

Maximizing Queueing Network Utility subject to Stability: Greedy Primal-Dual Algorithm

Alexander L. Stolyar
Bell Labs, Lucent Technologies
600 Mountain Ave., 2C-322
Murray Hill, NJ 07974
stolyar@research.bell-labs.com

Submitted August 4, 2004; revised February 4, 2005

Abstract

We study a model of controlled queueing network, which operates and makes control decisions in discrete time. An underlying random network mode determines the set of available controls in each time slot. Each control decision “produces” a certain vector of “commodities”; it also has associated “traditional” queueing control effect, i.e., it determines traffic (customer) arrival rates, service rates at the nodes, and random routing of processed customers among the nodes. The problem is to find a dynamic control strategy which maximizes a concave utility function $H(X)$, where X is the average value of commodity vector, subject to the constraint that network queues remain stable.

We introduce a dynamic control algorithm, which we call *Greedy Primal-Dual* (GPD) algorithm, and prove its asymptotic optimality. We show that our network model and GPD algorithm accommodate a wide range of applications. As one example, we consider the problem of congestion control of networks where both traffic sources and network processing nodes may be randomly time-varying and interdependent. We also discuss a variety of resource allocation problems in wireless networks, which in particular involve average power consumption constraints and/or optimization, as well as traffic rate constraints.

Key words and phrases: Queueing networks, Convex optimization, Primal-Dual algorithm, Stability, Congestion control, Resource Allocation, Scheduling, Wireless, Power and rate constraints

Abbreviated Title: Greedy Primal-Dual Algorithm

1 Introduction

1.1 The problem

We study a model of controlled queueing network, which operates and makes control decisions in discrete time $t = 0, 1, 2, \dots$. The key feature of the model is that each network control action has two effects. First, network has a finite set \mathcal{N}^p of “traffic” processing nodes, with queues, and each control action has associated “queueing control” which affects traffic (customer) arrival rates to processing nodes, their processing (service) rates, routing between processing nodes, etc. Second, network generates a finite number of “commodity (utility) flows,” forming set \mathcal{N}^u ; namely, each control action k generates amounts $b_n(k)$, $n \in \mathcal{N}^u$, of the commodities. In addition, the available set of control options depends on some underlying random network “mode,” modelled by an ergodic Markov chain.

Let $X = (X_n, n \in \mathcal{N}^u)$, be the vector of long-term average rates at which commodities are generated under a given control strategy. We seek to find a dynamic control strategy which maximizes a concave utility function $H(X)$ of average commodity rates, subject to the constraint that the network remains stable, that is, roughly speaking, queues at the processing nodes remain bounded. (The utility function H need *not* be strictly concave.) In this paper we introduce a dynamic control algorithm, called *Greedy Primal-Dual* (GPD) algorithm, which solves the above problem (asymptotically, as described shortly), under the natural assumption that, roughly speaking, it is feasible at all to keep the network stable (even if we ignore utility optimization). As we will see, the algorithm is very parsimonious, and naturally decomposes to become a decentralized algorithm in special cases when different network elements can make “their own” control decisions independently.

Since both commodity generation and queueing control actions depend on a chosen network control action, our model accommodates, in general, scenarios in which “currently available” choices of commodity generation rates, traffic arrival and service rates, and routing are mutually interdependent in arbitrary way. As we will see, this feature is very useful in modelling many systems arising in applications.

In applications, different commodity types may have different meanings, and some of the commodities may be “physical” and some “virtual.” In telecommunication systems a commodity may be a traffic flow, which may (or may not) need to go through and be processed by the processing network (this is modelled by “coupling” the generated commodity amounts and amounts of arrived traffic to some nodes). A commodity may also correspond to a monetary award (or penalty), associated with a control action. Or, going back to telecommunications, and in particular wireless systems, a commodity may be energy or power consumed by a control action. Thus, a commodity may be virtual in the sense that it serves simply to keep track of and optimize certain performance measures. For example, the GPD algorithm can be used to control a queueing network at the lowest average cost (or power consumption), while keeping queues stable. Some examples of applications are discussed in Section 5. We will demonstrate there that the GPD algorithm can create virtual processing

nodes as well; for example, to solve the stated optimization problem, subject to additional desired constraints on the average commodity generation rates. To summarize, our abstract view that network control actions have double effect of controlling queues on one hand, and generating some commodities on the other hand, allows the model to accommodate a large variety of applications and scenarios.

1.2 Informal description of the GPD algorithm

Formal definitions of the model and the GPD algorithm are given in Section 4. Here we describe the algorithm informally, to facilitate discussion.

Recall that $b_n(k)$, $n \in \mathcal{N}^u$, are the commodity amounts generated by control k . Denote by $\overline{\Delta Q_n}(k)$ the expected average (possibly negative) increment, or drift, of the queue length Q_n at a processing node $n \in \mathcal{N}^p$, caused by control k , assuming that all Q_n at all processing nodes were large enough for queues to not empty as a result of the control. (Drift values $\overline{\Delta Q_n}(k)$ incorporate the aggregate effect of new arrivals, service, and inter-node routing associated with control k .) Assume that the utility function H is continuously differentiable. The GPD algorithm is such that, from all the control actions k available at time t (with the set of available k depending on the underlying random network mode), it chooses a control

$$k(t) \in \arg \max_k \sum_{n \in \mathcal{N}^u} (\partial H(X(t))/\partial X_n) b_n(k) - \sum_{n \in \mathcal{N}^p} \beta Q_n(t) \overline{\Delta Q_n}(k), \quad (1)$$

where $\beta > 0$ is a (typically small) parameter, and $X(t)$ is the vector of the “current averages” of commodity rates, updated as follows:

$$X_n(t+1) = X_n(t) + \beta(b_n(k(t)) - X_n(t)), \quad n \in \mathcal{N}^u. \quad (2)$$

Using natural vector notations, and “ \cdot ” to denote scalar products, the GPD rule (1) can be concisely written as follows:

$$k(t) \in \arg \max_k \nabla H(X(t)) \cdot b(k) - \beta Q(t) \cdot \overline{\Delta Q}(k). \quad (3)$$

We prove asymptotic optimality of the GPD algorithm in the following sense:

(I) When parameter β is small, the rescaled trajectories $(X(t/\beta), \beta Q(t/\beta))$, $t \geq 0$, under GPD algorithm, are close to trajectories $(x(t), q(t))$, $t \geq 0$, of some dynamic system. (We call the latter *GPD-trajectories*.)

(II) If network structure and parameters are such that it is feasible at all to keep queues stable, then, as $t \rightarrow \infty$, all GPD-trajectories converge to points (v^*, q^*) such that $H(v^*)$ is the maximum possible value of network utility under the stability constraint.

Analysis of the dynamic system (in Section 3), including proof of the global uniform attraction property of GPD-trajectories (Theorem 2), is in fact the key part of this paper. The dynamic system is “primal-dual” - the attraction points (v^*, q^*) are optimal solutions to

some underlying convex optimization problem and its dual. To the best of our knowledge, the dynamic system is *not* of any known primal-dual type. (Previously studied primal-dual dynamic systems include, for example, the classical Arrow-Hurwicz-Uzawa system (cf. [13]) and systems within *passivity* framework of [26].)

1.3 Discussion

It is easy to observe that, if we remove all processing nodes from our model (so that there is no stability constraint), the GPD algorithm reduces to the Gradient algorithm [3, 21], always “greedily” choosing controls maximizing the drift of utility (“Lyapunov”) function $H(X(t))$, and shown to be asymptotically optimal as $\beta \downarrow 0$ in [3, 21]. If we remove all commodity flows (so that we are not concerned about any utility), the GPD algorithm becomes a “MaxWeight”-type algorithm (cf. [22, 23, 20, 4]), “greedily” pursuing minimization of Lyapunov function $\sum_{\mathcal{N}^p} (1/2)Q_n^2(t)$, and known to ensure stability of queueing networks if such is feasible at all. (Cf. [4] for recent results and review of previous work.)

The GPD rule (3) shows that, roughly speaking, the algorithm always “greedily” chooses controls maximizing the expected drift of function

$$H(X(t)) - \sum_{n \in \mathcal{N}^p} (1/2)\beta Q_n^2(t) ,$$

and thus can be viewed as a “naive” combination of the Gradient and MaxWeight algorithms. The fact that such naive combination produces an algorithm which is (asymptotically) optimal for our problem, is non-trivial and perhaps even surprising for the following reasons. For a MaxWeight algorithm, function $\sum (1/2)Q_n^2(t)$ *is* in fact a Lyapunov function which ensures stability. For the Gradient algorithm, $H(X(t))$ is “almost” a Lyapunov function - it becomes such after $X(t)$ gets close to a certain region (and convergence to the region is not very hard to establish [21]). In contrast, it is easy to see that for the GPD algorithm function $H(X(t)) - \sum (1/2)\beta Q_n^2(t)$ *cannot* serve as a Lyapunov function ensuring convergence to an optimal point. (See discussion at the beginning of Section 3.5.)

The proof of convergence is further complicated by the fact that we do not require utility function $H(X)$ to be strictly concave. As a result, optimal points (attractors) v^* may be non-unique. We avoid strict concavity assumption because it is too restrictive in many applications of interest. Utility functions which are linear in some of the X_n are an obvious example.

Papers [20] and [21], which, in particular, study convergence properties of dynamic systems arising as asymptotic limits under “pure MaxWeight” and “pure Gradient” policies, serve as starting points of our analysis.

Remark 1. *The role of parameter β .* Note that, the scaling $(X(t/\beta), \beta Q(t/\beta))$, which in the limit gives rise to GPD-trajectories $(x(t), q(t))$, is such that the space scaling by factor β is applied only to queue lengths Q , but not to average rates X . Given (2), this ensures that

both $x(t)$ and $q(t)$ change on the same scale in space. In the special case when there is no utility nodes, the first term under the arg max in (3) vanishes, making parameter β irrelevant, and leading (as already mentioned above) to MaxWeight algorithm, which is known to be optimal for the purposes of network stability (and not using parameter β). However, in the GPD algorithm (for the purposes of utility maximization subject to stability), we cannot avoid using small parameter β in (1) and (2), roughly for the following reason. For *any* algorithm to be “close to optimal,” a rule for choosing controls $k(t)$ *must* be “close” to the rule

$$k(t) \in \arg \max_k \nabla H(v^*) \cdot b(k) - q^* \cdot \overline{\Delta Q}(k),$$

where v^* and q^* are optimal solutions to the underlying convex optimization problem and its dual. In order for $\beta Q(t)$ to stay “close” to q^* (and $X(t)$ to be close to v^*), parameter β in (3) must be small. Then, $Q(t)$ will be large (of the order of $1/\beta$), and the increments of $\beta Q(t)$ small. Parameter β also determines the time scale on which $X(t)$ and $\beta Q(t)$ converge “close to” optimal values - the convergence time is of the order $1/\beta$. Thus, the smaller the β , the more “precise” the GPD algorithm is, but at the cost of larger queue lengths and larger convergence times.

Remark 2. *More general form of GPD algorithm, with “weighted queue lengths.”* All our results still hold if the queue lengths $Q_n(t)$ in (3) are weighted by arbitrary constants $\gamma_n > 0$. In other words, optimality of the GPD algorithm still holds if rule (3) is replaced with the following more general one:

$$k(t) \in \arg \max_k \nabla H(X(t)) \cdot b(k) - \sum_{n \in \mathcal{N}^p} \beta \gamma_n Q_n(t) \overline{\Delta Q}_n(k).$$

In this case the GPD-trajectories $(x(t), q(t))$ are such that $x(t)$ still converges to optimal points v^* , but it is the weighted vector $(\gamma_n q_n(t), n \in \mathcal{N}^p)$, not $q(t)$ itself, that converges to an optimal point q^* . (See Section 3.8.2.) This is a useful property for applications. As discussed in Remark 1 above, when system is close to optimal operating point, the (unscaled) queue lengths $Q_n(t)$ are “large,” of the order $1/\beta$. Introducing a weight $\gamma_n > 1$ for queue n will “scale down” a typical size of $Q_n(t)$ by factor γ_n . In other words, weights γ_n provide “knobs” allowing one to control queue length magnitudes. (But, of course, using GPD algorithm with all large weights γ_n will reduce its “precision,” due to reasons discussed in Remark 1.)

1.4 Applications

One application of our network model and the GPD algorithm, which we discuss in some detail in Sections 5.1-5.2, is to dynamic congestion control of communication networks. Recent advances in the theory of network congestion control originate from the discovery [10] that the congestion control mechanism employed by the Transport Control Protocol (TCP) in the Internet implicitly attempts to maximize the sum of traffic flow (concave) utilities, subject to the constraint that none of the network communication links are overloaded. (Note that the latter condition is necessary for stability of such networks.) This is a very large and

active field, which we do not attempt to review here. (See [11, 16, 14, 18] for recent reviews of the field.) We will only point out some aspects of congestion control modelling, to which we believe our results contribute, as illustrated by the application examples in Sections 5.1-5.2:

- We allow traffic sources to be randomly time-varying, interdependent, with possibly severely limited choices of instantaneously available transmission rates. This is the case when, for example, several network users inject traffic into a network via a shared wireless link. (Cf. [8, 1, 15] for basic features of transmission rate allocation over shared wireless links.)
- Similarly, we allow traffic processing nodes in the network to be randomly time-varying, interdependent, with possibly limited choices of instantaneously available service rates. (This aspect is also addressed in [6], which is an independent and contemporaneous work with present paper. See remark in Section 5.1.)
- Our models include routing and corresponding dynamics of the network queues. The latter means that (in our models) a data packet of a given traffic flow is forwarded to the next node on the flow route only after it is served in the previous node. (This is in contrast to assuming, explicitly or implicitly, that “an independent copy” of a packet is sent directly from a traffic source to each of the nodes on its route, and each copy leaves the network after being served by the corresponding single node.)
- Our models accommodate scenarios where both traffic sources and network processing nodes operate under average “power” (or other “cost”) consumption constraints.

Emphasizing the last point above, we demonstrate by several examples in Section 5.3 that GPD algorithm provides a solution to a wide class of resource allocation problems in wireless networks, which include power usage and traffic rate constraints. (Cf. [24, 27, 12] for some problems with power constraints.) We show that GPD algorithm easily and naturally accommodates different system models and objectives by using, if necessary, virtual commodities and/or virtual queues. The algorithms are dynamic and parsimonious, i.e., not requiring a priori system information - using only current system state and current control choices. As a result, GPD algorithm provides a unified (and provably optimal) framework for solving problems such as:

- “Optimize utility of the network traffic flows subject to stability and average power usage constraints,”
- “Keep network stable subject to average power usage constraints” (this is a special case of the previous problem),
- “Minimize average power consumption of the network, subject to stability.”

1.5 Layout of the rest of the paper

In Section 2 we introduce basic notations and conventions used in the paper. Analysis of the dynamic system, including key attraction property of GPD-trajectories (Theorem 2), is in Section 3. Our queueing network model and formal definition of GPD algorithm are in Section 4; this section also contains a formal proof that rescaled trajectories

$(X(t/\beta), \beta Q(t/\beta))$, $t \geq 0$, indeed converge to GPD-trajectories, as $\beta \rightarrow 0$. Section 5 discusses applications of GPD algorithm in some detail. (Section 3 and Sections 4-5 are virtually independent, besides the fact that Sections 4-5 refer to Theorem 2. However, we recommend that even if a reader is interested primarily in the network model and applications, he/she at least skims through Sections 3.1-3.3 before reading Sections 4-5.) Appendix contains some technical and auxiliary results.

2 Basic Notation

We denote by R , R_+ , R_- , and R_{--} , the sets of real, real non-negative, real non-positive, and real strictly negative numbers, respectively. Corresponding N -times product spaces are denoted R^N , R_+^N , R_-^N , and R_{--}^N . The space R^N is viewed as a standard vector-space, with elements $x \in R^N$ being row-vectors $x = (x_1, \dots, x_N)$. The scalar product of $x, y \in R^N$ is

$$x \cdot y \doteq \sum_{n=1}^N x_n y_n ;$$

and the norm of x is

$$\|x\| \doteq \sqrt{x \cdot x} .$$

We denote by

$$\rho(x, y) \doteq \|x - y\|$$

the *distance* between vectors x and y in R^N , and by

$$\rho(x, V) \doteq \inf_{y \in V} \rho(x, y)$$

the distance between a vector $x \in R^N$ and a set $V \subseteq R^N$. If $(x(t), t \geq 0)$ and V is a vector function and a set, respectively, in R^N , the convergence $x(t) \rightarrow V$ means that $\rho(x(t), V) \rightarrow 0$ as $t \rightarrow \infty$.

The angle between two non-zero vectors $x, y \in R^N$ is defined in the usual way as

$$\arccos \frac{x \cdot y}{\|x\| \|y\|} .$$

For a set V and a scalar function $W(v)$, $v \in V$,

$$\arg \max_{v \in V} W(v)$$

denotes the *subset* of vectors $v \in V$ which maximize $W(v)$. In all cases in this paper, V will be compact (in appropriate topology) and $W(v)$ continuous, so that the arg max will be non-empty.

