ \mathcal{O}\left(\sqrt{\log \log t/t}\right) & \text{if } a = 1/2, \\ \mathcal{O}(1/t^{1-a}) & \text{if } \frac{1}{2} < a < 1. \end{cases} \quad (21)$$

Proof: Simply note that $t^{\max\{a-1, -a\}}$ is the dominant term in (19) for all $a \neq 1/2$. Otherwise, for $a = 1/2$, the first two terms of (19) are both $\mathcal{O}(t^{-1/2})$ and are dominated by the third. ■

Theorem 2 and Corollary 1 make no assumptions for the underlying network or the magnitude of the stochastic perturbations affecting the system. However, the convergence they provide is in terms of the long-term average consumption $\bar{C}(t)$, not the instantaneous network state $X(t)$. As a matter of fact, in the presence of disturbances, $X(t)$ cannot converge with positive probability to a fully-mixing state x^* where a given source node utilizes several paths concurrently: even if x^* is a rest point of (BR), the noise term (18b) does not vanish at x^* , so x^* cannot be stationary

under (SBR). On the other hand, if (CM) admits a *strict, non-mixing* solution x^* (so each source node utilizes a *single* path at Nash equilibrium; cf. Definition 1), we show below that Boltzmann routing converges to x^* *independently of the magnitude of the noise*:

Theorem 3. *Assume that the global minimization problem (CM) admits a strict, non-mixing solution x^* . If (SBR) is run with sufficiently small (constant) η , $X(t)$ converges to x^* (a.s.).*

Remark 4. We should note here that if x^* is a strict, non-mixing solution of (CM), it is the *only* solution of (CM); as such, Theorem 3 shows that Boltzmann routing converges (a.s.) to the network’s *unique* globally efficient state.

The proof of Theorem 3 is fairly complicated, so we relegate it to Appendix D. In a nutshell, it consists of showing that *a)* if $X(t)$ remains close to x^* for all time, it is eventually attracted to it; and *b)* $X(t)$ gets trapped in arbitrarily small neighborhoods of x^* with controllably high probability (a much harder result which relies on an application of Girsanov’s theorem [19, 20] to estimate the probability that a Wiener process with positive drift attains a given negative level in finite time).

From a practical viewpoint, the importance of Theorem 3 is that it provides a global convergence result for constant (but small) η , irrespective of the noise level. This relaxes even further the “vanishing η ” requirement of Theorem 2 and is owed to the existence of a *strict, non-mixing* solution. On the flip side, if the network does not admit such a solution, $X(t)$ cannot converge with positive probability – even though the long-term average consumption $\bar{C}(t)$ does. Our next result shows that if the network admits an (isolated) interior solution x^* , learning with sufficiently small η allows $X(t)$ to remain arbitrarily close to x^* with arbitrarily high probability:

Theorem 4. *Let $x^* \in \text{int}(\mathcal{X})$ be an isolated, interior solution of (CM). Then, for all $\varepsilon, \delta > 0$, the scheme (SBR) can be run with a sufficiently small (constant) parameter $\eta \equiv \eta(\varepsilon, \delta)$ such that*

$$\mathbb{P}(\|X(t) - x^*\| \leq \delta \text{ for all sufficiently large } t) \geq 1 - \varepsilon. \quad (22)$$

Theorem 4 implies that if the inverse temperature η of (SBR) is small enough, then, after a certain amount of time, the system will be arbitrarily close to equilibrium with probability arbitrarily close to one. Thus, even though the network’s state does not converge almost surely, it will still spend most of the time close to equilibrium.

V. NUMERICAL RESULTS

In this section, we assess the performance of the proposed Boltzmann routing scheme in practical scenarios via numerical simulations. For clarity, we only present here a representative subset of these results but our conclusions apply to a wide range of optical network parameters and specifications.

Our setup is as follows (for an overview, see Table I): we consider an optical network deployed over the continental US with nodes and links as in Fig. 1 (for clarity, we only plotted the 50 largest metropolitan areas) [49]. At each link, data is transmitted over 80 wavelength-division multiplexing (WDM) channels spaced at 50 GHz, each with a carrying capacity of 10 Gb/s. To mitigate signal losses, each link carries an amplifier per 80 km of fiber length,

TABLE I
NETWORK SIMULATION PARAMETERS

Parameter	Value	Parameter	Value
Network topology	Continental US [49]	Source traffic	[200, 400] Gb/s
WDM channels / fiber	80	Channel capacity	10 Gb/s
Amplifier consumption	15 W	Amplifier density	1/(80 km)
Transponder consumption	35 W/chanel	Switch port consumption	0.8 W
DC consumption (zero load)	6.6 kW	DC consumption (full load)	13.2 kW
DC capacity	2 Tb/s	Cooling factor	2

and an additional amplifier at each end; in addition to the amplifiers, power is also consumed at the transponder and switch port levels (cf. Table I) [28]. Source nodes are drawn randomly in the network, each with a traffic rate between 200 Gb/s and 400 Gb/s; this traffic is then directed to a set of data centers, also drawn randomly from the network's nodes. Each data center has a capacity of 2 Tb/s and consumes between 6.6 kW and 13.2 kW, at zero and full load respectively [27]; for diversity, we also populate the network with a number of "legacy" data centers with lower capacity and higher power consumption specifications as indicated in [50].

In Fig. 2, we examine the performance of the proposed Boltzmann routing scheme in terms of power consumption, convergence speed and scalability. Specifically, we consider several different structures for the source nodes' choice sets \mathcal{A}_s : *a*) taking the shortest path to the closest destination ($|\mathcal{A}_s| = 1$); *b*) splitting traffic over the 4 paths with the lowest hop count to the closest destination ($|\mathcal{A}_s| = 4$); *c*) splitting traffic to the 5 closest destinations over shortest paths ($|\mathcal{A}_s| = 5$); and *d*) mixing shortest paths and closest destinations ($|\mathcal{A}_s| = 4 \times 5$). We then run the Boltzmann routing scheme (BR) with the Pigouvian pricing scheme of Proposition 3 ($\pi_e = c'_e$), a heating schedule of the form $\eta(t) \sim t^{-1/2}$ and all scores initialized uniformly at $Y(0) = 0$ (of course, in the single-path, single-destination case, there is nothing to learn).