For $\xi, \eta \in R$, we denote $\xi \wedge \eta \doteq \min\{\xi, \eta\}$, $\xi \vee \eta \doteq \max\{\xi, \eta\}$, $\xi^+ \doteq \max\{\xi, 0\}$; for $\xi \in R$ and $\eta \in R_+$, $[\xi]_\eta^+ \doteq \xi$ if $\eta > 0$, and $[\xi]_\eta^+ \doteq \xi^+$ if $\eta = 0$.

For a function $h(t)$ of a real variable t , $(d^+/dt)h(t)$ denotes the right derivative.

Abbreviation *u.o.c.* means *uniform on compact sets* convergence of functions. The term *almost everywhere (a.e.)* means almost everywhere with respect to Lebesgue measure. Abbreviation *w.p.1* means *with probability 1*. Symbol \xrightarrow{w} signifies *weak convergence of random processes*.

We denote by $D_{R^N}[0, \infty)$ the Skorohod space of functions with domain $[0, \infty)$, taking values in R^N , $N \geq 1$, which are right-continuous and have left-limits. The subspace of $D_{R^N}[0, \infty)$ consisting of continuous functions is denoted by $C_{R^N}[0, \infty)$; notation $C_{R^N_+}[0, \infty)$ is used for the subset of $C_{R^N}[0, \infty)$ consisting of functions with values in R^N_+ . (Topologies, σ -algebras, and norms on these spaces are specified later, where and when necessary.)

3 Greedy Primal-Dual Dynamic System: Global Attraction Property

In this section we formally define and study convergence properties of a dynamic system, which, in particular, arises as an asymptotic limit describing evolution of the queueing network model of Section 4. We believe that the dynamic system is of independent interest as well, as a “tool” for solving certain convex optimization problems. In fact, the main convergence result, Theorem 2, is somewhat more general than a result “needed” in Section 4. (See Remark 1 following Theorem 2 in Section 3.3.) Theorem 2 allows some further generalizations as well (see Section 3.8).

As mentioned earlier, this section is independent of Sections 4-5, besides the fact that those sections refer to Theorem 2. (In particular, reader should not look for direct connection between the model and notations we used in the Introduction, and those in this section. A close connection in fact exists, which will become clear in Section 4.) In some places, however, we refer to Section 4 for “physical interpretations” of some material in this section.

3.1 Optimization Problem

Consider a convex compact subset $V \subset R^N$ of a finite-dimensional space R^N , $N \geq 1$. (We will use notation $\mathcal{N} \doteq \{1, \dots, N\}$ for the set of indices of vectors $\xi = (\xi_1, \dots, \xi_N) \in R^N$.) Assume that $V \subset \tilde{V}$, where $\tilde{V} \subseteq R^N$ is open and convex, and we have a concave continuously differentiable (“utility”) function $H(v)$, $v \in \tilde{V}$.

Consider the following optimization problem:

$$\max_{v \in V} H(v) \tag{4}$$

subject to

$$v \in R_-^N. \quad (5)$$

(For a “queueing interpretation” of the problem (4)-(5), as well as that of the non-degeneracy condition (8) below, see Remark 2 in Section 3.2. For a more specific - “physical” - interpretation of (4)-(5) and (8), see the problem (52)-(53) and condition (55) for the queueing network control model of Section 4.)

The problem (4)-(5) can be equivalently written as

$$\max_{v \in V^{cond}} H(v), \quad \text{where } V^{cond} \doteq V \cap R_-^N.$$

Optimization problem (4)-(5) is feasible when

$$V \cap R_-^N = V^{cond} \neq \emptyset, \quad (6)$$

in which case we denote by $V^* \subseteq V$ the convex compact subset of its optimal solutions. (If H is *strictly* concave, the optimal solution is unique.) Also, we denote by Q^* the convex closed set of optimal solutions $q^* \in R_+^N$ to the following convex optimization problem, dual to the problem (4)-(5):

$$\min_{y \in R_+^N} [\max_{v \in V} (H(v) - y \cdot v)] \quad (7)$$

Next, in Section 3.2, we define a dynamic system, which, as we will show, “solves” problem (4)-(5) under the following non-degeneracy assumption, which is slightly stronger than (6):

$$V \cap R_{--}^N \neq \emptyset. \quad (8)$$

The dynamic system “solves” (4)-(5) in the sense that (assuming (8)) its trajectories converge to the (saddle) set $V^* \times Q^*$.

Note that, under assumption (8), set Q^* is compact. Indeed, the optimal value of the problem (4)-(5) is

$$H(v^*) = \max_{v \in V} (H(v) - q^* \cdot v) \quad (9)$$

for any $v^* \in V^*$ and any $q^* \in Q^*$. Set Q^* must be bounded, because otherwise, given condition (8), we could make RHS of (9) arbitrarily large by choosing $q^* \in Q^*$ with large norm $\|q^*\|$.

3.2 Dynamic System Definition

We define a *trajectory of the greedy primal-dual dynamic system*, or *GPD-trajectory*, as a triple of absolutely continuous functions $(x, q, f) = ((x(t), t \geq 0), (q(t), t \geq 0), (f(t), t \geq 0))$.

0)), each taking values in R^N , satisfying the following conditions:

(i) For all $t \geq 0$,

$$x(t) \in \tilde{V}, \quad (10)$$

and for almost all $t \geq 0$,

$$x'(t) = v(t) - x(t), \quad (11)$$

where

$$v(t) = (d/dt)f(t) \quad (12)$$

satisfies

$$v(t) \in \arg \max_{v \in V} (\nabla H(x(t)) - q(t)) \cdot v. \quad (13)$$

(ii) We have

$$q(0) \geq 0, \quad (14)$$

$$f_n(0) = q_n(0) \text{ and } q_n(t) = f_n(t) - [0 \wedge \inf_{0 \leq \xi \leq t} f_n(\xi)], \quad t \geq 0, \quad n \in \mathcal{N}. \quad (15)$$

The following lemma shows that a GPD-trajectory (x, q, f) can be equivalently, and perhaps more intuitively, defined via a pair (x, q) only. (We use the above definition because it is more convenient for the analysis. Sometimes, we call a pair (x, q) itself a GPD-trajectory, if there exists the corresponding GPD-trajectory (x, q, f) .)

Lemma 1 *If (x, q, f) is a GPD-trajectory, then the pair (x, q) satisfies the set of conditions (10), (11), (13), (14), and the following condition*

$$q_n(t) \geq 0, \forall t \geq 0, \text{ and } q_n'(t) = [v_n(t)]_{q_n(t)}^+ \text{ a.e. in } t \geq 0, \quad n \in \mathcal{N}. \quad (16)$$

Conversely, if a pair of absolutely continuous functions (x, q) satisfies conditions (10), (11), (13), (14) and (16), then there exists unique f such that (x, q, f) is a GPD-trajectory.

Proof of Lemma 1 is in Section 6.1 in Appendix.

Remark 1. *Interpretation in terms of optimization problem (4)-(5).* Functions $x(t)$ and $q(t)$ of a GPD-trajectory are “dynamically changing” primal and dual variables, respectively, for the problem (4)-(5). The term “primal-dual” in the name of the dynamic system reflects the fact that both $x(t)$ and $q(t)$ change simultaneously, and *not* in the way such that, say, $x(t)$ always optimizes the Lagrangian $H(x(t)) - q(t) \cdot x(t)$ given $q(t)$ (this would be a “dual” dynamic system), or vice versa. The term “greedy” refers to the key condition (13), which states that “control” $v(t)$ is always chosen within the “set of controls” V , so as to greedily

maximize $\nabla_x[H(x(t)) - q(t) \cdot x(t)] \cdot x'(t)$, i.e., the partial time derivative of the Lagrangian $H(x(t)) - q(t) \cdot x(t)$ with respect to $x(t)$ only. Another interpretation of condition (13), which is important for our analysis (and further justifies term “greedy”), is that control $v(t)$ always greedily maximizes the time derivative of function $H(x(t)) - (1/2) \sum_n q_n^2(t)$, with respect to both $x(t)$ and $q(t)$. (See Lemma 3.)

Remark 2. *Queueing interpretation.* Components of vector-function $q(t)$ can be viewed as dynamically changing *continuous* “queue lengths” (or, “amounts of fluid”) at a finite set of “nodes” $n \in \mathcal{N}$. Components of control $v(t)$ represent instantaneous rates of change of the queue lengths. (These rates can have any sign.) Queue lengths are not allowed to become negative (“reflected at 0”), that is, roughly speaking, if $q_n(t) = 0$ and $v_n(t) < 0$, then $q'_n(t) = 0$ (see (16)). Function $f(t)$ is simply an integral of $v(t)$, so that (14)-(15) is an integral form of (16). Function $x(t)$ is the exponentially weighted *average of past values of control* $v(t)$ (see (11), and also (39)-(40) below), and it determines the current “utility” $H(x(t))$ of the system. Condition (8) means that there exists a control which provides strictly negative “drift” to all the queues simultaneously. For any $v^* \in V^*$ and any $q^* \in Q^*$, $(x(t), q(t)) \equiv (v^*, q^*)$ is a stationary GPD-trajectory (which is verified directly). Such trajectory represents an optimal operating point of the system in the sense that it maximizes (in the long-term) system utility while keeping queues finite. This queueing interpretation of the dynamic system should *not* create a false impression that the system only models “parallel server”-type queueing systems, where “customers” (or, “traffic”) arrive in nodes only exogenously, and leave the system after being served by any node. In fact, as we will show in Section 4, GPD-trajectories arise as (asymptotic) limits for a very general *queueing network* model, where customers served by nodes may be routed for service in other nodes.

Theorem 1 *For any $x(0) \in \tilde{V}$ and any $q(0) \in R_+^N$, there exists a GPD-trajectory (x, q, f) having $(x(0), q(0), f(0) = q(0))$ as initial condition.*

Theorem 1 is a corollary (in fact - a part) of Lemma 12, formulated and proved in Section 3.6.

3.3 Global Attraction Result

The following theorem, showing that GPD-trajectories are such that $(x(t), q(t))$ is attracted to the saddle set $V^* \times Q^*$, is the main result of this paper.

Theorem 2 *Under the non-degeneracy condition (8), the following holds.*

(i) *For any GPD-trajectory (x, q, f) , as $t \rightarrow \infty$,*

$$x(t) \rightarrow V^*, \tag{17}$$

$$q(t) \rightarrow q^* \text{ for some } q^* \in Q^*. \tag{18}$$

(ii) Let compact subsets $V^\square \subset \tilde{V}$ and $Q^\square \subset R_+^N$ be fixed. Then, the convergence

$$(x(t), q(t)) \rightarrow V^* \times Q^* \text{ as } t \rightarrow \infty, \quad (19)$$

for GPD-trajectories is uniform with respect to initial conditions $(x(0), q(0)) \in V^\square \times Q^\square$.

Remark 1. In Theorem 2, the condition on set V is that it is convex and compact (as defined in Section 3.1). In particular, it is *not* required that V is a polyhedron. For the network model of Section 4, the corresponding “rate region” V (defined in Section 4.4) is in fact a (convex) compact polyhedron. Thus, Theorem 2, which is referred to in Section 4, is a more general result than Section 4 needs to establish asymptotic optimality of the GPD control algorithm introduced there.

Remark 2. Some generalizations of Theorem 2 are discussed in Section 3.8.

In the rest of Section 3, we first introduce (in Section 3.4) simple preliminaries needed to prove Theorem 2(i), and then prove it in Section 3.5. Then, in Section 3.6, we establish a set of (more involved) basic properties of the family of GPD-trajectories, including their existence (thus proving Theorem 1), which we need for the proof of Theorem 2(ii), presented in Section 3.7. (This layout is such that a reader interested only in the proof of convergence result of Theorem 2(i), can be spared the technical preliminaries needed for the proof of Theorem 2(ii) only.)

3.4 Preliminary Observations

We start with a characterization of optimal dual solutions, given below in Lemma 2, which follows from basic convex programming duality theory (cf. [17], Sections 28 and 30).

Suppose non-degeneracy condition (8) holds. By Kuhn-Tucker theorem, for any pair of optimal primal and dual solutions, $v^* \in V^*$ and $q^* \in Q^*$, the complementary slackness condition holds:

$$v^* \cdot q^* = 0. \quad (20)$$

Note that there always exists $v^* \in V^*$ for which the subset of $n \in \mathcal{N}$ with $v_n^* < 0$ is the maximum possible; this subset we denote by $\mathcal{N}^{(0)} \subseteq \mathcal{N}$. Thus, a vector $q^* \in R_+^N$ satisfies the complementary slackness condition (20) if and only if

$$q_n^* = 0 \text{ for } n \in \mathcal{N}^{(0)}. \quad (21)$$

For an optimal point $v^* \in V^*$, let $C^*(v^*)$ denote the normal cone to V at v^* . (It may have any dimension from 0 to N . A zero-dimensional cone is the one containing the single vector 0 - this is the case when v^* lies in the interior of V .) We know that $(v^*, q^*) \in V^* \times Q^*$ if and only if (v^*, q^*) is a saddle point of the Lagrangian $H(v) - y \cdot v$, $v \in V$, $y \in R_+^N$, for the pair of primal problem (4)-(5) and its dual (7). This implies the following property, recorded here for future reference.

Lemma 2 *Assume non-degeneracy condition (8). Then, the following holds for any fixed $v^* \in V^*$. Vector $q^* \in Q^*$ if and only if $q^* \in R_+^N$, $\nabla H(v^*) - q^* \in C^*(v^*)$ and the complementary slackness condition (20) (or, equivalently, (21)) holds. (Note that, by Proposition 1 in Appendix, $\nabla H(v^*)$ is same for all $v^* \in V^*$.)*

Next, we observe some basic properties of GPD-trajectories, which hold regardless of the condition (6) (or condition (8)).

Consider arbitrary fixed GPD-trajectory (x, q, f) . By Lemma 20 (in Appendix), the following properties hold:

$$\rho(x(t), V) \leq \rho(x(0), V)e^{-t}, \quad t \geq 0; \quad (22)$$

all three functions x , q and f are Lipschitz continuous in $[0, \infty)$; $x(\cdot)$ is uniformly bounded and, moreover, $x(t)$ for all $t \geq 0$ is contained within the (compact) convex hull of $V \cup \{x(0)\}$. The latter fact implies that $\nabla H(x(t))$ is uniformly bounded for all $t \geq 0$.

Let us call a time point $t > 0$ *regular* (for given GPD-trajectory (x, q, f)) if proper derivatives $x'(t)$, $q'(t)$, and $f'(t)$ exist, and conditions (11)-(13) hold for this t . (As shown in Section 6.1, for every regular t , (16) holds as well.) Almost all $t \geq 0$ are regular. To simplify notation, throughout the rest of this entire Section 3 we adopt a convention that any expression or statement involving any of the functions $x'(t)$, $q'(t)$, $f'(t)$, or $v(t)$, holds under the additional assumption that t is a regular point, even if we do not state it explicitly.

Let us introduce the following function:

$$F(v, y) = H(v) - \frac{1}{2} \sum_{n \in \mathcal{N}} y_n^2, \quad v \in \tilde{V}, \quad y \in R_+^N. \quad (23)$$

Lemma 3 *For any GPD-trajectory, at any (regular) $t \geq 0$,*

$$\frac{d}{dt}F(x(t), q(t)) = \nabla H(x(t)) \cdot (v(t) - x(t)) - q(t) \cdot v(t), \quad (24)$$

and

$$v(t) \in \arg \max_{v \in V} \nabla H(x(t)) \cdot (v - x(t)) - q(t) \cdot v. \quad (25)$$

(Expressions (24) and (25) imply that, given $x(t)$ and $q(t)$, $v(t)$ is a point in V maximizing $(d/dt)F(x(t), q(t))$.)

If, in addition, (6) holds (that is, V^* is non-empty), then for any $v^* \in V^*$

$$\frac{d}{dt}F(x(t), q(t)) \geq \nabla H(x(t)) \cdot (v^* - x(t)) \geq H(v^*) - H(x(t)). \quad (26)$$

Proof. We have:

$$\begin{aligned} \frac{d}{dt}F(x(t), q(t)) &= \nabla H(x(t)) \cdot (v(t) - x(t)) - \sum_n q_n(t)[v_n(t)]_{q_n(t)}^+ \\ &= \nabla H(x(t)) \cdot (v(t) - x(t)) - \sum_n q_n(t)v_n(t), \end{aligned}$$

which proves (24). Inclusion (25) is equivalent to (13). The first inequality in (26) follows from condition (25) and the fact that $q(t) \cdot v^* \leq 0$, and the second one from the subgradient inequality. \blacksquare

3.5 Proof of Theorem 2(i)

Throughout this Section 3.5, we always assume that we are in the conditions of Theorem 2(i), namely, that condition (8) holds and we consider a fixed GPD-trajectory (x, q, f) .

We start with a brief discussion of the key difficulties arising in the proof of Theorem 2(i). (This also gives an outline of the proof.) Lemma 3 shows that the “control” $v(t)$ is always chosen within “set of controls” V so as to maximize the derivative of $F(x(t), q(t))$. However, function $F(\cdot, \cdot)$ alone *cannot* serve as a Lyapunov function to prove the global attraction result, because it is *not* necessarily a non-decreasing function along a GPD-trajectory. We can claim that $(d/dt)F(x(t), q(t)) \geq 0$ when $x(t) \in V^{cond} = V \cap R_-^N$ (see (26)), but the latter condition on $x(t)$ clearly does not necessarily hold, roughly for the following two reasons. First, $x(t)$ is not necessarily within the set V at all. (In fact, it is easy to construct examples when $x(t)$ is never within V , and converges to optimal set $V^* \subseteq V$ “from outside” of V .) This problem is dealt with by using property (22), showing that $x(t)$ converges to V “very fast.” Second and more serious difficulty is that, even if $x(t) \in V$, it is not necessarily within R_-^N . To overcome this, we consider a different “Lyapunov” function F^* (see (27)), which is in fact non-decreasing as long as $x(t) \in V$ (see (29)). (Again, we only have $x(t) \rightarrow V$ as in (22), not $x(t) \in V$, but this turns out to be “good enough.”) Using F^* allows us to prove in Lemma 6 the convergence of $x(t)$ to a set $V^{max} \subseteq V$, which contains V^* , and convergence of $q(t)$ to some point within Q^* . (In the proof of Lemma 6, geometric reasoning plays important role.) Then, in Lemma 7, using Lemma 6 and some essentially queueing observations, we finally establish $x(t) \rightarrow R_-^N$, which implies $x(t) \rightarrow V^{cond} = V \cap R_-^N$. Only after that, we “finish off” in Lemma 8 by using F as a “Lyapunov” function. (Function F is also used at the very beginning, in Lemma 4, to provide crude estimates implying uniform boundedness of $q(t)$.)

Lemma 4 *Trajectory $(q(t), t \geq 0)$ is uniformly bounded:*

$$\sup_{t \geq 0} \|q(t)\| < \infty.$$

Proof. For any (regular) $t \geq 0$ we have:

$$\begin{aligned} \frac{d}{dt}F(x(t), q(t)) &= [\nabla H(x(t)) - q(t)] \cdot v(t) - \nabla H(x(t)) \cdot x(t) \\ &\geq [\nabla H(x(t)) - q(t)] \cdot v^{(1)} - \nabla H(x(t)) \cdot x(t) \\ &= -q(t) \cdot v^{(1)} + \nabla H(x(t)) \cdot (v^{(1)} - x(t)), \end{aligned}$$

where $v^{(1)}$ is some fixed element of $V \cap R_-^N$. (Such $v^{(1)}$ exists by (8).) Since both $\nabla H(x(t))$ and $x(t)$ are uniformly bounded, and all components of $v^{(1)}$ are strictly negative, we see that $(d/dt)F(x(t), q(t)) \geq \epsilon_1 > 0$ as long as $\|q(t)\| \geq C_1$, for some fixed constants ϵ_1 and C_1 . This implies (since $H(x(t))$ is uniformly bounded) that $(d/dt)F(x(t), q(t)) \geq \epsilon_2 > 0$ as long as $F(x(t), q(t)) \leq C_2$, for some fixed ϵ_2 and C_2 . This in turn means that $F(x(t), q(t))$ is uniformly bounded below, and therefore $q(t)$ is uniformly bounded. ■

Let us fix arbitrary optimal dual solution $q^* \in Q^*$, and associate with it the following function

$$F^*(v, y) \doteq H(v) - q^* \cdot v - \frac{1}{2} \sum_{n \in \mathcal{N}} (y_n - q_n^*)^2, \quad v \in \tilde{V}, \quad y \in R_+^N. \quad (27)$$

We also denote

$$H^*(v) \doteq H(v) - q^* \cdot v,$$

and so $F^*(v, y) = H^*(v) - (1/2) \sum_n (y_n - q_n^*)^2$. Function $H^*(v)$ is the Lagrangian $H(v) - y \cdot v$ of the problem (4)-(5), with the dual variable y equal to $q^* \in Q^*$. This implies that the convex compact set

$$V^{max} \doteq \arg \max_{v \in V} H^*(v)$$

contains all optimal solutions to the problem (4)-(5), i.e. $V^* \subseteq V^{max}$, and $H^*(v^*) = H(v^*)$ for any $v^* \in V^*$. By Proposition 1 (in Appendix), the gradient $\nabla H^*(v) = \nabla H(v) - q^*$ is same for all $v \in V^{max}$. Consequently, $\nabla H(v)$ is same for all $v \in V^{max}$.

Lemma 5 *Consider $F^*(\cdot, \cdot)$ associated with arbitrary $q^* \in Q^*$. Then, for all (regular) $t \geq 0$,*

$$\frac{d}{dt}F^*(x(t), q(t)) \geq [\nabla H(x(t)) - q^*] \cdot (v(t) - x(t)) - (q(t) - q^*) \cdot v(t) \quad (28)$$

and

$$x(t) \in V \text{ implies } \frac{d}{dt}F^*(x(t), q(t)) \geq 0. \quad (29)$$

Proof. Indeed,

$$\begin{aligned} \frac{d}{dt}F^*(x(t), q(t)) &= [\nabla H(x(t)) - q^*] \cdot (v(t) - x(t)) - (q(t) - q^*) \cdot q'(t) \\ &= [\nabla H(x(t)) - q^*] \cdot (v(t) - x(t)) - \sum_n (q_n(t) - q_n^*) [v_n(t)]_{q_n(t)}^+ \end{aligned}$$

$$\geq [\nabla H(x(t)) - q^*] \cdot (v(t) - x(t)) - \sum_n (q_n(t) - q_n^*) v_n(t),$$

where the inequality holds because, for any n , $q_n(t)[v_n(t)]_{q_n(t)}^+ = q_n(t)v_n(t)$ and $q_n^*[v_n(t)]_{q_n(t)}^+ \geq q_n^*v_n(t)$. This proves (28). To prove (29) we notice that, by condition (13), the RHS of (28) is maximized by $v(t)$ among all $v \in V$. Thus, for any fixed $v^* \in V^*$, the RHS of (28) is no less than

$$\begin{aligned} & [\nabla H(x(t)) - q^*] \cdot (v^* - x(t)) - (q(t) - q^*) \cdot v^* \\ & \geq [\nabla H(x(t)) - q^*] \cdot (v^* - x(t)) \\ & = \nabla H^*(x(t)) \cdot (v^* - x(t)) \geq H^*(v^*) - H^*(x(t)) \geq 0, \end{aligned}$$

where the very last inequality is because v^* maximizes $H^*(v)$ over V . ■

Lemma 6 Consider $F^*(\cdot, \cdot)$ associated with arbitrary fixed $q^* \in Q^*$. Let $v^* \in V^*$ be fixed such that $v_n^* < 0$ for each $n \in \mathcal{N}^{(0)}$. The following properties hold, as $t \rightarrow \infty$.

- (i) $x(t) \rightarrow V^{max}$, which in particular implies that $\nabla H(x(t)) \rightarrow \nabla H(v^*)$.
- (ii) $q_n(t) \rightarrow 0$ for every $n \in \mathcal{N}^{(0)}$.
- (iii) $[\nabla H(x(t)) - q(t)] \rightarrow C^*(v^*)$.
- (iv) Both $F^*(x(t), q(t))$ and $H^*(x(t))$ converge (to some constants). Consequently, $\sum_n (q_n(t) - q_n^*)^2 = \|q(t) - q^*\|^2$ converges.
- (v) $q(t) \rightarrow q^{**}$, for some fixed element $q^{**} \in Q^*$.

Proof. All statements of the lemma will follow from a number of auxiliary results, which we first derive.

For any (regular) $t \geq 0$,

$$\begin{aligned} \frac{d}{dt} F^*(x(t), q(t)) & \geq [\nabla H(x(t)) - q^*] \cdot (v(t) - x(t)) - (q(t) - q^*) \cdot v(t) \\ & = [\nabla H(x(t)) - q^*] \cdot (v^* - x(t)) - (q(t) - q^*) \cdot v^* + [\nabla H(x(t)) - q(t)] \cdot (v(t) - v^*) \quad (30) \\ & = B_1(t) + B_2(t) + B_3(t), \end{aligned}$$

where $B_i(t)$ is the i th term in the RHS of (30).

We always have

$$B_2(t) = -(q(t) - q^*) \cdot v^* = -q(t) \cdot v^* \geq 0.$$

Moreover, for any $\epsilon_1 > 0$ there exists sufficiently small $\epsilon_2 > 0$ such that

$$B_2(t) \geq \epsilon_2 \quad \text{as long as} \quad q_n(t) \geq \epsilon_1 \quad \text{for at least one } n \in \mathcal{N}^{(0)}. \quad (31)$$

Analogous properties hold for $B_3(t)$. Namely,

$$B_3(t) = [\nabla H(x(t)) - q(t)] \cdot (v(t) - v^*) \geq 0,$$

because $v(t)$ maximizes $[\nabla H(x(t)) - q(t)] \cdot v$ over all $v \in V$. Moreover, for any $\epsilon_1 > 0$ there exists sufficiently small $\epsilon_2 > 0$ such that

$$B_3(t) \geq \epsilon_2 \quad \text{as long as} \quad \rho([\nabla H(x(t)) - q(t)], C^*(v^*)) \geq \epsilon_1. \quad (32)$$

The proof of (32) requires some technical details - it is presented in Section 6.3 in Appendix. (Note, that if V is a polyhedron, property (32) is a rather simple observation. But, in the setting of Theorem 2, V is not necessarily polyhedral.)

Now, let us turn to $B_1(t)$. Let us denote by $x^*(t)$ the normal projection of $x(t)$ onto V ; namely, $x^*(t)$ is the (unique) point of V which is the closest to $x(t)$. According to (22),

$$\|x(t) - x^*(t)\| \leq \|x(0) - x^*(0)\|e^{-t}, \quad t \geq 0.$$

We have

$$\begin{aligned} B_1(t) &= \nabla H^*(x(t)) \cdot (v^* - x(t)) \\ &\geq H^*(v^*) - H^*(x(t)) = B_{11}(t) + B_{12}(t), \end{aligned}$$

where $B_{11}(t) = H^*(v^*) - H^*(x^*(t))$ and $B_{12}(t) = H^*(x^*(t)) - H^*(x(t))$.

For $B_{11}(t)$ we have

$$B_{11}(t) \geq 0,$$

and, moreover, for any $\epsilon_1 > 0$ there exists sufficiently small $\epsilon_2 > 0$ such that

$$B_{11}(t) \geq \epsilon_2 \quad \text{as long as} \quad \rho(x^*(t), V^{max}) \geq \epsilon_1. \quad (33)$$

For the function $B_{12}(t), t \geq 0$, the following estimate holds for any pair $0 \leq t_1 \leq t_2 \leq \infty$:

$$\begin{aligned} \int_{t_1}^{t_2} |B_{12}(s)| ds &\leq \int_{t_1}^{t_2} C_1 \|x^*(s) - x(s)\| ds \\ &\leq C_1 \int_{t_1}^{t_2} \|x^*(0) - x(0)\| e^{-s} ds = C_1 \|x^*(0) - x(0)\| [e^{-t_1} - e^{-t_2}] \\ &\leq C_1 \|x^*(0) - x(0)\| < \infty, \end{aligned} \quad (34)$$

where $C_1 > 0$ is a uniform upper bound on $\|\nabla H^*(x(s))\|$ and $\|\nabla H^*(x^*(s))\|$ over all $s \geq 0$. (For example, C_1 can be chosen as the maximum of $\|\nabla H^*(\xi)\|$ over all ξ in the convex hull of $V \cup \{x(0)\}$.)

Proof of (i). It is easy to see that, since V is convex, $\|x^*(t_2) - x^*(t_1)\| \leq \|x(t_2) - x(t_1)\|$ for any $t_1, t_2 \geq 0$. Then, Lipschitz continuity of $x(\cdot)$ implies that $x^*(\cdot)$ is also Lipschitz. The above estimates for $B_{11}(t)$ and $B_{12}(t)$ (along with the Lipschitz continuity of $x^*(\cdot)$ and non-negativity of $B_2(\cdot)$ and $B_3(\cdot)$) show that $\rho(x^*(t), V^{max})$ (and then $\rho(x(t), V^{max})$ as well) must converge to 0. (Otherwise we would have $\int_0^\infty \frac{d}{dt} F^*(x(t), q(t)) dt = \infty$ which is impossible since $F^*(x(t), q(t))$ is uniformly bounded.) This proves (i).

Property (ii) is obtained similarly to (i), relying on estimates (31) and (34) (along with the Lipschitz continuity of $q(\cdot)$ and non-negativity of $B_{11}(\cdot)$ and $B_3(\cdot)$). Property (iii) is also

obtained the same way, relying on estimates (32) and (34) (along with the convergence of $\nabla H(x(t))$, Lipschitz continuity of $q(\cdot)$ and non-negativity of $B_{11}(\cdot)$ and $B_2(\cdot)$).

Proof of (iv). Convergence of $H^*(x(t))$ follows directly from (i). Function $F^*(x(t), q(t))$ converges because it is absolutely continuous (in fact - Lipschitz) and bounded, its derivative $(d/dt)F^*(x(t), q(t)) \geq B_{12}(t)$, and due to the estimate (34).

Proof of (v). Consider any limiting point q^{**} of the trajectory $(q(t), t \geq 0)$, which exists since the trajectory is bounded. We must have $[\nabla H(v^*) - q^{**}] \in C^*(v^*)$, by (i) and (iii). In addition $q_n^{**} = 0$ for every $n \in \mathcal{N}^{(0)}$, by (ii), implying $v^* \cdot q^{**} = 0$. Therefore, by Lemma 2, $q^{**} \in Q^*$. Note that properties (i)-(iv) of this Lemma hold for any a priori fixed $q^* \in Q^*$, including q^{**} . Therefore, by (iv), $\|q(t) - q^{**}\|$ converges, and it can only converge to 0. ■

Lemma 7 *The following property holds:*

$$\limsup_{t \rightarrow \infty} x_n(t) \leq 0, \quad n \in \mathcal{N}.$$

Consequently, as $t \rightarrow \infty$, $x(t) \rightarrow V^{max} \cap R_-^N \subseteq V^{cond}$.

Proof. We have the following lower bound for any $n \in \mathcal{N}$ and any pair $0 \leq t_1 \leq t_2 < \infty$:

$$q_n(t_2) - q_n(t_1) \geq \int_{t_1}^{t_2} v_n(t) dt = x_n(t_2) - x_n(t_1) + \int_{t_1}^{t_2} x_n(t) dt.$$

If we further assume that $x_n(t) \geq \epsilon_1 > 0$ for all $t \in [t_1, t_2]$, then

$$q_n(t_2) - q_n(t_1) \geq x_n(t_2) - x_n(t_1) + \epsilon_1(t_2 - t_1), \quad (35)$$

and if, in addition, $x_n(t_1) = \epsilon_1 < x_n(t_2) = \epsilon_2$, then

$$q_n(t_2) - q_n(t_1) \geq \epsilon_2 - \epsilon_1. \quad (36)$$

Thus, for any fixed positive levels $\epsilon_1 < \epsilon_2$, we have the following. First, by (35) and boundedness of $x_n(\cdot)$, if $x_n(s) > \epsilon_1$ at some time s , it must hit level ϵ_1 within finite time after s - otherwise $q_n(t) \rightarrow \infty$ which is impossible. Second, the trajectory of $x_n(t)$ can move from level ϵ_1 to ϵ_2 only finite number of times - otherwise, by (36), $q_n(t)$ would never converge, which would contradict to Lemma 6(v). ■

Lemma 8 *We have*

$$x(t) \rightarrow V^* \quad \text{as } t \rightarrow \infty.$$

Proof. We know that $q(t)$ converges to some $q^* \in Q^*$ and that $x(t) \rightarrow V^{cond}$. Therefore, it will suffice to show that $F(x(t), q(t)) \rightarrow F(v^*, q^*)$, where the function $F(\cdot, \cdot)$ is defined in (23)

and v^* is an arbitrarily picked element of V^* . Obviously, $\limsup_t F(x(t), q(t)) \leq F(v^*, q^*)$. Let us show that

$$\liminf_t F(x(t), q(t)) \geq F(v^*, q^*). \quad (37)$$

Let us fix arbitrary $\epsilon > 0$. Since $q(t)$ converges, for all sufficiently large t , $F(x(t), q(t)) \leq F(v^*, q^*) - 2\epsilon$ implies $H(v^*) - H(x(t)) > \epsilon$. We have

$$\begin{aligned} \frac{d}{dt}F(x(t), q(t)) &= \nabla H(x(t)) \cdot (v(t) - x(t)) - q(t) \cdot v(t) \\ &\geq \nabla H(x(t)) \cdot (v^* - x(t)) - q(t) \cdot v^* \geq [H(v^*) - H(x(t))] - q(t) \cdot v^* . \end{aligned} \quad (38)$$

In the RHS of (38), first, $q(t) \cdot v^* \rightarrow q^* \cdot v^* = 0$ and, second, for all large t , $H(v^*) - H(x(t)) > \epsilon$ as long as $F(x(t), q(t)) \leq F(v^*, q^*) - 2\epsilon$. This means that, for large t , $(d/dt)F(x(t), q(t))$ is strictly positive and bounded away from 0 as long as $F(x(t), q(t)) \leq F(v^*, q^*) - 2\epsilon$. Since this holds for any $\epsilon > 0$, (37) follows. ■

Thus, (18) and (17) are proved by Lemma 6(v) and Lemma 8, respectively. The proof of Theorem 2(i) is complete. ■

3.6 Further Basic Properties of GPD-trajectories

Results of this section hold regardless of the condition (6) (and therefore regardless of (8)).