As shown in Fig. 2(a), Boltzmann routing converges very rapidly (within a few iterations) to a socially efficient network state that solves the corresponding consumption minimization problem (CM). Moreover, we see that splitting traffic over several data centers (as opposed to exclusively targeting the closest data center) leads to significant performance gains, of the order of 100 kW/(Gb/s) – approximately a 40% decrease in power consumption with respect to closest-destination routing. Finally, in Fig. 2(b), we examine the scalability of these performance gains as the number of source nodes increases: in so doing, we observe that *a*) the relative percentage gain of splitting traffic over multiple data centers becomes more prominent for large numbers of sources; and *b*) the Boltzmann routing scheme retains its convergence speed and attains a socially efficient state within 4–5 iterations, even for $S = 100$ source nodes. This means that, in the presence of intermittent traffic fluctuations, the system will be capable of adapting fast enough to changes in the system.

In Fig. 3, we examine the algorithm's robustness in the presence of stochastic fluctuations and capacity constraints. First, in Fig. 3(a), we consider a network with $S = 50$ source nodes and $D = 10$ data centers with power consumption and capacity attributes as in Table I. We further assume that consumption costs are subject to stochastic disturbances

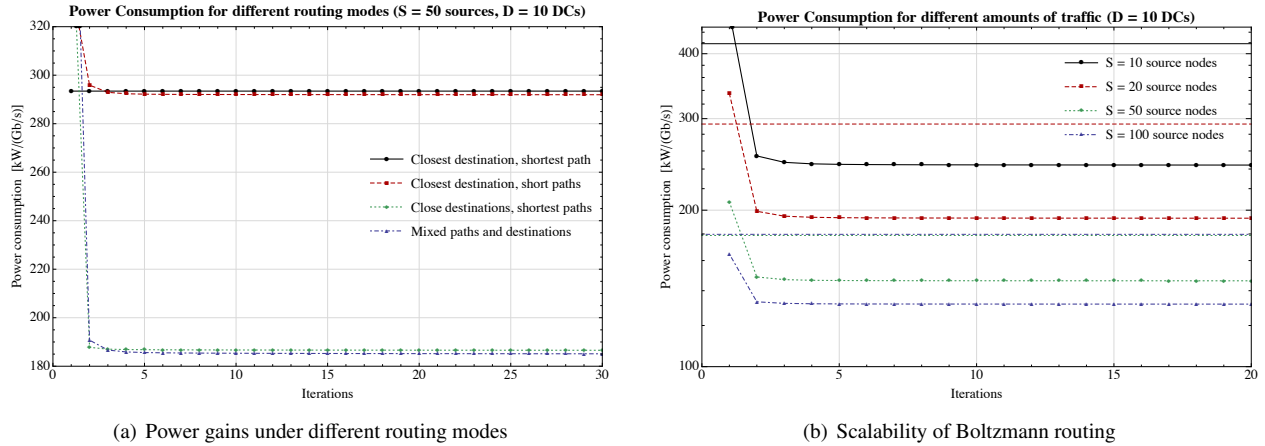


Fig. 2. Gains in power consumption for different routing modes and different numbers of sources. Fig. 2(a) shows that splitting traffic to several destinations leads to a 100 kW/(Gb/s) decrease in power consumption. As we see in Fig. 2(b), these gains become more pronounced for larger number of source nodes (marked/unmarked lines correspond to mixed/closest-destination routing respectively); moreover, Boltzmann routing attains 99% of its overall gain within 4–5 iterations, even for $S = 100$ source nodes.

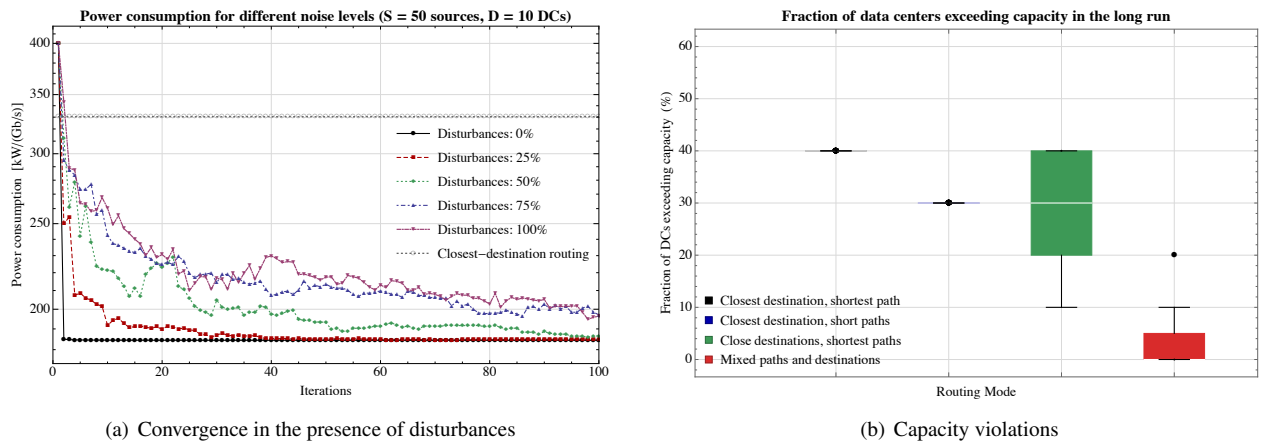


Fig. 3. Long-term behavior of Boltzmann routing in the presence of disturbances (Fig. 3(a)) and capacity constraints (Fig. 3(b)).

and noise with volatility coefficient σ_* equal to a fraction z of the costs' mean value (ranging between 0% and 100%). Remarkably, Boltzmann routing reaches a socially efficient state even under high degrees of uncertainty (up to $z = 100\%$); however, as could be expected, the algorithm's convergence speed decreases in the presence of large disturbances.