It follows directly from its definition that every GPD-trajectory (x, q, f) (if exists) is such that f is Lipschitz continuous in $[0, \infty)$ and, moreover, $f'(t) \in V$ almost everywhere. We now introduce and study some operators which allow us to characterize GPD-trajectories (x, q, f) as essentially fixed points of a multivalued operator, mapping a set of Lipschitz continuous functions into itself. This will allow us to establish existence and also obtain some properties of the family of GPD-trajectories which we will need in Section 3.7 to prove Theorem 2(ii).

We remind that $C_{R^N}[0, \infty)$ [resp. $C_{R_+^N}[0, \infty)$] denotes the set of continuous functions on $[0, \infty)$, taking values in R^N [resp. R_+^N]. On $C_{R^N}[0, \infty)$ [and $C_{R_+^N}[0, \infty)$] we always consider the topology induced by u.o.c. convergence. The space $C_{R^N}[0, \infty)$ can be viewed as a Banach space (that is, a complete normed space), if it is equipped with the following norm

$$\|f\|^* = \sum_{i=1}^{\infty} 2^{-i} \frac{\|f\|_i}{1 + \|f\|_i},$$

where

$$\|f\|_c \doteq \sup_{0 \leq t \leq c} \|f(t)\|.$$

Note that the topology on $C_{R^N}[0, \infty)$ induced by $\|\cdot\|^*$ -norm coincides with the topology of u.o.c. convergence.

Consider the subset $L_{R^N} \subset C_{R^N}[0, \infty)$, consisting of Lipschitz continuous functions f such that $f(0) \in R_+^N$ and $f'(t) \in V$ almost everywhere. (Since V is convex, the latter condition is equivalent to $[f(t_2) - f(t_1)]/[t_2 - t_1] \in V$ for all $0 \leq t_1 < t_2$.) For a fixed $\xi \in R^N$, we denote by $L_{R^N}(\xi) \doteq \{f \in L_{R^N} \mid f(0) = \xi\}$ the restriction of L_{R^N} to functions f with $f(0) = \xi$. The following result is easily verified and presented here for future reference.

Lemma 9 *The set L_{R^N} is a closed subset of $C_{R^N}[0, \infty)$. Moreover, for any compact set $A \subset R_+^N$, set*

$$\{f \in L_{R^N} \mid f(0) \in A\} = \cup_{\xi \in A} L_{R^N}(\xi)$$

is compact.

Now, let us define operator \mathcal{A}_1 , which takes $(f, x(0)) \in L_{R^N} \times \tilde{V}$ into $\mathcal{A}_1(f, x(0)) = (x, q) \in C_{R^N}[0, \infty) \times C_{R_+^N}[0, \infty)$, as follows:

(a) function $x \in C_{R^N}[0, \infty)$ is the unique solution of the differential equation

$$x'(t) = f'(t) - x(t), \quad t \geq 0, \quad \text{a.e.}, \quad (39)$$

with fixed initial condition $x(0)$, that is,

$$x(t) = x(0)e^{-t} + \int_0^t f'(t - \xi)e^{-\xi}d\xi = x(0)e^{-t} + [f(t) - f(0)]e^{-t} - \int_0^t f(t - \xi)e^{-\xi}d\xi, \quad t \geq 0; \quad (40)$$

(b) components of function $q \in C_{R_+^N}[0, \infty)$ are obtained from the components of f by the reflection mapping at 0 (as in (15)):

$$q_n(t) = f_n(t) - [0 \wedge \inf_{0 \leq \xi \leq t} f_n(\xi)], \quad t \geq 0, \quad n \in \mathcal{N}. \quad (41)$$

Note that (41) implies $q(0) = f(0) \in R_+^N$.

Lemma 10 *Operator \mathcal{A}_1 has the following properties.*

(i) *Mapping \mathcal{A}_1 is continuous.*

(ii) *Image functions q are Lipschitz continuous, uniformly across all possible f .*

(iii) *Let a compact convex set V^\square , such that $V \subseteq V^\square \subset \tilde{V}$, be fixed. Then, uniformly on $x(0) \in V^\square$ and all possible f , the image functions x are Lipschitz continuous and are such that $x(t) \in V^\square$ for all $t \geq 0$.*

Proof. Statements (i) and (ii) are verified directly, and (iii) follows from Lemma 20 (in Appendix). ■

Denote $C_{\tilde{V}}[0, \infty) \doteq \{x \in C_{R^N}[0, \infty) \mid x(t) \in \tilde{V}, t \geq 0\}$. We define a multivalued operator \mathcal{A}_2 , which takes $(x, q) \in C_{\tilde{V}}[0, \infty) \times C_{R_+^N}[0, \infty)$ to the set $\mathcal{A}_2(x, q) \subset L_{R^N}(q(0))$, as follows: function $f \in \mathcal{A}_2(x, q)$ if and only if $f \in L_{R^N}(q(0))$ (which implies $f(0) = q(0)$) and conditions (12)-(13) hold, namely,

$$f'(t) \in \arg \max_{v \in V} (\nabla H(x(t)) - q(t)) \cdot v, \quad t \geq 0, \quad \text{a.e..}$$

Lemma 11 *Operator \mathcal{A}_2 has the following properties.*

(i) *Mapping \mathcal{A}_2 is closed. Namely, if sequences of $(x^{(i)}, q^{(i)})$ and $f^{(i)}$, with $i = 1, 2, \dots$, are such that $(x^{(i)}, q^{(i)}) \rightarrow (x, q) \in C_{\tilde{V}}[0, \infty) \times C_{R_+^N}[0, \infty)$, $f^{(i)} \in \mathcal{A}_2(x^{(i)}, q^{(i)})$ and $f^{(i)} \rightarrow f$, as $i \rightarrow \infty$, then $f \in \mathcal{A}_2(x, q)$.*

(ii) *For any (x, q) , the image set $\mathcal{A}_2(x, q)$ is a compact convex subset of L_{R^N} .*

Proof. (i) It suffices to show that at any point $t > 0$, where the derivative $f'(t)$ exists,

$$f'(t) \in V(t) \doteq \arg \max_{v \in V} (\nabla H(x(t)) - q(t)) \cdot v. \quad (42)$$

It follows from Proposition 2 (in Section 6.4), that for any $\epsilon > 0$, there exists a sufficiently small $\delta > 0$ such that for almost all $\xi \in [t, t + \delta]$ and all sufficiently large indices i , the set $\arg \max_{v \in V} (\nabla H(x^{(i)}(\xi)) - q^{(i)}(\xi)) \cdot v$ is a subset of $V(t; \epsilon) \doteq \{v \in V \mid \rho(v, V(t)) \leq \epsilon\}$. This implies that

$$[f(t + \delta) - f(t)]/\delta \in V(t; \epsilon).$$

Since $V(t; \epsilon) \downarrow V(t)$ as $\epsilon \downarrow 0$, and we can choose a sequence of pairs (ϵ, δ) converging to $(0, 0)$, we obtain (42).

(ii) Since $\mathcal{A}_2(x, q)$ is a subset of a compact set $L_{R^N}(q(0))$ (see Lemma 9) its compactness follows from its closedness. Convexity is verified directly. \blacksquare

Lemma 12 (i) *A triple (x, q, f) is a GPD-trajectory if and only if $f \in L_{R^N}$, $x(0) \in \tilde{V}$, $(x, q) = \mathcal{A}_1(f, x(0))$ and $f \in \mathcal{A}_2(x, q)$.*

(ii) *For any $x(0) \in \tilde{V}$ and any $q(0) \in R_+^N$, there exists a GPD-trajectory (x, q, f) having $(x(0), q(0), f(0) = q(0))$ as initial condition.*

(iii) *The set of GPD-trajectories (x, q, f) is such that:*

(iii.1) *functions f and q are Lipschitz continuous, uniformly across all GPD-trajectories,*
(iii.2) *for arbitrary compact convex set V^\square , $V \subseteq V^\square \subset \tilde{V}$, uniformly on $x(0) \in V^\square$, the functions x are Lipschitz continuous and are such that $x(t) \in V^\square$ for all $t \geq 0$.*

(iv) *The set of GPD-trajectories is closed.*

(v) *Let compact sets $V^\square \subset \tilde{V}$ and $Q^\square \subset R_+^N$ be fixed. Then, the set of the GPD-trajectories with $x(0) \in V^\square$ and $q(0) = f(0) \in Q^\square$ is compact.*

Proof. Statement (i) follows directly from the definitions of operators \mathcal{A}_1 and \mathcal{A}_2 , and Lipschitz properties of their images, as described in Lemmas 10 and 11.

Proof of (ii). According to (i), a GPD-trajectory with initial state $(x(0), q(0), f(0) = q(0))$ exists if and only if the multivalued operator

$$\mathcal{A}_{x(0)}(f) \doteq \mathcal{A}_2 \mathcal{A}_1(f, x(0)),$$

which takes $f \in L_{R^N}(x(0))$ into a subset $\mathcal{A}_{x(0)}(f) \subseteq L_{R^N}(x(0))$, has a fixed point, that is $f \in \mathcal{A}_{x(0)}(f)$ for some f . Operator $\mathcal{A}_{x(0)}$ maps a compact convex subset $L_{R^N}(x(0))$ of the Banach space $C_{R^N}[0, \infty)$ into the set of subsets of $L_{R^N}(x(0))$. From Lemmas 10 and 11 we see that operator $\mathcal{A}_{x(0)}$ is closed, and its images $\mathcal{A}_{x(0)}(f)$ are compact and convex. Thus, by Kakutani theorem (cf. [9], Theorem XVI.5.1), operator $\mathcal{A}_{x(0)}$ has a fixed point.

Statements (iii) and (iv) follow from the properties of operators \mathcal{A}_1 and \mathcal{A}_2 described in Lemmas 10 and 11.

Proof of (v). Without loss of generality, we assume that V^\square is a compact convex set such that $V \subseteq V^\square \subset \tilde{V}$, as in (iii.2). By (iii), any sequence $(x^{(i)}, q^{(i)}, f^{(i)})$, $i = 1, 2, \dots$, of GPD-trajectories with $x^{(i)}(0) \in V^\square$ and $q^{(i)}(0) = f^{(i)}(0) \in Q^\square$ is such that, uniformly on i , all three component functions $x^{(i)}$, $q^{(i)}$ and $f^{(i)}$ are Lipschitz. This means that we can find a subsequence, converging u.o.c. to a triple (x, q, f) with $x(0) \in V^\square$ and $q(0) = f(0) \in Q^\square$. By (iv), (x, q, f) is a GPD-trajectory. \blacksquare

We will also need a shift invariance property of the family of GPD-trajectories. For a GPD-trajectory (x, q, f) and a constant $\tau \geq 0$, let us define its shifted version

$$\Theta_\tau(x, q, f) = (\Theta_\tau x, \Theta_\tau q, \Theta_\tau f)$$

as follows:

$$\begin{aligned} (\Theta_\tau x)(t) &= x(\tau + t), \quad (\Theta_\tau q)(t) = q(\tau + t), \quad t \geq 0, \\ (\Theta_\tau f)(t) &= q(\tau) + [f(\tau + t) - f(\tau)], \quad t \geq 0. \end{aligned}$$

The following fact is verified directly.

Lemma 13 *If (x, q, f) is a GPD-trajectory, then, for any $\tau \geq 0$, $\Theta_\tau(x, q, f)$ is also a GPD-trajectory.*

3.7 Proof of Theorem 2(ii)

Throughout this Section 3.7, we always assume that we are in the conditions of Theorem 2(ii), in particular, that condition (8) holds. Without loss of generality, we assume that V^\square is a compact convex set such that $V \subseteq V^\square \subset \tilde{V}$, as in Lemma 12(iii.2). We denote by Φ^\square the set of GPD-trajectories (x, q, f) with $(x(0), q(0)) \in V^\square \times Q^\square$.

Lemma 14 *There exists a compact subset $Q^{\square\square} \subset R_+^N$, such that, for all GPD-trajectories $(x, q, f) \in \Phi^{\square}$,*

$$(x(t), q(t)) \in V^{\square} \times Q^{\square\square}, \quad t \geq 0.$$

Proof. The fact that $x(t) \in V^{\square}$ for all t , uniformly on Φ^{\square} , follows from Lemma 12(iii.2). Given that both $x(t)$ and $\nabla H(x(t))$ are uniformly bounded, the proof of Lemma 4 goes through without change to show that $\|q(t)\|$ is uniformly bounded as well. ■

Lemma 15 *For any $\epsilon > 0$, there exists $T(\epsilon) > 0$ such that, uniformly on $(x, q, f) \in \Phi^{\square}$ and on $\tau \geq 0$,*

$$\inf\{t \geq \tau \mid \rho[(x(t), q(t)), V^* \times Q^*] < \epsilon\} < \tau + T(\epsilon).$$

Proof. Suppose not. Then, for some fixed $\epsilon > 0$, there exist sequences $\{(x^{(i)}, q^{(i)}, f^{(i)}) \in \Phi^{\square}\}$, $\{\tau^{(i)} \geq 0\}$ and $\{T^{(i)} > 0\}$, with index $i = 1, 2, \dots$, such that $T^{(i)} \uparrow \infty$ and

$$\inf_{\tau^{(i)} \leq t \leq \tau^{(i)} + T^{(i)}} \rho[(x(t), q(t)), V^* \times Q^*] \geq \epsilon.$$

The sequence of shifted GPD-trajectories $\Theta_{\tau^{(i)}}(x^{(i)}, q^{(i)}, f^{(i)})$ is such that their initial states $(\Theta_{\tau^{(i)}}x^{(i)}(0), \Theta_{\tau^{(i)}}q^{(i)}(0)) \in V^{\square} \times Q^{\square\square}$. Therefore, we can find a subsequence converging to a GPD-trajectory (x, q, f) such that $(x(t), q(t))$ never converges to $V^* \times Q^*$. This contradiction completes the proof. ■

Lemma 15 shows that to prove Theorem 2(ii), it would suffice to prove that a GPD-trajectory with initial state “close to $V^* \times Q^*$ ” stays close to this set. However, as we will see shortly, we will only need the following, somewhat weaker, property.

Lemma 16 *Consider a point $(v^*, q^*) \in V^* \times Q^*$. Let $\Phi(v^*, q^*; \epsilon)$ be the set of all GPD-trajectories such that $\rho[(x(0), q(0)), (v^*, q^*)] \leq \epsilon$ and $x(0) \in V$. Then,*

$$\sup_{\Phi(v^*, q^*; \epsilon)} \sup_{t \geq 0} \rho[(x(t), q(t)), V^* \times q^*] \downarrow 0, \quad \text{as } \epsilon \downarrow 0.$$

Proof. Within this proof, we will say that a (scalar or vector) variable h “is $o(1; \epsilon)$ ” if $\|h\| \leq \delta(\epsilon)$, for some positive function $\delta(\epsilon)$ of $\epsilon > 0$ vanishing as $\epsilon \downarrow 0$.

Consider a GPD-trajectory $(x, q, f) \in \Phi(v^*, q^*; \epsilon)$, and function $F^*(x(t), q(t))$ associated with q^* . Since $x(0) \in V$, the entire trajectory of $x(t)$ lies in V (see (22)). Then, by Lemma 5, $F^*(x(t), q(t))$ is non-decreasing and

$$F^*(x(0), q(0)) \leq F^*(x(t), q(t)) \leq F^*(v^*, q^*) = H^*(v^*) = H(v^*),$$

which implies

$$F^*(x(0), q(0)) \leq F^*(x(t), q(t)) \leq H^*(x(t)) \leq H^*(v^*) = H(v^*).$$

We see that, uniformly on $(x, q, f) \in \Phi(v^*, q^*; \epsilon)$ and $t \geq 0$, both $|F^*(x(t), q(t)) - H(v^*)|$ and $|H^*(x(t)) - H(v^*)|$, and consequently $\|q(t) - q^*\|$, are $o(1; \epsilon)$. This implies that, uniformly on $\Phi(v^*, q^*; \epsilon)$, $0 \leq t_1 \leq t_2$ and $n \in \mathcal{N}$, the increment $\|q_n(t_2) - q_n(t_1)\|$ is $o(1; \epsilon)$. Then, the argument and the estimates we used in the proof of Lemma 7 show that, uniformly on $\Phi(v^*, q^*; \epsilon)$ and $t \geq 0$, $\rho(x(t), R_-^N)$ (and then $\rho(x(t), V \cap R_-^N)$) is $o(1; \epsilon)$.