Finally, in Fig. 3(b), we examine the behavior of Boltzmann routing with respect to the capacity constraints imposed by the network's links and data centers. Specifically, we run (SBR) for different routing modes (cf. Fig. 2) with a 25% disturbance level and we record the number of links and/or data centers that exceed capacity at any given time. The results are plotted in the box-and-whisker charts of Fig. 3(b) which show that routing to a single destination (over a single or several paths) consistently leads to relatively high violation percentages, between 30% and 40%. On the other hand, except for a few outliers, splitting traffic over several paths and destinations yields a much lower capacity

violation percentage, of the order of 5% (a result which is statistically significant within 5 interquartile ranges). This behavior stems from the fact that the Pigouvian pricing scheme of Proposition 3 also penalizes users when a link (or data center) approaches capacity because the slope of the ε -adjusted cost function c_ε^e is equal to $1/\varepsilon$ in that case (in our simulations, we took $\varepsilon = 10^{-4}$; recall also the relevant discussion in Section II-B).

VI. CONCLUSIONS AND PERSPECTIVES

The prolific increase of traffic in optical networks is forcing network operators to seek energy-efficient routing methods while facing capacity limits and random fluctuations. This paper is an attempt to propose such a scheme: specifically, we analyzed a distributed approach for energy-efficient routing in optical data networks based on the implementation of a Pigouvian pricing scheme and a learning method derived from the Boltzmann distribution of statistical mechanics. The Nash equilibria of the proposed pricing scheme coincide with the network's socially optimum states (in terms of total energy consumption), and the resulting Boltzmann routing method converges to such states exponentially fast in realistic network conditions, leading to gains of up to 40% in energy consumption over simple, shortest-path routing schemes.

Motivated by the highly volatile nature of traffic and power consumption in optical networks, we also examine the behavior of Boltzmann routing in the presence of uncertainty and random disturbances. In this regard, we showed that the network's long-term average consumption converges within ε of its minimum value in time which is at most $\tilde{O}(\varepsilon^{-2})$, irrespective of the fluctuations' magnitude. Moreover, if the network admits a strict, non-mixing optimum state, the network's state converges to it almost surely (again, no matter the noise level); instead, if the network admits an interior, fully-mixing optimum state, the proposed scheme remains arbitrarily close to the said state with probability arbitrarily close to 1, provided that the scheme's inverse temperature parameter η is taken sufficiently small. These theoretical results were complemented by extensive numerical simulations showing that, in realistic network conditions, Boltzmann routing provides significant advantages over closest-destination/shortest-path routing choices. In addition, the proposed method scales up to hundreds of source nodes with minimal impact on its convergence speed, and it respects the network's capacity constraints, even under very high degrees of volatility and uncertainty.

The results presented in this paper apply to a wide range of congestion-limited network scenarios, ranging from latency minimization and throughput maximization in data networks to urban traffic management in road networks; the detailed analysis of these applications is relegated to future work. Additionally, a key question that arises is what happens in the presence of non-zero-mean fluctuations (random or not), which could lead the system's optimum/equilibrium states to evolve over time in an arbitrary fashion (so there is no fixed underlying state to target). In this case, efficient routing protocols would have to be capable of tracking the system's evolving optimum state in a flexible, dynamic manner; we intend to explore this issue in future work.

APPENDIX A

PRELIMINARY RESULTS

For simplicity, in what follows, we take all rates equal to $\varrho_s = 1$; this assumption is only done for notational clarity and does not otherwise affect our results. For posterity, define also the Gibbs map $y \mapsto G(y)$ with $G_a(y) =$

$\exp(y_\alpha) / \sum_{\beta \sim \alpha} \exp(y_\beta)$, so the update step of (BR) can be written as $x = G(y)$. Finally, with a fair degree of hindsight, introduce the (negative) entropy function $h(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha \log x_\alpha$, $x \in \mathcal{X}$, and its convex conjugate [34]

$$h^*(y) = \max_{x \in \mathcal{X}} \{\langle y|x \rangle - h(x)\} = \sum_{s \in \mathcal{S}} \log \sum_{\alpha \in \mathcal{A}_s} e^{y_\alpha}, \quad y \in \mathbb{R}^A. \quad (\text{A.1})$$

Then, an easy differentiation shows that $G(y) = \nabla h^*(y)$ and, furthermore, for all $\beta, \gamma \sim \alpha$:

$$\frac{\partial G_\alpha}{\partial x_\beta} = x_\alpha (\delta_{\alpha\beta} - x_\beta), \quad \frac{\partial^2 G_\alpha}{\partial x_\beta \partial x_\gamma} = x_\alpha (\delta_{\alpha\beta} - x_\beta) (\delta_{\beta\gamma} - 2x_\beta), \quad (\text{A.2})$$

where $\delta_{\alpha\beta}$ denote Kronecker's delta symbols. With these preliminaries at hand, we have:

Proof of Proposition 4: Applying Itô's formula [19, 20] to the dynamics (SBR), we get:

$$dX_\alpha = \eta \sum_{\beta \sim \alpha} \frac{\partial G_\alpha}{\partial x_\beta} \Big|_{\eta Y} dY_\beta + \dot{\eta} \sum_{\beta \sim \alpha} Y_\beta \frac{\partial G_\alpha}{\partial x_\beta} \Big|_{\eta Y} dt + \frac{\eta^2}{2} \sum_{\beta, \gamma \sim \alpha} \frac{\partial^2 G_\alpha}{\partial x_\beta \partial x_\gamma} \Big|_{\eta Y} dY_\beta \cdot dY_\gamma. \quad (\text{A.3})$$

The system (18) is then obtained by substituting dY from (SBR), using the stochastic differential multiplication formula (16) and the derivative calculations (A.2) above; this a straightforward algebraic manipulation, so we leave the details to the reader. ■

To continue, we will require a further piece of technical machinery that will act as a Lyapunov function for the dynamics (BR)/(SBR). Specifically, following [43, 48], consider the so-called *Fenchel coupling*

$$F(x, y) = h(x) + h^*(y) - \langle y|x \rangle. \quad (\text{A.4})$$

By the convexity of h , $F(x, y)$ is convex in both x and y ; furthermore, by Fenchel's inequality [34], we have $F(x, y) \geq 0$ with equality if and only if $x = \nabla h^*(y)$, i.e. if and only if $x = G(y)$. Moreover, using the explicit log-sum-exp formula (A.1), it is easy to see that

$$F(x, y) = D_{\text{KL}}(x, G(y)) \quad \text{for all } x \in \mathcal{X}, y \in \mathbb{R}^A. \quad (\text{A.5})$$

As a result, for fixed $x^* \in \mathcal{X}$, we will have $x(t) \rightarrow x^*$ if and only if $F(x^*, y(t)) \rightarrow 0$.

With this in mind, we turn to the behavior of F under (SBR):

Lemma A.1. *Let $Y(t)$ be given by (SBR). Then, for all $x^* \in \mathcal{X}$, we have:*

$$dF(x^*, Y) = \sum_{\alpha} (X_\alpha - x_\alpha^*) dY_\alpha + \frac{1}{2} \sum_{\alpha} X_\alpha \left(\sigma_{\alpha\alpha}^2 - \sum_{\beta \sim \alpha} X_\beta \sigma_{\alpha\beta}^2 \right) dt \quad (\text{A.6})$$

Proof: Simply apply Itô's lemma [19, 20] to $F(x^*, Y(t))$ and use the differentiation formula (A.2) to rewrite $\partial_\alpha \partial_\beta h^*(y) = \partial_\beta G_\alpha(y)$ in terms of x ; the details are straightforward and are left to the reader. ■

Finally, we will require the following growth estimate for martingales with bounded volatility:

Lemma A.2. *Let $W(t)$ be a Wiener process in \mathbb{R}^n and let $Z(t)$ be a bounded, continuous process in \mathbb{R}^n . Then, for every positive function f such that $f(t) / \sqrt{t \log \log t} \rightarrow \infty$ as $t \rightarrow \infty$, we have:*

$$f(t) + \int_0^t Z(s) \cdot dW(s) \sim f(t) \quad \text{as } t \rightarrow \infty \text{ (a.s.)}. \quad (\text{A.7})$$

Proof: Let $\xi(t) = \sum_{i=1}^n \int_0^t Z_i(s) dW_i(s)$. Then, the quadratic variation $\rho = [\xi, \xi]$ of ξ satisfies

$$d[\xi, \xi] = d\xi \cdot d\xi = \sum_{i=1}^n Z_i Z_j \delta_{ij} dt \leq M dt, \quad (\text{A.8})$$

where $M = \sup_{t \geq 0} \|Z(t)\|^2 < \infty$. By the time-change theorem for martingales [19, Cor. 8.5.4], there exists a Wiener process $\tilde{W}(t)$ such that $\xi(t) = \tilde{W}(\rho(t))$, so $[f(t) + \xi(t)]/f(t) = 1 + \tilde{W}(\rho(t))/f(t)$. If $\lim_{t \rightarrow \infty} \rho(t) < \infty$, we have $\limsup_{t \rightarrow \infty} |\tilde{W}(\rho(t))| < \infty$ (a.s.); otherwise, if $\lim_{t \rightarrow \infty} \rho(t) = \infty$, our claim follows from the law of the iterated logarithm [19] and the fact that $\rho(t) \leq Mt$. \blacksquare

APPENDIX B

DETERMINISTIC ANALYSIS

Proof of Theorem 1: Let $x(t)$ be an interior solution orbit of the Boltzmann routing dynamics (BR) and let $\mathcal{X}^* = \arg \min C$ denote the solution set of (CM). By Lemma A.1 and the convexity of the objective function C , it follows that the function $V(t) = F(x^*, y(t))$ is decreasing for any $x^* \in \mathcal{X}^*$, so Lyapunov's theorem ensures that $x(t) = G(y(t))$ converges to \mathcal{X}^* (note here that all stochastic terms in (A.6) vanish when $\sigma = 0$). Assume now that $x(t)$ has two distinct ω -limit points $x^*, x^{**} \in \mathcal{X}^*$ such that $x(t_n) \rightarrow x^*$ and $x(t'_n) \rightarrow x^{**}$ for two time sequences $t_n, t'_n \nearrow \infty$. It then follows that $F(x^*, y(t_n)) \rightarrow 0$ and $F(x^{**}, y(t'_n)) \rightarrow 0$; however, with both $F(x^*, y(t))$ and $F(x^{**}, y(t))$ decreasing, we obtain $\lim F(x^*, y(t)) = \lim F(x^{**}, y(t)) = 0$. This only holds if $x^* = x^{**}$, so we conclude that $x(t)$ has a unique ω -limit. \blacksquare

APPENDIX C

CONVERGENCE OF LONG-TERM AVERAGES

Proof of Theorem 2: Let x^* be a solution of (CM) and consider the η -deflated Fenchel coupling

$$H \equiv \frac{1}{\eta} F(x^*, \eta Y) = \frac{1}{\eta} \cdot [h(x^*) + h^*(\eta Y) - \langle \eta Y | x^* \rangle]. \quad (\text{C.1})$$

Then, the Itô formula of Lemma A.1 readily gives:

$$dH = -\frac{\dot{\eta}}{\eta} H dt + \frac{\dot{\eta}}{\eta} \langle Y | X - x^* \rangle dt + \langle dY | X - x^* \rangle + \frac{\eta}{2} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt. \quad (\text{C.2})$$