We know that (uniformly on $\Phi(v^*, q^*; \epsilon)$) $|F(x(0), q(0)) - F(v^*, q^*)|$ is $o(1; \epsilon)$. Let us show that $[F(v^*, q^*) - F(x(t), q(t))] \vee 0$ is $o(1; \epsilon)$, uniformly on $\Phi(v^*, q^*; \epsilon)$ and $t \geq 0$. (This means that the lower bound of $F(x(t), q(t))$ is close to $F(v^*, q^*)$ when ϵ is small.) Suppose not. Then, there exists $\epsilon_1 > 0$ such that, no matter how small $\epsilon > 0$ is, there exists a GPD-trajectory $(x, q, f) \in \Phi(v^*, q^*; \epsilon)$ such that in some time interval the value of $F(x(t), q(t))$ changes from $F(v^*, q^*) - 2\epsilon_1$ to $F(v^*, q^*) - 3\epsilon_1$. This is impossible, however, because when ϵ (and then $\|q(t) - q^*\|$ for all t) is small enough, inequality $F(x(t), q(t)) \leq F(v^*, q^*) - 2\epsilon_1$ implies $H(x(t)) \leq H(v^*) - \epsilon_1$, in which case $(d/dt)F(x(t), q(t)) > 0$ (by Lemma 3).

Finally, since $[F(v^*, q^*) - F(x(t), q(t))] \vee 0$ is $o(1; \epsilon)$, so is $[H(v^*) - H(x(t))] \vee 0$; and so is $\rho(x(t), V \cap R_-^N)$ as we established earlier. This implies that $\rho(x(t), V^*)$ is $o(1; \epsilon)$, uniformly on $\Phi(v^*, q^*; \epsilon)$ and $t \geq 0$, which completes the proof. \blacksquare

Remainder of the proof of Theorem 2(ii). Suppose the statement is not true. Then, for some fixed $\epsilon_1 > 0$ the following property holds. There exist sequences $\{(x^{(i)}, q^{(i)}, f^{(i)}) \in \Phi^\square\}$ and $\{\eta^{(i)} \geq 0\}$ such that $\eta^{(i)} \uparrow \infty$ and

$$\rho[(x^{(i)}(\eta^{(i)}), q^{(i)}(\eta^{(i)})), V^* \times Q^*] \geq \epsilon_1.$$

Let us choose arbitrary ϵ , $0 < \epsilon < \epsilon_1$. By Lemma 15, for all large i the following time

$$\tau^{(i)} = \max\{t \in [0, \eta^{(i)}] \mid \rho[(x^{(i)}(t), q^{(i)}(t)), V^* \times Q^*] \leq \epsilon\}$$

is well defined, and moreover sequence $\{T^{(i)} = \eta^{(i)} - \tau^{(i)} > 0\}$ is bounded above (by Lemma 15) and is bounded below away from 0 (because all functions $x^{(i)}$ and $q^{(i)}$ are uniformly Lipschitz). Recall also that, uniformly on $(x, q, f) \in \Phi^\square$, $x(t) \rightarrow V$ as $t \rightarrow \infty$ (by (22)) and both x and q are bounded (by Lemma 14). Then, we can choose a subsequence of indices i along which

$$\Theta_{\tau^{(i)}}(x^{(i)}, q^{(i)}, f^{(i)}) \rightarrow (x, q, f),$$

where (x, q, f) is a GPD-trajectory such that $x(0) \in V$, $\rho[(x(0), q(0)), (v^*, q^*)] \leq \epsilon$ for some $(v^*, q^*) \in V^* \times Q^*$, and $\rho[(x(T), q(T)), V^* \times Q^*] \geq \epsilon_1$ for some $T > 0$. Since for the fixed ϵ_1 we can choose arbitrarily small ϵ , we obtain (using compactness of $V^* \times Q^*$) a contradiction with Lemma 16. \blacksquare

3.8 Generalizations of Theorem 2

3.8.1 Free Variables in the Optimization Problem

Consider the following generalization of problem (4)-(5):

$$\max_{v \in V} H(v) \quad (43)$$

subject to

$$v_n \in R, \quad n \in \mathcal{N}^{(f)}, \quad \text{and} \quad v_n \in R_-, \quad n \in \mathcal{N} \setminus \mathcal{N}^{(f)}, \quad (44)$$

where $\mathcal{N}^{(f)}$ is a fixed subset of \mathcal{N} . (Superscript (f) stands for “free variables.”) Let us denote by V^* the set of all optimal solutions of (43)-(44), and by Q^* the set of all vectors $q^* \in R^N$ such that q_n^* , $n \in \mathcal{N} \setminus \mathcal{N}^{(f)}$, are optimal dual variables and $q_n^* = 0$ for $n \in \mathcal{N}^{(f)}$ by convention.

Now let us define generalized GPD-trajectories, corresponding to problem (43)-(44). In terms of a pair (x, q) , they are defined as trajectories satisfying conditions (10), (11), (13), (14) and the following generalized form of (16):

$$q_n(t) \geq 0, \quad q_n'(t) = [v_n(t)]_{q_n(t)}^+, \quad n \in \mathcal{N} \setminus \mathcal{N}^{(f)}, \quad \text{and} \quad q_n(t) \equiv 0, \quad n \in \mathcal{N}^{(f)}. \quad (45)$$

We claim that *Theorem 2 holds as is for the optimization problem (43)-(44), if GPD-trajectories are defined more generally as solutions to (10), (11), (13), (14), (45).*

To see this, for a fixed constant $c > 0$, let us consider the following (shift) transform σ of R^N : $(\sigma\xi)_n = \xi_n - c$ if $n \in \mathcal{N}^{(f)}$, and $(\sigma\xi)_n = \xi_n$ otherwise. Let us choose c large enough so that for all $v \in V$ and $n \in \mathcal{N}^{(f)}$, the values of $v_n - c$ are negative and bounded away from 0. (We can do this because V is bounded.) Consider the following optimization problem:

$$\max_{u \in \sigma V} H(\sigma^{-1}u) \quad \text{subject to} \quad u \in R_-^N. \quad (46)$$

It is easy to observe that the sets of optimal primal and dual solutions to the problem (46) are σV^* and Q^* , respectively. Moreover, the following simple correspondence is easily verified directly: (x, q) is a generalized GPD-trajectory (as defined in this Section 3.8.1) for the optimization problem (43)-(44) if and only if $(\sigma x, q)$ is a “conventional” GPD-trajectory (as defined in Section 3.2), with $q_n(0) = 0$, $n \in \mathcal{N}^{(f)}$, for the problem (46). But, Theorem 2 applies to the latter. This proves our claim that Theorem 2 still holds for problem (43)-(44) and generalized GPD-trajectories.

3.8.2 Weighted “Queue Lengths”

The following generalization of Theorem 2 also holds, and can be naturally combined with the generalization described in Section 3.8.1.

Consider arbitrary fixed (weighting) vector $\gamma \in R^N$ with all $\gamma_n > 0$. Consider a generalized definition of a GPD-trajectory, with condition (13) replaced by

$$v(t) \in \arg \max_{v \in V} (\nabla H(x(t)) - [\gamma q(t)]) \cdot v, \quad (47)$$

where $[\gamma q(t)]$ denotes vector $(\gamma_n q_n(t), n \in \mathcal{N})$. (GPD-trajectories generalized this way arise as asymptotic limits for the GPD control algorithm, which uses weighed queue lengths. See Remark 2 in Section 1.3.)

We claim that *Theorem 2 still holds, if we replace $q(t)$ by $[\gamma q(t)]$ in (13) (to obtain (47)), (18) and (19)*. It is important to emphasize that q^* and Q^* in Theorem 2 formulation *do not change*.

One way to prove this claim is to slightly adjust the proof of Theorem 2, in particular, use functions F and F^* defined as

$$F(v, y) = H(v) - \frac{1}{2} \sum_{n \in \mathcal{N}} \gamma_n y_n^2, \quad v \in \tilde{V}, y \in R_+^N,$$

$$F^*(v, y) \doteq H(v) - q^* \cdot v - \frac{1}{2} \sum_{n \in \mathcal{N}} \frac{1}{\gamma_n} (\gamma_n y_n - q_n^*)^2, \quad v \in \tilde{V}, y \in R_+^N,$$

as opposed to the definitions (23) and (27). Another (more neat) way is to map the “new” formulation of Theorem 2 into the original one by, roughly speaking, changing variables $q_n(t)$ to $\sqrt{\gamma_n} q_n(t)$. This, however, requires some redefining of the set of controls V . We do not provide details to save space.

4 Utility Based Resource Allocation in a Queueing Network subject to Stability Constraint

In this section we define and study a controlled queueing network. The model is such that, roughly speaking, each control action results in (a) network generating certain amounts of commodities and (b) a “queue processing action”, consisting in turn of (b.1) certain amounts of exogenous traffic arriving to and queued at the network nodes and (b.2) nodes processing certain amounts of traffic from their queues, with the processed traffic either leaving the network or being routed to other nodes. The “utility” of the network is a concave function H of the long-term average rates at which commodities are generated. Informally, the problem is to find a network control strategy which maximizes utility subject to the constraint that the network remains stable, namely, the queues at the network nodes remain bounded. In this section we, first, introduce the model and the optimization problem formally (in Sections 4.1-4.5). Then, we define a dynamic control policy, which we call GPD algorithm (Section 4.6). Finally, in Section 4.8 we prove asymptotic optimality of the algorithm, namely, that, roughly speaking, the algorithm becomes arbitrarily close to optimal as one of its parameters approaches 0.

We note that Sections 4.3 and 4.7 introduce equivalent forms of the network model and GPD algorithm, respectively, solely for the purposes of problem formulation and analysis. As far as applications are concerned, the basic forms of the model (in Sections 4.1 and 4.2) and the algorithm (in Section 4.6) suffice.

4.1 The Network Model

We consider a network consisting of a finite set of nodes $\mathcal{N} = \{1, 2, \dots, N\}$, $N \geq 1$. The nodes can be of two different types: N_u *utility nodes* form the subset $\mathcal{N}^u = \{1, 2, \dots, N_u\}$ and $N_p = N - N_u$ *processing nodes* form the subset $\mathcal{N}^p = \{N_u + 1, \dots, N\}$. Thus, $\mathcal{N} = \mathcal{N}^u \cup \mathcal{N}^p$. (The cases when either \mathcal{N}^u or \mathcal{N}^p is empty are allowed.) Each processing node $n \in \mathcal{N}^p$ has associated *queue*, formed by *customers* waiting for *processing* (or *service*) by the node. There is no predefined limit on the *queue length* (i.e., the number of customers in the queue) at any node.

The system operates in discrete time $t = 0, 1, 2, \dots$ as follows. (By convention, we will identify an integer time t with the unit time interval $[t, t + 1)$, which will sometimes be referred to as the *time slot* t .) The network has a finite set of *modes* M . The sequence of modes $m(t)$, $t = 0, 1, 2, \dots$, forms an irreducible (finite) Markov chain with stationary distribution $\{\pi_m, m \in M\}$, where all $\pi_m > 0$ and $\sum \pi_m = 1$. (The mode process $m(t)$ models the underlying randomly changing network “environment,” and is *not* affected by any network control.) When the network mode is $m \in M$, a finite number of *controls* is available, which form set $K(m)$. (We denote by $K \doteq \cup_m K(m)$ the finite set of all possible controls across all modes $m \in M$.) If a control $k \in K(m)$ is chosen at time t , then the following occurs (sequentially, as listed below):

- (a) each utility node $n \in \mathcal{N}^u$ generates an integer amount $b_n(k)$ of certain commodity. (The commodity may be, for example, the amount of “traffic”, although it can be anything, and $b_n(k)$ need not be non-negative);
- (b) each processing node $n \in \mathcal{N}^p$ serves integer number $\mu_n(k) \geq 0$ of customers from its queue (or the entire queue n content, if it is less than $\mu_n(k)$), which are then randomly and independently routed to other processing nodes (including possibly self) with probabilities $p_{nj}(k)$, $j \in \mathcal{N}^p$, $\sum_{j \in \mathcal{N}^p} p_{nj}(k) \leq 1$, or leave the network with probability $1 - \sum_{j \in \mathcal{N}^p} p_{nj}(k)$;
- (c) an integer number $\lambda_n(k) \geq 0$ of exogenous customers arrive into each processing node queue $n \in \mathcal{N}^p$.

If we denote by $Q_n(t)$ the queue size (number of queued customers) at a processing node $n \in \mathcal{N}^p$ at time t , and $k = k(t) \in K(m(t))$ is the control chosen at time t , then, according to steps (b) and (c) above,

$$Q_n(t + 1) = Q_n(t) - [Q_n(t) \wedge \mu_n(k)] + A_n(t) + \lambda_n(k) , \quad (48)$$

where $A_n(t)$ is the total number of customers routed to this (processing) node at time t after service at other processing nodes.

We make the following (non-restrictive in many applications) assumption on the set of con-

trols. For any mode $m \in M$, along with any control $k \in K(m)$, the set $K(m)$ also contains any control k' that is identical to k in every respect, except that $\mu_n(k') = 0$ for some (possibly empty) subset of processing nodes n (and $\mu_n(k') = \mu_n(k)$ for all other processing nodes). More precisely, $k \in K(m)$ implies $k' \in K(m)$, if the following holds: $b_n(k') = b_n(k)$ for all $n \in \mathcal{N}^u$, $\lambda_n(k') = \lambda_n(k)$ and $p_{nj}(k') = p_{nj}(k)$ for all $n, j \in \mathcal{N}^p$, and for each $n \in \mathcal{N}^p$ either $\mu_n(k') = \mu_n(k)$ or $\mu_n(k') = 0$.

Informally, the problem we are going to address is as follows. Let $u_n, n \in \mathcal{N}^u$, denote long-term average value of $b_n(k(t))$ for a utility node n , under a given dynamic control policy, that is, a policy of choosing control $k(t)$ depending on the network mode and state. (Thus, u_n is the average rate at which the commodity is produced by a utility node n .) We would like to find a dynamic control policy which maximizes some concave utility function $H(u_1, \dots, u_{N_u})$, subject to the constraint that the network of processing nodes remains stable, that is, (informally speaking) the processes $Q_n(t), t \geq 0$, for all $n \in \mathcal{N}^p$ remain bounded.

The network model described above is rather abstract. A reader who wishes to gain a more concrete understanding of it, may wish at this point to look through an application example described at the beginning of Section 5.1.

4.2 Remarks on the Model Assumptions

Assumption that $\lambda_n(k)$ and $\mu_n(k)$ for the processing nodes $n \in \mathcal{N}^p$ are fixed numbers is not essential - these quantities can be non-negative random variables with finite means. (In this case they will be i.i.d. in time but in general dependent across n at any given time.) Moreover, for the processing nodes which never forward served customers to other nodes, these random variables can be real rather than integer.

For utility nodes $n \in \mathcal{N}^u$, assumption that $b_n(k)$ are integers is not essential. (All “integer” assumptions are related to forwarding of “discrete customers” between nodes.) Also, $b_n(k)$ can be real random variables with finite means, *as long as utility function $H(u_1, \dots, u_{N_u})$ is well defined on the convex hull of all possible realizations of vectors $(b_1(k), \dots, b_{N_u}(k))$ for all k .* (See definition of set B at the beginning of Section 4.5.) The latter condition is often true “automatically,” for example when H is defined on the entire space R^{N_u} .

Adjustments of our results and proofs to accommodate the above generalizations are straightforward.

4.3 Equivalent Network Model

To simplify the formal problem statement (and simplify notations throughout analysis), it will be convenient to consider a slightly enhanced model, *which is equivalent to the model of Section 4.1*, and is such that both utility and processing nodes are treated in a “more

unified” way. The enhancement is as follows. First, for each utility node $n \in \mathcal{N}^u$ we denote $\mu_n(k) \doteq -b_n(k)$. Without loss of generality, we will assume that the values of $\mu_n(k)$ for all utility nodes $n \in \mathcal{N}^u$ and all controls $k \in K$ are strictly positive and uniformly bounded away from 0:

$$\min_{n \in \mathcal{N}^u, k \in K} \mu_n(k) \geq c_* > 0. \quad (49)$$

Indeed, suppose an arbitrary constant c is fixed, and we replace $b_n(k)$ with $b_n(k) - c$ for all $n \in \mathcal{N}^u$ and all k , and replace the utility function $H(u_1, \dots, u_{N_u})$ with $H(u_1 + c, \dots, u_{N_u} + c)$. Then, the new model is equivalent to the original one, *as long as a network control algorithm we are going to consider is invariant with respect to this change of variables*. The latter condition is easily seen to hold for the GPD control algorithm, which we introduce in Section 4.6. Finally, since the set of $b_n(k)$ for all $n \in \mathcal{N}^u$ and all k is uniformly bounded, we can choose $c > 0$ large enough so that (49) holds.