Hence, letting $v = -\pi$ and combining the definition of H with the dynamics (SBR), we get:

$$dH = -\frac{\dot{\eta}}{\eta^2} [h(x^*) - h(X)] dt + \langle v | X - x^* \rangle dt + \sum_{\beta} (X_{\beta} - x_{\beta}^*) \sigma_{\beta} dW_{\beta} + \frac{\eta}{2} \sum_{\beta} \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2 dt, \quad (\text{C.3})$$

where we used the fact that $h^*(\eta Y) = \langle \eta Y | X \rangle - h(X)$. Thus, after rearranging and integrating:

$$\int_0^t \langle v(s) | x^* - X(s) \rangle ds = H(0) - H(t) \quad (\text{C.4a})$$

$$- \int_0^t \frac{\dot{\eta}(s)}{\eta^2(s)} [h(x^*) - h(X(s))] ds \quad (\text{C.4b})$$

$$+ \sum_{\beta} \int_0^t (X_{\beta}(s) - x_{\beta}^*) \sigma_{\beta}(s) dW_{\beta}(s) \quad (\text{C.4c})$$

$$+ \frac{1}{2} \sum_{\beta} \int_0^t \eta(s) \frac{\partial^2 h^*}{\partial y_{\beta}^2} \sigma_{\beta}^2(s) ds. \quad (\text{C.4d})$$

We now proceed to bound each term of (C.4):

a) Since $H \geq 0$, the term (C.4a) is bounded from above as follows:

$$H(0) \leq \frac{h(x^*) + h^*(Y(0))}{\eta(0)} = \frac{h(x^*) - \min_{x \in \mathcal{X}} h(x)}{\eta(0)} + \mathcal{O}(1) = \frac{\sum_s \log A_s}{\eta(0)} + \mathcal{O}(1). \quad (\text{C.5})$$

b) For (C.4b), we have $h(x^*) - h(X(s)) \leq \sum_s \log A_s$ by definition; hence, with $\dot{\eta} \leq 0$, we get:

$$(\text{C.4b}) \leq - \sum_s \log A_s \int_0^t \frac{\dot{\eta}(s)}{\eta^2(s)} ds = \sum_s \log A_s \left(\frac{1}{\eta(t)} - \frac{1}{\eta(0)} \right). \quad (\text{C.6})$$

c) By Lemma A.2 and the law of the iterated logarithm, (C.4c) is bounded from above by $2\sigma_*^2 \sqrt{t \log \log t}$.

d) Finally, using the derivative calculations (A.2) for $\partial_\beta^2 h^*(y) = \partial_\beta G_\beta(y)$, we immediately deduce that (C.4d) is bounded from above by $\frac{1}{2}\sigma_*^2 \int_0^t \eta(s) ds$.

The bound (19) then follows by recalling that $v = -\pi = -\partial c$, integrating the convexity relation $C(X(t)) \leq C^* + \langle v(X(t)) | x^* - X(t) \rangle$, and combining the above terms. \blacksquare

APPENDIX D

CONVERGENCE TO NON-MIXING OPTIMUM STATES

To prove Theorem 3 we will need a series of auxiliary results. We begin by showing that neighborhoods of a strict, non-mixing optimum state are recurrent under (SBR):

Proposition D.1. *Let x^* be a strict, non-mixing optimum state of (CM). If η is chosen sufficiently small, there exists a (random) sequence of times $t_n \nearrow \infty$ such that $\|X(t_n) - x^*\| < \delta$ (a.s.).*

Proof: Suppose there exists some t_0 such that $\|X(t) - x^*\| \geq \delta$ for all $t \geq t_0$. Then, (C.3) yields

$$\begin{aligned} H(t) &= H(t_0) + \int_{t_0}^t \langle v(X(s)) | X(s) - x^* \rangle ds + \frac{\eta}{2} \int_{t_0}^t \frac{\partial^2 h^*}{\partial y_\beta^2} \sigma_\beta^2 ds + \xi(t) \\ &\leq H(t_0) - \mu\delta(t - t_0) + \frac{\eta\sigma^2}{2K}(t - t_0) + \xi(t) \leq H(t_0) - \left(\mu\delta - \frac{\eta\sigma^2}{2K} - \frac{\xi(t)}{t - t_0} \right) (t - t_0), \end{aligned} \quad (\text{D.1})$$

where, $v = -\pi$ as before, and we set $\xi(t) = \sum_{i=1}^d \int_{t_0}^t (X_i(s) - x_i^*) dZ_i(s)$, and $\mu = \min_{\beta \sim \text{supp}(x^*)} \{v(x^*) - v_\beta(x^*)\} > 0$ (recall that x^* is a strict equilibrium). By Lemma A.2, it follows that $\xi(t)/t \rightarrow 0$ (a.s.), so the above estimate yields $\lim_{t \rightarrow \infty} H(t) = -\infty$ if $\eta\sigma^2 < 2\mu\delta K$, a contradiction (recall that $H(t) \geq 0$ for all $t \geq 0$). This shows that $t_0 = \infty$ (a.s.) and our claim follows. \blacksquare

We now show that if $X(t)$ begins close enough to a strict equilibrium x^* , then it remains nearby and eventually converges with arbitrarily high probability:

Proposition D.2. *Let x^* be a strict, non-mixing optimum state of (CM). Then, for every $\varepsilon > 0$ and for every neighborhood U_0 of x^* in \mathcal{X} , there exists a neighborhood $U \subseteq U_0$ of x^* such that, if $X(0) \in U$, then:*

$$\mathbb{P}(X(t) \in U_0 \text{ for all } t \geq 0 \text{ and } \lim_{t \rightarrow \infty} X(t) = x^*) \geq 1 - \varepsilon. \quad (\text{D.2})$$

Proof: Write $x^* = (\alpha_1^*, \dots, \alpha_s^*)$ and let $\mathcal{A}_s^* \equiv \mathcal{A}_s \setminus \{\alpha_s^*\}$ denote the set of non-equilibrium paths of the s -th source. Moreover, for all $\alpha \in \mathcal{A}_s^*$, let

$$Z_\alpha = \eta(Y_\alpha - Y_{\alpha_s^*}), \quad (\text{D.3})$$

so that, by definition, $X(t) \rightarrow x^*$ if and only if $Z_\alpha(t) \rightarrow -\infty$ for all $\alpha \in \mathcal{A}_s^*$. With this in mind, let $M > 0$ be sufficiently large so that $X(t) \in U$ if $Z_\alpha(t) \leq -M$ for all $\alpha \in \mathcal{A}_s^*$, and assume that

$$M > \mu\eta\sigma_*^2 \log(N/\varepsilon), \quad (\text{D.4})$$

with μ defined as in the proof of Proposition D.1. We will show that if $Z_\alpha(0) \leq -2M$, then $X(t) \in U$ for all $t \geq 0$ and $Z_\alpha(t) \rightarrow -\infty$ with probability at least $1 - \varepsilon$.