We also set by convention $\lambda_n(k) = 0$ for all utility nodes $n \in \mathcal{N}^u$ and all $k \in K$. Second, we introduce a queue (of infinite capacity) associated with each utility node n , and assume that (upon control k at time t) $Q_n(t) \wedge \mu_n(k)$ customers are served from this queue and leave the system (without being routed to any other node), and no customers ever arrive in the queue. In other words, the same rule (48) of the queue size $Q_n(t)$ dynamics applies to utility nodes $n \in \mathcal{N}^u$ (as well as for the processing nodes), but with both $A_n(t) = 0$ and $\lambda_n(t) = 0$ for $n \in \mathcal{N}^u$, any time t and any $k \in K$. (Clearly, the queue length process $Q_n(t), t \geq 0$, for any utility node has trivial behavior: it decreases at the rate $-c_*$ or less until it hits 0, and then stays at 0; in particular, if $Q_n(0) = 0$ then $Q_n(t) = 0$ for all $t \geq 0$, and so this queue is trivially stable.) Also, we extend the definition of routing matrices to $P(k) = \{p_{nj}(k), n, j \in \mathcal{N}\}$, using a natural convention that $p_{nj}(k) = 0$ when either $n \in \mathcal{N}^u$ or $j \in \mathcal{N}^u$. From now on in this section 4 we consider the enhanced version of the model.

Further, without loss of generality, we assume the following construction governing random routing of customers after their service at processing nodes. Suppose that, for every time $t = 0, 1, 2, \dots$, every possible routing matrix $\tilde{P} = \{\tilde{p}_{nj}(k), n, j \in \mathcal{N}\}$ (i.e., such that $\tilde{P} = P(k)$ for at least one $k \in K$), and every processing node $n \in \mathcal{N}^p$, there is an i.i.d. sequence of random variables $\xi_1(t, \tilde{P}, n), \xi_2(t, \tilde{P}, n), \dots$. This sequence determines where the “first,” “second,” and so on, customer leaving processing node n is going after being served by node n at time t , given the control k chosen at time t was such that $P(k) = \tilde{P}$. Thus, each random variable $\xi_i(t, \tilde{P}, n)$ takes values $j \in \mathcal{N}^p$ with probabilities \tilde{p}_{nj} , and value, say, 0 (indicating that a customer leaves the network) with probability $1 - \sum_{j \in \mathcal{N}^p} \tilde{p}_{nj}$. The sequences corresponding to different triples (t, \tilde{P}, n) are mutually independent.

4.4 System Rate Region

In this section we define the system rate region $V \subset R^N$. The meaning of the rate region is simple, and can be informally described as follows. Suppose, all initial queue sizes $Q_n(0), n \in \mathcal{N}$, are “infinitely large,” so that they “never” hit 0. Then, elements $v \in V$ represent all

possible vectors of long-term average drifts of the processes $Q_n(t)$, which can be induced by different control policies. We now proceed with the formal definitions.

For each control $k \in K$ and time $t = 0, 1, 2, \dots$, consider random vector $b(k, t) = (b_1(k, t), \dots, b_N(k, t))$, where each component $b_n(k, t)$ is equal to the (random) increment $Q_n(t+1) - Q_n(t) = -\mu_n(k) + A_n(t) + \lambda_n(k)$ of node n queue length, *assuming that k is the control chosen at time t and $Q_j(t) \geq \mu_j(k)$ for all $j \in \mathcal{N}$* . (See (48).) We call $b_n(k, t)$ the *nominal increment* of queue n upon control k at time t . Note that, for utility nodes $n \in \mathcal{N}^u$, $b_n(k, t) = -\mu_n(k) = b_n(k)$ for all t .

Random vectors $b(k, t)$ are well defined by our construction governing routing. (See Section 4.3.) The components of each $b(k, t)$ are in general mutually dependent. Moreover, if controls k and k' have equal routing matrix $P(k) = P(k')$, random vectors $b(k, t)$ and $b(k', t)$ are coupled, again via the routing process construction. However, the families of vectors $\{b(k, t), k \in K\}$ corresponding to different times t are mutually independent and identically distributed.

For each $k \in K$, let us denote $\bar{b}(k) = (\bar{b}_1(k), \dots, \bar{b}_N(k)) \doteq Eb(k, t)$ (this expectation does not depend on t); its n -th component

$$\bar{b}_n(k) = \lambda_n(k) - \mu_n(k) + \sum_{j \in \mathcal{N}} \mu_j(k) p_{jn}(k), \quad n \in \mathcal{N}, \quad (50)$$

will be called the *drift* of queue n upon control k . Using vector notations $\mu(k) = (\mu_1(k), \dots, \mu_N(k))$ and $\lambda(k) = (\lambda_1(k), \dots, \lambda_N(k))$, we can rewrite (50) as

$$\bar{b}(k) = \lambda(k) - \mu(k)(I - P(k)).$$

For utility nodes $n \in \mathcal{N}^u$ we have simply $\bar{b}_n(k) = b_n(k) = -\mu_n(k) \leq -c_* < 0$.

Suppose, for each network mode $m \in M$, a probability distribution $\phi_m = (\phi_{mk}, k \in K(m))$ is fixed, which means that $\phi_{mk} \geq 0$ for all $k \in K(m)$, and $\sum_{k \in K(m)} \phi_{mk} = 1$. For such a set of distributions $\phi \doteq (\phi_m, m \in M)$, consider the following vector

$$v(\phi) = \sum_{m \in M} \pi_m \sum_{k \in K(m)} \phi_{mk} \bar{b}(k).$$

If we interpret ϕ_{mk} as the long-term average fraction of time slots when control $k \in K(m)$ is chosen among the slots when the network mode is m , then $v(\phi)$ is the corresponding vector of long-term average drifts of the queue lengths $Q_n(t)$ (assuming queue lengths never hit 0).

The system *rate region* V is defined as the set of all possible vectors $v(\phi)$ corresponding to all possible ϕ . Clearly, V is a convex compact (in fact - polyhedral) subset of R^N , as a linear image of the compact polyhedral set of all possible values of ϕ . Rate region V may turn out to be degenerate (i.e., have dimension less than N).

Note that for any $v \in V$ and $n \in \mathcal{N}^u$ we automatically have $v_n \leq -c_* < 0$.

Also, for future reference, note that the following equality holds for any fixed vector $\zeta \in R^N$:

$$\max_{v \in V} \zeta \cdot v = \sum_{m \in M} \pi_m \max_{k \in K(m)} \zeta \cdot \bar{b}(k). \quad (51)$$

4.5 Underlying Optimization Problem

Let us denote by $B \subseteq R^{N_u}$ the convex hull of the set $\{(b_1(k), \dots, b_{N_u}(k)), k \in K\}$. Suppose a set $\tilde{V} = \tilde{B} \times R^{N_p} \subseteq R^N$ is fixed, where $\tilde{B} \subseteq R^{N_u}$ is an open convex set containing B . Clearly, $V \subset \tilde{V}$.

Suppose a continuously differentiable concave *utility function* $H(v)$ is defined on \tilde{V} , and it is such that $H(v)$ depends only on (v_1, \dots, v_{N_u}) , and *not* on (v_{N_u+1}, \dots, v_N) . Since, in essence, H is a function of $(v_1, \dots, v_{N_u}) \in \tilde{B}$, we sometimes write $H(v_1, \dots, v_{N_u})$ to mean $H(v_1, \dots, v_{N_u}, v_{N_u+1}, \dots, v_N)$ with arbitrary $v_n \in R$, $n \in \mathcal{N}^p$.

Consider the following optimization problem:

$$\max_{v \in V} H(v) \quad (52)$$

subject to

$$v \in R_-^N. \quad (53)$$

Problem (52)-(53) is feasible when

$$V \cap R_-^N \neq \emptyset, \quad (54)$$

in which case we denote by V^* the compact convex set of optimal solutions of (52)-(53), and by $Q^* \subseteq R_+^N$ the closed convex set of optimal solutions to the problem dual to (52)-(53).

For a given network control algorithm, let (u_1, \dots, u_{N_u}) denote the vector of long-run average values of $b_n(k(t)) = b_n(k(t), t)$ for utility nodes (that is, average rates at which commodities are generated by utility nodes). We seek to find a dynamic control algorithm such that, when problem (52)-(53) is feasible, utility $H(u_1, \dots, u_{N_u})$ is equal to the optimal value of the problem (52)-(53) *and* the queues at the processing nodes remain stable. *Such an algorithm would solve the problem of maximizing $H(u_1, \dots, u_{N_u})$, subject to the constraint that queues at the processing nodes remain stable*, and, in essence, “solve” problem (52)-(53). Indeed, if a control policy is able to keep processing node queues stable, then this policy is easily mapped into an equivalent policy (using “virtual controls” introduced later in the paper) which produces exactly same u_n for utility nodes, and produces non-positive (in fact - zero) long-term average values of $b_n(k(t), t)$, also denoted by u_n , for all processing nodes. This means that the resulting vector $u = (u_1, \dots, u_N)$ is a feasible solution of problem (52)-(53) and, consequently, $H(u)$ cannot exceed the optimal value of the problem. Thus, a dynamic control that we seek, would “produce” $u = (u_1, \dots, u_N) \in V^*$, i.e., an optimal solution of (52)-(53).

In the next section 4.6 we introduce an algorithm (called GPD algorithm), which is (asymptotically) optimal in the sense that it (asymptotically) achieves the goal described above, under the following non-degeneracy assumption, which is slightly stronger than feasibility condition (54):

$$V \cap R_{--}^N \neq \emptyset. \quad (55)$$

Since we automatically have $v_n < 0$, $n \in \mathcal{N}^u$, for all $v \in V$, assumption (55) means that there exists a control strategy which is able to provide *strictly negative* average drift to all processing node queues. (Under (55), Q^* is a compact set, as well as V^* . See Section 3.1.)

Remark. Non-degeneracy assumption (55), or even a weaker feasibility assumption (54), are *not* needed for *any* of the results of this Section 4, which are concerned with network dynamics under GPD algorithm. Assumption (55) is only invoked to apply Theorem 2 (which says that the dynamic system in fact converges to an optimal state), and thus establish asymptotic optimality of the GPD algorithm. (See also the discussion at the beginning of Section 4.8.)

4.6 Greedy Primal-Dual Algorithm

Consider the following control policy. (Recall that $H(\cdot)$ is the utility function defined in Section 4.5.)

Greedy Primal-Dual (GPD) algorithm. *At time t choose a control*

$$k(t) \in \arg \max_{k \in K(m(t))} \sum_{n \in \mathcal{N}^u} (\partial H(X(t))/\partial x_n) \bar{b}_n(k) - \sum_{n \in \mathcal{N}^p} \beta Q_n(t) \bar{b}_n(k), \quad (56)$$

where running averages $X_n(t)$ of the values of $b_n(k(t))$ for utility nodes are updated as follows:

$$X_n(t+1) = (1 - \beta)X_n(t) + \beta b_n(k(t)), \quad n \in \mathcal{N}^u, \quad (57)$$

with $\beta > 0$ being a (small) parameter, and where the queue lengths $Q_n(t)$, $n \in \mathcal{N}^p$, for the processing nodes are updated according to (48). The initial values $(X_1(0), \dots, X_{N_u}(0)) \in \tilde{B}$ and $Q_n(0) \geq 0$, $n \in \mathcal{N}^p$, are fixed.

Remark. The initial condition $(X_1(0), \dots, X_{N_u}(0)) \in \tilde{B}$ and the update rule (57) imply (by induction) that $(X_1(t), \dots, X_{N_u}(t)) \in \tilde{B}$ for all $t \geq 0$. Therefore, the (random) system evolution is well defined for all $t \geq 0$, because all partial derivatives and the arg max in (56) are well defined.

It is very important to emphasize that the GPD algorithm defined above uses and maintains variables $X_n(t)$ for utility nodes only and queue lengths $Q_n(t)$ for processing nodes only.

4.7 Equivalent Form of GPD Algorithm

Solely for the purposes of the GPD algorithm analysis, we will now introduce some definitions and conventions, which will allow us to treat both utility and processing nodes in a unified

way, and thus bring the definition of GPD algorithm to an equivalent form convenient for the analysis. We wish to emphasize that these definitions and conventions (as well as those in Section 4.3) are *not* utilized in any way (neither explicit nor implicit) by the GPD algorithm itself, as it is defined in Section 4.6.

First, let us consider an extended set of controls. Let us denote by $k' \preceq k$ the following partial order relation between controls k' and k (which may or may not be within K): $0 \leq \mu_n(k') \leq \mu_n(k)$, $\lambda_n(k') = \lambda_n(k)$, $p_{nj}(k') = p_{nj}(k)$ for all $n, j \in \mathcal{N}^p$, and $b_n(k') = b_n(k)$ for $n \in \mathcal{N}^u$. Simply put, $k' \preceq k$ means that k' is identical to k in every respect, except it serves smaller or equal number of customers from each processing queue n . For each mode m , consider the extended set of controls $\hat{K}(m) \doteq \{k' \mid k' \preceq k \in K(m)\}$, and denote $\hat{K} = \cup_m \hat{K}(m)$. Clearly, $K(m) \subseteq \hat{K}(m)$ for each m , and $K \subseteq \hat{K}$. We extend the definitions of (random) nominal increment vectors $b(k, t)$ and drift vectors $\bar{b}(k)$ on all $k \in \hat{K}$. Observe that if we were to define system rate region as in Section 4.4, but via extended control sets $\hat{K}(m)$, we would obtain the same region V . In addition, (51) can be strengthened as follows: for any fixed vector $\zeta \in R^N$,

$$\max_{v \in V} \zeta \cdot v = \sum_{m \in M} \pi_m \max_{k \in K(m)} \zeta \cdot \bar{b}(k) = \sum_{m \in M} \pi_m \max_{k \in \hat{K}(m)} \zeta \cdot \bar{b}(k). \quad (58)$$

Note also that replacing K with \hat{K} does not change the convex hull B defined in Section 4.5.

Now we will introduce the notion of a virtual control at time t . Suppose the *actual control* (that is, the control actually chosen by a control algorithm) at time t is $k = k(t) \in K(m(t))$. Then, the *virtual control* $\hat{k} = \hat{k}(t) \in \hat{K}(m)$ at time t is the control defined as follows: $\hat{k} \preceq k$ and

$$\mu_n(\hat{k}) = \mu_n(k) \wedge Q_n(t), \quad n \in \mathcal{N}^p.$$

Simply put, the only difference between \hat{k} and k is that, for those processing nodes n with $Q_n(t) < \mu_n(k)$, $\mu_n(k)$ is “replaced” with $\mu_n(\hat{k})$ equal to the actual number of customers $Q_n(t)$ which are served in slot t .

It is easy to see that if we were to replace actual control $k(t)$ at time t with the corresponding virtual control $\hat{k}(t)$, this will have *no effect* on the way $X_n(t)$ for utility nodes and queue lengths $Q_n(t)$ for processing nodes are updated at time t .

We extend the definition of nominal increments $b_n(k, t)$, given at the beginning of Section 4.4, for all $k \in \hat{K}$. The key reason for introducing the notion of virtual control is that, for processing nodes, we always have:

$$Q_n(t+1) - Q_n(t) = -\mu_n(\hat{k}(t)) + A_n(t) + \lambda_n(\hat{k}(t)) = b_n(\hat{k}(t), t), \quad n \in \mathcal{N}^p, \quad (59)$$

that is, for the virtual control $\hat{k}(t)$ the actual queue increment for each processing node is equal to the corresponding nominal increment.

We already assumed (as part of the enhanced model formulation) that queues $Q_n(t)$ are maintained for utility nodes (as well as for processing nodes); let us further assume that for utility nodes $Q_n(0) = 0$, and therefore $Q_n(t) \equiv 0$, $t \geq 0$, for $n \in \mathcal{N}^u$. Using (59) and the

fact that for utility nodes $-\mu_n(k(t)) = b_n(k(t)) = b_n(\hat{k}(t), t) \leq -c_* < 0$, the unified queue lengths update rule (48), which holds for all $n \in \mathcal{N}$, can be equivalently rewritten as

$$Q_n(t+1) = [Q_n(t) + b_n(\hat{k}(t), t)]^+, \quad n \in \mathcal{N}. \quad (60)$$

We will denote $Q(t) = (Q_1(t), \dots, Q_N(t))$.

We also assume that variables $X_n(t)$ are maintained for processing nodes $n \in \mathcal{N}^p$ (as well as for utility nodes), with arbitrary initial values $X_n(0) \in R$, and with the update rule

$$X_n(t+1) = (1 - \beta)X_n(t) + \beta b_n(\hat{k}(t), t), \quad n \in \mathcal{N}^p. \quad (61)$$

Using the fact that for utility nodes $b_n(k(t)) = b_n(\hat{k}(t), t)$, we can combine (57) and (61) into a “unified” update rule for the vector $X(t) = (X_1(t), \dots, X_N(t))$ under GPD algorithm:

$$X(t+1) = (1 - \beta)X(t) + \beta b(\hat{k}(t), t). \quad (62)$$

Using the conventions and notations introduced above, we see that the GPD algorithm can be equivalently defined as follows. (The equivalence is in the sense that both produce the same behavior of $X_n(t)$, $n \in \mathcal{N}^u$, and $Q_n(t)$, $n \in \mathcal{N}^p$.)