To that end, let $\tau_U = \inf\{t > 0 : X(t) \notin U\}$ denote the first exit time from U ; then, for $t \leq \tau_U$, (SBR) gives

$$Z_\alpha(t) = Z_\alpha(0) + \eta \int_0^t [v_\alpha(X(s)) - v_{\alpha_s^*}(X(s))] ds + \eta\xi_s(t) \leq -2M - \eta[\mu t - \xi_s(t)], \quad (\text{D.5})$$

where we have set

$$\xi_s(t) = \int_0^t \sigma_\alpha(X) dW_\alpha - \int_0^t \sigma_{\alpha_s^*}(X) dW_{\alpha_s^*}. \quad (\text{D.6})$$

We will first show that $\mathbb{P}(\tau_U < \infty) \leq \varepsilon$. Indeed, as in the proof of Lemma A.2, there exists a rescaled Wiener process $\widetilde{W}_s(t)$ such that $\xi_s(t) = \widetilde{W}_s(\rho_s(t))$ where $\rho_s = [\xi_s, \xi_s]$ is the quadratic variation of ξ_s . We thus conclude that $Z_\alpha(t) \leq -M$ whenever $\mu t - \widetilde{W}_s(\rho_s(t)) \geq -M/\eta$. Accordingly, with $\rho_s(t) \leq 2\sigma_*^2 t$, it suffices to show that the hitting time

$$\tau_0 = \inf \left\{ t > 0 : \widetilde{W}_s(t) = \frac{\mu t}{2\sigma_*^2} + \frac{M}{\eta} \text{ for some } k \in k \right\} \quad (\text{D.7})$$

is finite with probability not exceeding ε .

Now, if a trajectory of \widetilde{W}_s has $\widetilde{W}_s(t) \leq \mu t/(2\sigma_*^2) + M/\eta$ for all $t \geq 0$, we also get

$$\widetilde{W}_s(\rho_s(t)) \leq \frac{\mu\rho_s(t)}{2\sigma_*^2} + \frac{M}{\eta} \leq \mu t + \frac{M}{\eta}, \quad (\text{D.8})$$

so τ_U is infinite whenever τ_0 is infinite. Thus, if we write E_s for the event that $\widetilde{W}_s(t) \geq \mu t/(2\sigma_*^2) + M/\eta$ for some finite $t \geq 0$, the hitting-time estimate [51, p. 197] for a Brownian motion with drift yields $\mathbb{P}(E_s) = e^{-\lambda_s M}$ with $\lambda_s = \mu/(\eta\sigma_*^2)$. Therefore, by construction of M , we obtain:

$$\mathbb{P}(\tau_0 < +\infty) = \mathbb{P}\left(\bigcup_s E_s\right) \leq \sum_s \mathbb{P}(E_s) = \sum_s e^{-\lambda_s M} \leq \varepsilon. \quad (\text{D.9})$$

By Lemma A.2, Eq. (D.5) then yields $Z_\alpha(t) \leq -2M - \eta[\mu t - \xi_s(t)] \sim -\eta\mu t \rightarrow -\infty$ (a.s.). We conclude that $Z_\alpha(t) \rightarrow -\infty$ for all $\alpha \in \mathcal{A}_s^*$, so $\lim_{t \rightarrow \infty} X(t) = x^*$, as claimed. ■

With all this at hand, we are finally in a position to prove Theorem 3:

Proof of Theorem 3: By Proposition D.1, if η is chosen sufficiently small, $X(t)$ will visit any neighborhood U of x^* infinitely often. Thus, if $\varepsilon > 0$ is fixed and U is chosen as in the statement of Proposition D.2, $X(t)$ will stay in U and converge to x^* with probability exceeding $1 - \varepsilon$. Since $X(t)$ is recurrent, the probability that $X(t)$ exits and then re-enters U for a strictly positive amount of time at least n times does not exceed ε^n . Our claim then follows by taking $n \rightarrow \infty$. ■

APPENDIX E

CONCENTRATION NEAR INTERIOR OPTIMUM STATES

Proof of Theorem 4: Assume that x^* is an isolated interior solution of (CM) and define $H(t) = \eta^{-1}F(x^*, \eta Y(t))$ as in the proof of Theorem 3. Then, taking a sufficiently small $B > 0$ such that $C(x) - C(x^*) \geq \frac{1}{2}B\|x - x^*\|^2$ (recall that x^* is an isolated, interior minimum point), (C.4) yields

$$H(t) \leq H(0) - \frac{1}{2}B \int_0^t \|X(s) - x^*\|^2 ds + \frac{\eta}{2}\sigma_*^2 t + \xi(t), \quad (\text{E.1})$$

where the martingale term $\xi(t)$ is defined as in the proof of Proposition D.1. Thus, dividing by t and taking expectations yields the bound

$$\mathbb{E}\left[\frac{1}{t} \int_0^t \|X(s) - x^*\|^2 ds\right] \leq \frac{\eta\sigma_*^2}{B} + \frac{H(0)}{t}. \quad (\text{E.2})$$

Our goal will be to show that $X(t)$ admits an invariant measure which can be estimated from the above bound to yield (22).

To that end, consider the process $\Psi(X(t))$ where $\Psi_{s\beta}(x) = \log(x_{s\beta}/x_{s\alpha}) = Y_{s\beta} - Y_{s\alpha}$ for some fixed $\alpha \in \mathcal{A}_s, \beta \sim \alpha, s \in \mathcal{S}$. Itô's formula then yields:

$$d\Psi_\beta = dY_\beta - dY_\alpha = \mathcal{L}\Psi_\beta(X) dt + \sum_e \Pi'_{e\beta} \sigma_e dW_e, \quad (\text{E.3})$$

where $\mathcal{L}\Psi$ denotes the deterministic part of (D.5) and $\Pi'_{e\beta} = \Pi_{e\beta} - \Pi_{e\alpha}$ (the source index s has been dropped for convenience). Following [40], we claim that the infinitesimal generator of Ψ is (uniformly) elliptic [19, 20]. For this, let $A_{\beta e} = \Pi'_{e\beta} \sigma_e, \beta \in \coprod_s \mathcal{A}_s^*$ denote the coefficient matrix of the martingale term of (E.3); then, it suffices to show that the matrix AA^T is positive-definite on \mathcal{X} . Indeed, for all $z \in \text{TC}_{\mathcal{X}}(x^*)$, we have:

$$\langle Az, Az \rangle = \sum_{\beta, \gamma} \sum_e \Pi'_{e\beta} \Pi'_{e\gamma} \sigma_e^2 z_\beta z_\gamma = \sum_e \sigma_e^2 (\Pi' z)_e^2 > 0, \quad (\text{E.4})$$

because Π' is necessarily invertible on \mathcal{X} (otherwise x^* would not be isolated). This proves our ellipticity assertion, so [52, Lemma 3.4] shows that Ψ is recurrent.

Since Ψ is recurrent, it follows that $X(t)$ is also recurrent; consequently, the transition probabilities of $X(t)$ converge in total variation to an invariant probability measure ν on \mathcal{X} [20, 53]. Thus, if B_δ is a δ -ball centered at x^* , ergodicity gives $\nu(B_\delta) = \lim_{t \rightarrow \infty} \mathbb{E}_x \left[\frac{1}{t} \int_0^t \chi_{B_\delta}(X(s)) ds \right]$, where χ_{B_δ} is the indicator function of B_δ .

Now, with $\|x - x^*\| \geq \delta$ outside B_δ , we readily obtain

$$\mathbb{E}\left[\frac{1}{t} \int_0^t \chi_{B_\delta}(X(s)) ds\right] \geq \mathbb{E}\left[\frac{1}{t} \int_0^t (1 - \|X(s) - x^*\|^2/\delta^2) ds\right]. \quad (\text{E.5})$$

Thus, letting $t \rightarrow \infty$ in (E.2), we get $\nu(B_\delta) \geq 1 - \eta\sigma_*^2/(B\delta^2)$. Taking η sufficiently small, this bound can be assumed greater than or equal to $1 - \varepsilon$; our assertion then follows from the convergence in total variation of the transition probabilities of $X(t)$ to the invariant measure ν . ■

REFERENCES

- [1] P. Mertikopoulos and A. L. Moustakas, “Selfish routing revisited: Degeneracy, evolution and stochastic fluctuations,” in *ValueTools '11: Proceedings of the 5th International Conference on Performance Evaluation Methodologies and Tools*, 2011.
- [2] ———, “A learning approach to efficient routing and the effect of stochastic fluctuations,” <http://arxiv.org/abs/0912.4012>.
- [3] R. W. Tkach, “Scaling optical communications for the next decade and beyond,” *Bell Labs Technical Journal*, vol. 14, no. 4, pp. 3–9, 2010.
- [4] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermuelen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, “Worldwide energy needs for ICT: The rise of power-aware networking,” in *ANTS '08: Proceedings of the 2008 IEEE International Conference on Advanced Networks and Telecommunication Systems*, 2008.
- [5] A. Plepys, “The grey side of ICT,” *Journal of Environmental Impact Assessment Review*, vol. 22, no. 5, pp. 509–523, 2002.
- [6] A. Vukovic, “Network power density challenges,” *ASHRAE journal*, vol. 47, no. 4, p. 55, 2005.
- [7] G. Fettweis and E. Zimmermann, “ICT energy consumption-trends and challenges,” in *Proceedings of the 11th International Symposium on Wireless Personal Multimedia Communications*, vol. 2, no. 4, 2008, p. 6.
- [8] W. H. Sandholm, “Pigouvian pricing and stochastic evolutionary implementation,” *Journal of Economic Theory*, vol. 132, pp. 367–382, 2007.
- [9] V. S. Borkar and P. R. Kumar, “Dynamic Cesaro-Wardrop equilibration in networks,” *IEEE Trans. Autom. Control*, vol. 48, no. 3, pp. 382–396, March 2003.
- [10] A. Blum, E. Even-Dar, and K. Ligett, “Routing without regret: on convergence to Nash equilibria of regret-minimizing in routing games,” in *PODC '06: Proceedings of the 25th annual ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing*, 2006.
- [11] W. Krichene, B. Drighès, and A. Bayen, “On the convergence of no-regret learning in selfish routing,” in *ICML '14: Proceedings of the 31st International Conference on Machine Learning*, 2014.
- [12] S. Fischer and B. Vöcking, “On the evolution of selfish routing,” in *Proceedings of the 12th European Symposium on Algorithms*, 2004.
- [13] P. Mertikopoulos and A. L. Moustakas, “The emergence of rational behavior in the presence of stochastic perturbations,” *The Annals of Applied Probability*, vol. 20, no. 4, pp. 1359–1388, 2010.
- [14] P. Mertikopoulos, E. V. Belmega, A. L. Moustakas, and S. Lasaulce, “Distributed learning policies for power allocation in multiple access channels,” *IEEE J. Select. Areas Commun.*, vol. 30, no. 1, pp. 96–106, January 2012.
- [15] S. D’Oro, P. Mertikopoulos, A. L. Moustakas, and S. Palazzo, “Interference-based pricing for opportunistic multi-carrier cognitive radio systems,” *IEEE Trans. Wireless Commun.*, vol. 14, no. 12, pp. 6536–6549, December 2015.
- [16] J. Zhou and X. Yuan, “A study of dynamic routing and wavelength assignment with imprecise network state information,” in *Parallel Processing Workshops, 2002. Proceedings. International Conference on*. IEEE, 2002, pp. 207–213.
- [17] X. Masip-Bruin, S. Sanchez-Lopez, and D. Colle, “Routing and wavelength assignment under inaccurate routing information in networks with sparse and limited wavelength conversion,” in *Global Telecommunications Conference, 2003. GLOBECOM'03. IEEE*, vol. 5. IEEE, 2003, pp. 2575–2579.
- [18] M. Yannuzzi, M. Quagliotti, G. Maier, E. Marín-Tordera, X. Masip-Bruin, S. Sánchez-López, J. Sole-Pareta, W. Erangoli, and G. Tamiri, “Performance of translucent optical networks under dynamic traffic and uncertain physical-layer information,” in *Optical Network Design and Modeling, 2009. ONDM 2009. International Conference on*. IEEE, 2009, pp. 1–6.
- [19] B. Øksendal, *Stochastic Differential Equations*, 6th ed. Berlin: Springer-Verlag, 2007.
- [20] H.-H. Kuo, *Introduction to Stochastic Integration*. Berlin: Springer, 2006.
- [21] D. P. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 1992.
- [22] R. W. Rosenthal, “A class of games possessing pure-strategy Nash equilibria,” *International Journal of Game Theory*, vol. 2, pp. 65–67, 1973.
- [23] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. Cambridge, MA, USA: MIT Press, 2005.
- [24] W. H. Sandholm, *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press, 2010.
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [26] M. Beckmann, C. B. McGuire, and C. Winsten, *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [27] X. Fan, W.-D. Weber, and L. A. Barroso, “Power provisioning for a warehouse-sized computer,” *SIGARCH Computer Architecture News*,