GPD algorithm (equivalent form): *At time t choose a control*

$$k(t) \in \arg \max_{k \in K(m(t))} [\nabla H(X(t)) - \beta Q(t)] \cdot \bar{b}(k), \quad (63)$$

where $\beta > 0$ is a (small) parameter, $X(t)$ is updated according to (62) with $X(0) \in \tilde{V}$, and $Q(t)$ is updated according to (60) (or, equivalently, (48)) with $Q_n(0) \geq 0$ for $n \in \mathcal{N}^p$ and $Q_n(0) = 0$ for $n \in \mathcal{N}^u$.

Random process $S = \{S(t), t = 0, 1, \dots\}$, where $S(t) \doteq (X(t), Q(t), m(t))$, describes evolution of the network. Given our model assumptions, S is a discrete time Markov chain with state space $R^N \times Z_+^N \times M$, where Z_+ is the set of non-negative integers.

4.8 Asymptotic Optimality of GPD Algorithm

The main result of this section (Theorem 3) is that, as $\beta \downarrow 0$, the “fluid-scaled” processes $\{(X(t/\beta), t \geq 0), (\beta Q(t/\beta), t \geq 0)\}$ converge to a random process $\{(x(t), t \geq 0), (q(t), t \geq 0)\}$ with sample paths being GPD-trajectories studied in Section 3, that is, trajectories satisfying conditions (10), (11), (13), (14), (16). (This result, as well as all other results of Section 4, does *not* use assumption (55), or even (54).) But, according to Theorem 2, under the non-degeneracy assumption (55), for all GPD-trajectories, as time $t \rightarrow \infty$, we have $(x(t), q(t)) \rightarrow V^* \times q^*$ for some fixed $q^* \in Q^*$, where V^* and Q^* are the sets of optimal solutions of the problem (52)-(53) and its dual, respectively. In this sense, Theorems 3 and 2 demonstrate *asymptotic optimality* of the GPD algorithm.

4.8.1 Asymptotic Regime. Fluid Scaled Processes.

First, we need to define the asymptotic regime formally. From this point on in the paper, we consider a sequence of processes $S^\beta = (X^\beta, Q^\beta, m^\beta)$, indexed by the value of parameter β , with $\beta \downarrow 0$ along a sequence $\mathcal{B} = \{\beta_j, j = 1, 2, \dots\}$ such that $\beta_j > 0$ for all j . The initial state $S^\beta(0) = (X^\beta(0), Q^\beta(0), m^\beta(0))$ is fixed for each $\beta \in \mathcal{B}$, and it satisfies the conditions specified in the GPD algorithm definition in Section 4.6. (Here and below, the processes and variables pertaining to a fixed parameter β will be supplied the upper index β . Expression $\beta \downarrow 0$ means that β converges to 0 along the sequence \mathcal{B} , unless otherwise specified.)

The probability law of the Markov chain $m^\beta(\cdot)$ describing the system mode process is same for each β .

Before we introduce fluid-scaled version of the process (for each $\beta \in \mathcal{B}$), we need to augment the definition of the process itself. First, let us extend the definition of $X^\beta(t)$ to continuous time $t \in R_+$ by adopting the convention that $X^\beta(t)$ is constant within each time slot $[l, l+1)$. We do analogous domain extension for $Q^\beta(t)$. For each β let us define some additional random functions, associated with the system evolution. We define them for continuous time $t \in R_+$ as well, although they (just as $X^\beta(t)$ and $Q^\beta(t)$) are constant within each time slot $[l, l+1)$. Let us denote

$$\bar{F}_n^\beta(t) \doteq Q_n^\beta(0) + \sum_{l=1}^{\lfloor t \rfloor} b_n^\beta(k(l-1), l-1), \quad t \geq 0, \quad n \in \mathcal{N}, \quad (64)$$

$$F_n^\beta(t) \doteq Q_n^\beta(0) + \sum_{l=1}^{\lfloor t \rfloor} b_n^\beta(\hat{k}(l-1), l-1), \quad t \geq 0, \quad n \in \mathcal{N}. \quad (65)$$

Note that the difference between the definitions of $\bar{F}_n^\beta(t)$ and $F_n^\beta(t)$ is that in (64) we sum up the values of b_n^β corresponding to actual controls chosen in the interval $[0, t-1]$, while in (65) we sum up b_n^β for the corresponding virtual controls.

For future reference, we record the following simple facts. For the processing nodes n we have (59), and therefore

$$F_n^\beta(t) \equiv Q_n^\beta(t), \quad t \geq 0, \quad n \in \mathcal{N}^p. \quad (66)$$

For utility nodes

$$Q_n^\beta(t) \equiv 0, \quad t \geq 0, \quad n \in \mathcal{N}^u, \quad (67)$$

and

$$F_n^\beta(l) - F_n^\beta(l-1) = b_n^\beta(\hat{k}(l-1), l-1) = b_n^\beta(k(l-1), l-1) \leq -c_* < 0, \quad l = 1, 2, \dots, \quad n \in \mathcal{N}^u. \quad (68)$$

Denote by $G_m^\beta(t)$ the total number of time slots by (and including) time $t-1$, when the system mode was m . Denote by $\bar{G}_{mk}^\beta(t)$ (respectively, $G_{mk}^\beta(t)$) the number of time slots by (and including) time $t-1$ when the mode was m and the actual (respectively, virtual) control was $k \in \hat{K}(m)$. (Since actual controls are chosen from $K(m)$, $\bar{G}_{mk}^\beta(t) \equiv 0$ if $k \in \hat{K}(m) \setminus K(m)$.)

Clearly, for any $n \in \mathcal{N}$, $m \in M$, and $k \in \hat{K}(m)$, we have:

$$\bar{F}_n^\beta(0) = F_n^\beta(0) = Q_n^\beta(0), \quad G_m^\beta(0) = 0, \quad \bar{G}_{mk}^\beta(0) = G_{mk}^\beta(0) = 0,$$

and we have the following relations:

$$G_m^\beta(t) = \sum_{k \in \hat{K}(m)} \bar{G}_{mk}^\beta(t), \quad t \geq 0,$$

$$G_m^\beta(t) = \sum_{k \in \hat{K}(m)} G_{mk}^\beta(t), \quad t \geq 0.$$

We define process Z^β , describing system evolution, as

$$Z^\beta = (X^\beta, Q^\beta, F^\beta, \bar{F}^\beta, G^\beta, \bar{G}^\beta),$$

where

$$\begin{aligned} X^\beta &= (X^\beta(t) = (X_1^\beta(t), \dots, X_N^\beta(t)), \quad t \geq 0), \\ Q^\beta &= (Q^\beta(t) = (Q_1^\beta(t), \dots, Q_N^\beta(t)), \quad t \geq 0), \\ F^\beta &= (F^\beta(t) = (F_1^\beta(t), \dots, F_N^\beta(t)), \quad t \geq 0), \\ \bar{F}^\beta &= (\bar{F}^\beta(t) = (\bar{F}_1^\beta(t), \dots, \bar{F}_N^\beta(t)), \quad t \geq 0), \\ G^\beta &= ((G_m^\beta(t), m \in M), t \geq 0), (G_{mk}^\beta(t), k \in \hat{K}(m), m \in M), t \geq 0)), \\ \bar{G}^\beta &= ((\bar{G}_{mk}^\beta(t), m \in M, k \in \hat{K}(m)), t \geq 0). \end{aligned}$$

For each β consider the following process z^β , which is a *fluid-scaled* version of process Z^β :

$$z^\beta = (x^\beta, q^\beta, f^\beta, \bar{f}^\beta, g^\beta, \bar{g}^\beta),$$

where x^β is obtained by time scaling only:

$$x^\beta(t) \doteq X^\beta(t/\beta), \tag{69}$$

and the other components by time and space scaling:

$$q^\beta(t) \doteq \beta Q^\beta(t/\beta), \quad f^\beta(t) \doteq \beta F^\beta(t/\beta), \quad \bar{f}^\beta(t) \doteq \beta \bar{F}^\beta(t/\beta), \tag{70}$$

$$g^\beta(t) \doteq \beta G^\beta(t/\beta), \quad \bar{g}^\beta(t) \doteq \beta \bar{G}^\beta(t/\beta). \tag{71}$$

Note that the component functions of z^β are piece-wise constant, with the ‘‘time slot’’ of length β .

4.8.2 Fluid Sample Paths under GPD Algorithm

We now define *fluid sample paths*, which are fixed trajectories arising as possible limits of sequences (on β) of realizations of z^β , given those realizations satisfy some functional law of large numbers (FLLN) type conditions.

Definition. A fixed set of functions $z = (x, q, f, \bar{f}, g, \bar{g})$ we will call a fluid sample path (FSP) if there exists a sequence \mathcal{B}_0 of positive values of β , such that $\beta \downarrow 0$, and a sequence of sample paths (of the corresponding processes) $\{z^\beta\}$ such that (as $\beta \downarrow 0$ along sequence \mathcal{B}_0)

$$z^\beta \rightarrow z, \quad u.o.c. \quad ,$$

and in addition

$$\|x(0)\| \in \tilde{V} \quad , \quad \|q(0)\| < \infty \quad ,$$

$$(g_m^\beta(t), t \geq 0) \rightarrow (\pi_m t, t \geq 0) \quad u.o.c. \quad , \quad \forall m \in M, \quad (72)$$

$$(\bar{f}_n^\beta(t), t \geq 0) \rightarrow (q_n(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) \bar{g}_{mk}(t), t \geq 0) \quad u.o.c. \quad , \quad \forall n \in \mathcal{N}, \quad (73)$$

$$(f_n^\beta(t), t \geq 0) \rightarrow (q_n(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) g_{mk}(t), t \geq 0) \quad u.o.c. \quad , \quad \forall n \in \mathcal{N}. \quad (74)$$

Remark. A sequence \mathcal{B}_0 whose existence is required in the above definition may be completely unrelated to the sequence \mathcal{B} we introduced earlier. For an FSP z , any sequence $\{z^\beta\}$ of sample paths satisfying conditions of the above definition, will be called a sequence defining this FSP z .

Lemma 17 For any fluid sample path z , all its component functions are Lipschitz continuous in $[0, \infty)$, with the Lipschitz constant $C + \|x(0)\|$, where $C > 0$ is a fixed constant depending only on the system parameters. In addition, functions $f_n(\cdot)$, $\bar{f}_n(\cdot)$, $g_m(\cdot)$, $g_{mk}(\cdot)$, $\bar{g}_{mk}(\cdot)$, are non-decreasing, satisfying the following relations:

$$g_m(t) = \pi_m t \quad , \quad t \geq 0, \quad m \in M \quad , \quad (75)$$

$$g_m(t) = \sum_{k \in \hat{K}(m)} g_{mk}(t), \quad t \geq 0, \quad m \in M \quad , \quad (76)$$

$$g_m(t) = \sum_{k \in \hat{K}(m)} \bar{g}_{mk}(t), \quad t \geq 0, \quad m \in M \quad , \quad (77)$$

$$\bar{f}(t) = q(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}(k) \bar{g}_{mk}(t), \quad t \geq 0, \quad (78)$$

$$f(t) = q(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}(k) g_{mk}(t), \quad t \geq 0. \quad (79)$$

Proof is in Section 4.8.4.

Lemma 18 *Any FSP z satisfies the following additional properties:*

$$f_n(t) \equiv q_n(t), \quad t \geq 0, \quad n \in \mathcal{N}^p, \quad (80)$$

$$q_n(t) \equiv 0, \quad t \geq 0, \quad n \in \mathcal{N}^u, \quad (81)$$

$$f'_n(t) \leq -c_* < 0 \text{ for almost all } t \geq 0, \quad n \in \mathcal{N}^u. \quad (82)$$

Proof is in Section 4.8.4.

Lemma 19 *Any FSP z satisfies the following additional properties.*

(i) *For all $t \geq 0$*

$$x(t) \in \tilde{V}, \quad (83)$$

and for almost all $t \geq 0$ we have:

$$x'(t) = v(t) - x(t), \quad (84)$$

where

$$v(t) \doteq (d/dt)f(t) \in V \quad (85)$$

satisfies the condition

$$v(t) \in \arg \max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot v. \quad (86)$$

(ii) *We have*

$$q(0) \geq 0, \quad (87)$$

$$f(0) = q(0) \quad \text{and} \quad q_n(t) = f_n(t) - [0 \wedge \inf_{0 \leq \xi \leq t} f_n(\xi)], \quad t \geq 0, \quad n \in \mathcal{N}. \quad (88)$$

Proof is in Section 4.8.4.

4.8.3 Fluid Scaled Processes Converge to Processes Concentrated on Fluid Sample Paths

Let $N_* = 4N + \sum_{m \in M} (1 + 2|\hat{K}(m)|)$, where $|\hat{K}(m)|$ is cardinality of set $\hat{K}(m)$, denote the total number of all scalar component functions comprising $z^\beta = (x^\beta, q^\beta, f^\beta, \bar{f}^\beta, g^\beta, \bar{g}^\beta)$. We will view *random processes* z^β as processes with realizations in the Skorohod space $D_{R^{N_*}}[0, \infty)$ of functions with domain $[0, \infty)$, taking values in R^{N_*} , which are right-continuous and have left-limits. The Skorohod topology and corresponding Borel σ -algebra on $D_{R^{N_*}}[0, \infty)$ are defined in the usual way. (Cf. [5] for the definitions.)

Theorem 3 *Consider the sequence of processes $\{z^\beta\}$ with $\beta \downarrow 0$ along set \mathcal{B} . Assume that $z^\beta(0) \rightarrow z(0)$, where $z(0) \in R^{N_*}$ is a fixed vector such that $(x(0), q(0)) \in \tilde{V} \times R_+^N$ and $q_n(0) = 0$, $n \in \mathcal{N}^u$. Then, the sequence $\{z^\beta\}$ is relatively compact and any weak limit of this sequence (i.e., a process obtained as a weak limit of a subsequence of $\{z^\beta\}$) is a process with sample paths being FSPs (with initial state $z(0)$) with probability 1.*

Theorem 3 and its proof (in Section 4.8.5) are analogous to the process convergence result in [21].

4.8.4 Proofs of Lemmas 17-19

Proof of Lemma 17. Consider a fixed FSP z and a sequence of paths z^β defining it. For each β , let Z^β be the corresponding unscaled path from which z^β is obtained.

To prove Lipschitz property of the components of z , recall that by (62) we have

$$X_n^\beta(l) = \beta b_n^\beta(\hat{k}(l-1), l-1) + (1-\beta)X_n^\beta(l-1). \quad (89)$$

Applying (89) iteratively for $l = 1, 2, \dots$, we obtain:

$$X_n^\beta(l) = \tilde{X}_n^\beta(l) + (1-\beta)^l X_n^\beta(0), \quad (90)$$

where

$$\tilde{X}_n^\beta(l) \doteq \sum_{j=0}^{l-1} \beta(1-\beta)^j b_n^\beta(\hat{k}(l-1-j), l-1-j).$$

We see that for any $l \geq 1$

$$\tilde{X}_n^\beta(l) \leq c^*, \quad (91)$$

where constant $c^* > 0$ is a uniform upper bound of the (random) values of $|b_n^\beta(k, l)|$ across all n and k (and l). Indeed, $\tilde{X}_n^\beta(l)$ is a weighted sum of the values of $b_n^\beta(\hat{k}(0), 0), \dots, b_n^\beta(\hat{k}(l-1), l-1)$ with positive weights summing up to at most 1.

Let us rewrite (89) as follows:

$$X_n^\beta(l) - X_n^\beta(l-1) = \beta b_n^\beta(\hat{k}(l-1), l-1) - \beta X_n^\beta(l-1). \quad (92)$$

We see that the one time slot increments of $X_n^\beta(\cdot)$ are uniformly bounded:

$$|X_n^\beta(l) - X_n^\beta(l-1)| \leq \beta(|b_n^\beta(\hat{k}(l-1), l-1)| + |\tilde{X}_n^\beta(l-1)| + |X_n^\beta(0)|) \leq \beta(2c^* + X_n^\beta(0)). \quad (93)$$

It is trivial to see that all other component functions of Z^β also have uniformly bounded one slot increments. Therefore, recalling the rescaling formulas (69)-(71), the Lipschitz continuity of all components of FSP z , with the specified Lipschitz constant (if we choose C large enough), easily follows.

The fact that the components $g_m(\cdot)$, $g_{mk}(\cdot)$, $\bar{g}_{mk}(\cdot)$, $f_n(\cdot)$, $\bar{f}_n(\cdot)$, are non-decreasing, and relations (75)-(79) follow directly from the definitions involved. ■

Proof of Lemma 18. Relations (80)-(82) immediately follow from the corresponding pre-limit relations (66)-(68), the definition of an FSP, and Lipschitz continuity of the components of z . ■

Proof of Lemma 19. As in the proof of Lemma 17, consider a fixed FSP z and a sequence of paths z^β defining it, along with their unscaled versions Z^β .