- vol. 35, no. 2, pp. 13–23, 2007.
- [28] J. Buysse, K. Georgakilas, A. Tzanakaki, M. de Leenheer, B. Dhoedt, and C. Develder, “Energy-efficient resource provisioning algorithms for optical clouds,” *IEEE J. Opt. Commun. Netw.*, vol. 5, no. 3, pp. 226–239, 2013.
- [29] M. P. Anastasopoulos, A. Tzanakaki, and K. Georgakilas, “Evolutionary optimization for energy efficient service provisioning in it and optical network infrastructures,” *Optics express*, vol. 19, no. 26, pp. B496–B502, 2011.
- [30] C. D. Aliprantis and K. C. Border, *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, 2nd ed. Berlin: Springer, 1999.
- [31] A. C. Pigou, *The Economics of Welfare*. London, UK: Macmillan, 1920.
- [32] J. G. Wardrop, “Some theoretical aspects of road traffic research,” in *Proceedings of the Institute of Civil Engineers, Part II*, vol. 1, 1952, pp. 325–78.
- [33] S. C. Dafermos and F. T. Sparrow, “The traffic assignment problem for a general network,” *Journal of Research of the National Bureau of Standards*, vol. 73B, no. 2, pp. 91–118, 1969.
- [34] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton University Press, 1970.
- [35] P. D. Taylor and L. B. Jonker, “Evolutionary stable strategies and game dynamics,” *Mathematical Biosciences*, vol. 40, no. 1-2, pp. 145–156, 1978.
- [36] L. E. Blume, “The statistical mechanics of strategic interaction,” *Games and Economic Behavior*, vol. 5, pp. 387–424, 1993.
- [37] R. Durrett, *Stochastic Calculus*. Boca Raton, FL: CRC Press, 1996.
- [38] D. Fudenberg and C. Harris, “Evolutionary dynamics with aggregate shocks,” *Journal of Economic Theory*, vol. 57, no. 2, pp. 420–441, August 1992.
- [39] A. Cabrales, “Stochastic replicator dynamics,” *International Economic Review*, vol. 41, no. 2, pp. 451–81, May 2000.
- [40] L. A. Imhof, “The long-run behavior of the stochastic replicator dynamics,” *The Annals of Applied Probability*, vol. 15, no. 1B, pp. 1019–1045, 2005.
- [41] J. Hofbauer and L. A. Imhof, “Time averages, recurrence and transience in the stochastic replicator dynamics,” *The Annals of Applied Probability*, vol. 19, no. 4, pp. 1347–1368, 2009.
- [42] P. Mertikopoulos and Y. Viosat, “Imitation dynamics with payoff shocks,” *International Journal of Game Theory*, vol. 45, no. 1-2, pp. 291–320, March 2016.
- [43] M. Bravo and P. Mertikopoulos, “On the robustness of learning in games with stochastically perturbed payoff observations,” <http://arxiv.org/abs/1412.6565>, 2016.
- [44] W. H. Sandholm, “Potential games with continuous player sets,” *Journal of Economic Theory*, vol. 97, pp. 81–108, 2001.
- [45] J. Hofbauer and W. H. Sandholm, “Stable games and their dynamics,” *Journal of Economic Theory*, vol. 144, pp. 1710–1725, 2009.
- [46] S. Kullback and R. A. Leibler, “On information and sufficiency,” *Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [47] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [48] P. Mertikopoulos and W. H. Sandholm, “Learning in games via reinforcement and regularization,” *Mathematics of Operations Research*, 2016, to appear.
- [49] R. D. Doverspike, K. K. Ramakrishnan, and C. Chase, “Structural overview of ISP networks,” in *Guide to Reliable Internet Services and Applications*, R. C. Kalmanek, S. Misra, and Y. Yang, Eds. Springer London, 2010.
- [50] J. E. Pultz, “Data center redesign yields an 80%-plus reduction in energy usage,” Gartner, Inc., Tech. Rep. G00213049, August 2011.
- [51] I. Karatzas and S. E. Shreve, *Brownian Motion and Stochastic Calculus*. Berlin: Springer-Verlag, 1998.
- [52] R. N. Bhattacharya, “Criteria for recurrence and existence of invariant measures for multidimensional diffusions,” *The Annals of Probability*, vol. 6, pp. 541–553, 1978.
- [53] R. Z. Khasminskii, *Stochastic Stability of Differential Equations*, 2nd ed., ser. Stochastic Modelling and Applied Probability. Berlin: Springer-Verlag, 2012, no. 66.