Update rule (89) implies (by induction) that, for all $l = 0, 1, \dots$, vector $(X_1^\beta(l), \dots, X_{N_u}^\beta(l))$ lies within the (compact) convex hull of the set $B \cup \{(X_1^\beta(0), \dots, X_{N_u}^\beta(0))\}$. Since $(X_1^\beta(0), \dots, X_{N_u}^\beta(0)) \rightarrow (x_1(0), \dots, x_{N_u}(0))$ as $\beta \downarrow 0$, we easily obtain that the entire trajectory $(x_1(t), \dots, x_{N_u}(t))$, $t \geq 0$, lies within the (compact) convex hull of the set $B \cup \{(x_1(0), \dots, x_{N_u}(0))\}$. This, in particular, implies (83).

Since all component functions of an FSP are Lipschitz, they are absolutely continuous, and therefore almost all points $t > 0$ are such that all component functions of z have proper derivatives (that is, both right and left derivatives exist and are equal). Within the rest of this proof we will call such points $t > 0$ *regular*.

We now derive the basic integral equation for $x(\cdot)$ (it is equation (94) below), which will imply (84)-(85). Let us sum up equations (92) for $l = 1, \dots, j$. We obtain

$$X_n^\beta(j) - X_n^\beta(0) = \beta \sum_{l=1}^j b_n^\beta(\hat{k}(l-1), l-1) - \sum_{l=1}^j \beta X_n^\beta(l-1).$$

Switching to scaled processes, for all integer $j \geq 0$ we have

$$x_n^\beta(\beta j) - x_n^\beta(0) = f_n^\beta(\beta j) - q_n^\beta(0) - \int_0^{\beta j} x_n^\beta(\xi) d\xi.$$

Using the fact that all *limiting* functions x_n and f_n are Lipschitz continuous and we have u.o.c. convergence $z^\beta \rightarrow z$, we finally obtain the desired integral equation

$$x_n(t) - x_n(0) = f_n(t) - q_n(0) - \int_0^t x_n(\xi) d\xi, \quad t \in R_+. \quad (94)$$

Since both $x_n(\cdot)$ and $f_n(\cdot)$ are Lipschitz continuous, (84) and (85) hold for every regular point $t > 0$.

Let us prove (86). Consider a fixed regular point $t > 0$. It will suffice to prove that (86) holds for this t . To do that, we will first prove a similar property for $\bar{v}(t) \doteq (d/dt)\bar{f}(t) \in V$, namely,

$$\bar{v}(t) \in \arg \max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot v, \quad (95)$$

and then show that (95) implies (86).

The function $\nabla H(\eta) - \zeta$ of $(\eta, \zeta) \in \tilde{V} \times R_+^N$ is continuous and bounded in some open neighborhood of point $(x(t), q(t))$. Given this, the rest of the proof of (95) is analogous to the proof of Lemma 5(ii) in [20]. Namely, since $(x^\beta(\xi), q^\beta(\xi))$ is close to $(x(t), q(t))$ when $|\xi - t|$ and β are small, the following observation is true. (It follows from the form of GPD rule and Proposition 2 in Section 6.4.)

For any $\epsilon > 0$, there exists a sufficiently small $\Delta > 0$, such that for all unscaled paths Z^β with sufficiently small β , we have the following property. For any (positive integer) time slot l within interval $[t/\beta, (t + \Delta)/\beta]$,

$$|[\nabla H(X^\beta(l-1)) - \beta Q^\beta(l-1)] \cdot \bar{b}(k(l-1)) - c_{m(l-1)}| \leq \epsilon,$$

where we use notation

$$c_m \doteq \max_{k \in K(m)} [\nabla H(x(t)) - q(t)] \cdot \bar{b}(k), \quad m \in M.$$

(The values of c_m depend on $(x(t), q(t))$ of course. We do not indicate this dependence explicitly to simplify notation.)

From the above observation we have

$$\left| \sum_{t/\beta \leq l \leq (t+\Delta)/\beta} [\nabla H(X^\beta(l-1)) - \beta Q^\beta(l-1)] \cdot \bar{b}(k(l-1)) - \sum_{t/\beta \leq l \leq (t+\Delta)/\beta} c_{m(l-1)} \right| \leq \epsilon \Delta / \beta + O(1),$$

where $O(1)$ denotes a term with absolute value bounded above by some fixed constant $C > 0$ as $\beta \rightarrow 0$. From the last display, multiplied by β , it is easy to obtain the following estimate for the fluid-scaled paths:

$$\left| \int_t^{t+\Delta} [\nabla H(x^\beta(\xi)) - q^\beta(\xi)] \cdot d\bar{f}^\beta(\xi) - \sum_{m \in M} \int_t^{t+\Delta} c_m dg_m^\beta(\xi) \right| \leq \epsilon \Delta + \beta O(1).$$

Taking the limit on $\beta \rightarrow 0$ and using (58), we obtain

$$\left| \int_t^{t+\Delta} [\nabla H(x(\xi)) - q(\xi)] \cdot d\bar{f}(\xi) - \sum_{m \in M} c_m \pi_m \Delta \right|$$

$$= \left| \int_t^{t+\Delta} [\nabla H(x(\xi)) - q(\xi)] \cdot d\bar{f}(\xi) - (\max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot v) \Delta \right| \leq \epsilon \Delta .$$

Since Δ can be chosen arbitrarily small (for a given fixed ϵ) and $\nabla H(x(\xi)) - q(\xi)$ is continuous at $\xi = t$, we have

$$|[\nabla H(x(t)) - q(t)] \cdot \bar{v}(t) - \max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot v| \leq \epsilon .$$

Finally, since ϵ can be chosen arbitrarily small, $[\nabla H(x(t)) - q(t)] \cdot \bar{v}(t) = \max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot v$, which completes the proof of (95).

To show that (95) implies (86), we need to verify that

$$[\nabla H(x(t)) - q(t)] \cdot v(t) = [\nabla H(x(t)) - q(t)] \cdot \bar{v}(t) .$$

Recall that the structure of our model is such that:

$v_n(t) = \bar{v}_n(t)$ for all utility nodes $n \in \mathcal{N}^u$,

and $(\partial/\partial \xi_n)H(\xi) = 0$ for all processing nodes $n \in \mathcal{N}^p$. This means it will suffice to verify the following fact:

$$\text{If } q_n(t) > 0 \text{ for some } n \in \mathcal{N}^p, \text{ then } v_n(t) = \bar{v}_n(t). \quad (96)$$

To demonstrate (96), consider any $n \in \mathcal{N}^p$ for which $q_n(t) > 0$. Since $q^\beta(\xi)$ is close to $q(t)$ when $|\xi - t|$ and β are small, the following observation (which relies in part on the form of GPD rule) is true.

There exists a sufficiently small $\Delta > 0$, such that for all unscaled paths Z^β with sufficiently small β , and any (positive integer) time slot l within interval $[t/\beta, (t + \Delta)/\beta]$, we have the following properties:

- (a) $\mu_j(\hat{k}(l)) = \mu_j(k(l))$ for any processing node j (including node n) for which $q_j(t) > 0$;
- (b) if $q_j(t) = 0$ and $p_{jn}(k(l)) > 0$ for some processing node $j \neq n$, then $\mu_n(\hat{k}(l)) = \mu_n(k(l)) = 0$;
- (c) as a corollary of (a) and (b),

$$b_n^\beta(\hat{k}(l), l) = b_n^\beta(k(l), l).$$

The above observation easily implies (after switching to scaled paths and taking the $\beta \downarrow 0$ limit), that

$$f_n(\xi) - f_n(t) = \bar{f}_n(\xi) - \bar{f}_n(t), \quad \forall \xi \in [t, t + \Delta],$$

which implies $v_n(t) = \bar{v}_n(t)$. This proves (96), and thus completes the proof of (86) and statement (i) of the lemma.

Statement (ii). Inequality $q(0) \geq 0$ holds by the definition of an FSP, and $f(0) = q(0)$ follows from (79). Using these relations and properties (80)-(82), the second part of (88) is verified directly. ■

4.8.5 Proof of Theorem 3

Recall that the probability law of the Markov chain $m^\beta(\cdot)$ describing the network mode process is same for each β , and therefore, for each $m \in M$, sequence $\{g_m^\beta\}$ satisfies the law of large numbers (LLN), namely

$$g_m^\beta(t) \rightarrow \pi_m t \quad \text{in probability} \quad \forall t \geq 0. \quad (97)$$

The following asymptotic relations, for each $n \in \mathcal{N}$, also easily follow from the LLN:

$$|\bar{f}_n^\beta(t) - (q_n^\beta(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) \bar{g}_{mk}^\beta(t))| \rightarrow 0 \quad \text{in probability} \quad \forall t \geq 0, \quad (98)$$

$$|f_n^\beta(t) - (q_n^\beta(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) g_{mk}^\beta(t))| \rightarrow 0 \quad \text{in probability} \quad \forall t \geq 0. \quad (99)$$

Using (97) and the uniform bound (93) on the increments of (unscaled paths) $X_n^\beta(\cdot)$, it is easy to see that the sequence $\{z^\beta\}$ is *asymptotically Lipschitz*, namely, for some fixed $C > 0$,

$$P\{\|z^\beta(t_2) - z^\beta(t_1)\| \leq C(t_2 - t_1) + \epsilon\} \rightarrow 1, \quad \forall 0 \leq t_1 \leq t_2, \quad \forall \epsilon > 0. \quad (100)$$

The above asymptotic Lipschitz property and the fact that the initial states $z^\beta(0)$ remain uniformly bounded (in fact they converge), imply that the sequence of processes $\{z^\beta\}$ is tight. Therefore, by Prohorov theorem (cf. Theorem 3.2.2 in [5]), the sequence $\{z^\beta\}$ is relatively compact.

Consider any subsequence of indices β , along which $z^\beta \xrightarrow{w} z$, where $z = (x, q, f, \bar{f}, g, \bar{g})$ is a *random process*. It easily follows from the asymptotic Lipschitz property (100) (and the fact of the weak convergence) that, w.p.1, the sample paths of z are Lipschitz continuous. Moreover, using asymptotic properties (97), (98) and (99), it is easy to see that, w.p.1, sample paths of z are such that

$$\begin{aligned} g_m(t) &= \pi_m t, \quad t \geq 0, \quad \forall m \in M, \\ \bar{f}_n(t) &= q_n(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) \bar{g}_{mk}(t), \quad t \geq 0, \quad \forall n \in \mathcal{N}, \\ f_n(t) &= q_n(0) + \sum_{m \in M} \sum_{k \in \hat{K}(m)} \bar{b}_n(k) g_{mk}(t), \quad t \geq 0, \quad \forall n \in \mathcal{N}. \end{aligned}$$

Using the above properties of the sample paths of z and the definition of an FSP, we obtain the fact that sample paths of z are FSPs w.p.1 (which completes the proof of the lemma). This can be done, for example, analogously to the proof of Theorem 7.1 in [19]. Alternatively, we can use *Skorohod representation* (cf. [5]), which allows us to construct all the processes z^β (of the converging subsequence) and process z on a common probability space so that the convergence $z^\beta \rightarrow z$ is u.o.c. w.p.1, and then use the definition of an FSP. \blacksquare

4.9 Comments on Notions of Asymptotic Optimality and Related Issues

The notion of asymptotic optimality of GPD algorithm, as parameter $\beta \downarrow 0$, described at the beginning of Section 4.8 (let us label this notion (AO-1)), is not the only possible. Perhaps a more “conventional” asymptotic optimality statement would be as follows.

(AO-2) *Under GPD algorithm, for each parameter β , let u_n , $n \in \mathcal{N}^u$, be expected value of $b_n(k(t))$, when Markov chain S is in stationary regime. Then, as $\beta \downarrow 0$, $H(u_1, \dots, u_{N_u})$ converges to the optimal value of (52)-(53).*

We believe that the above property (AO-2) in fact holds, and can be proved following steps similar to those in [21], where asymptotic optimality of the Gradient algorithm was established in both forms (AO-1) and (AO-2), with the latter derived from the former. We do not pursue the “derivation” of (AO-2) from (AO-1) here to save space and because the focus of this paper is on dynamic properties of the network under GPD algorithm. (The fact that convergence of GPD-trajectories to the optimal set is uniform, as in Theorem 2(ii), is used in such derivation.)

We note that, under condition (55), stochastic stability of the network (roughly, the existence of a stationary regime) under GPD algorithm *with any fixed parameter β* can be established the same way as that for MaxWeight type algorithms. (Cf. [22, 23, 20, 4] and references therein.) This is because, given our model assumptions (with, essentially, $X(t)$ being confined to a compact set), when $\|Q(t)\|$ is large the term $\beta Q(t)$ in (63) dominates in norm the term $\nabla H(X(t))$, and so GPD “behaves like” a MaxWeight algorithm for large $\|Q(t)\|$, ensuring stability of the queue length process. The latter circumstance also suggests that, in the *heavy traffic* asymptotic regime (cf. [20] and references therein), in which $\|Q(t)\|$ is in fact “typically” large, the network behavior under GPD algorithm (with any fixed β) should be same (in the limit) as that under the corresponding MaxWeight algorithm. In terms of our model, the system is in heavy traffic if, roughly speaking, the rate region V is such that the minimum possible average drift c , which can be provided *simultaneously* to all processing queues $Q_n(t)$, is close to 0.

5 Applications

5.1 Optimal Congestion Control of a Network of Time-Varying Switches

In this section we consider a network which is a special case of the network studied in Section 4, and is a model of a complex heterogeneous communication network. Suppose utility nodes $n \in \mathcal{N}^u$ model “network users,” each generating data traffic. Processing nodes model buffers in the network, where data messages are queued while awaiting service from the

network elements, such as communication links, data routers, wireless links, etc. Complexity and heterogeneity of the system will be captured by the fact that the users' capabilities to inject traffic into the network are randomly time-varying and mutually dependent; similarly, available rates at which data from different buffers can be processed by the network, as well as available routing decisions, are mutually dependent and randomly time-varying. The goal is to find a network control algorithm which maximizes aggregate utility of the generated traffic flows (a function of average traffic rates), subject to the constraint that the network remains stable, that is, data queues in the buffers do not "run away" to infinity. Formally, the special case we consider here is as follows.

The set of utility nodes (users) \mathcal{N}^u is broken down into non-intersecting subsets \mathcal{N}_i indexed by $i \in \mathcal{I}^u$: $\cup_{i \in \mathcal{I}^u} \mathcal{N}_i = \mathcal{N}^u$. Each such subset i we will call a (*traffic*) *generating switch*, so that \mathcal{I}^u is the set of (indices of) these switches. The mode $m_i(t)$ of switch $i \in \mathcal{I}^u$ is an independent irreducible aperiodic Markov chain with the finite state space M_i . Each switch $i \in \mathcal{I}^u$ can make control (traffic generation) decisions independently. When the switch mode is $m_i \in M_i$, the available controls k_i form a finite set $K_i(m_i)$. If generating switch i chooses control $k_i \in K_i(m_i)$ at time t , then each node $n \in \mathcal{N}_i$ generates the amount of data (the number of data bits, or customers) $b_n(k_i) \geq 0$, and the amounts of data $\hat{\lambda}_{nj}(k_i) \geq 0$ are sent to the processing nodes $j \in \mathcal{N}^p$. (We do *not* need to assume that $\sum_{j \in \mathcal{N}^p} \hat{\lambda}_{nj}(k_i) \leq b_n(k_i)$.)

The system utility function is $H(x_1, \dots, x_{N_u})$, where each x_n is interpreted as (long-term) average value of $b_n(k_i(t))$. Assume that the utility function decomposes as

$$H(x_1, \dots, x_{N_u}) = \sum_{i \in \mathcal{I}^u} H^{(i)}(x^{(i)}),$$

where $x^{(i)} = \{x_n, n \in \mathcal{N}_i\}$. (In other words, the system utility is the sum of utilities of individual generating switches.) Assume that each $H^{(i)}(x^{(i)})$ is concave continuously differentiable function on $(-\delta, \infty)^{|\mathcal{N}_i|}$, where $\delta > 0$ is a fixed constant and $|\mathcal{N}_i|$ is the cardinality of \mathcal{N}_i .

The processing nodes are also grouped into independent "switches." Namely, the set of processing nodes (servers) \mathcal{N}^p is broken down into non-intersecting subsets \mathcal{N}_i indexed by $i \in \mathcal{I}^p$: $\cup_{i \in \mathcal{I}^p} \mathcal{N}_i = \mathcal{N}^p$. Each such subset i we will call a (*traffic*) *processing switch*, so that \mathcal{I}^p is the set of (indices of) these switches. The mode $m_i(t)$ of switch $i \in \mathcal{I}^p$ is an independent irreducible aperiodic Markov chain with the finite state space M_i . Each switch $i \in \mathcal{I}^p$ can make control (traffic generation) decisions independently. When the switch mode is $m_i \in M_i$, the available controls k_i form a finite set $K_i(m_i)$. If processing switch i chooses control $k_i \in K_i(m_i)$ at time t , then each node $n \in \mathcal{N}_i$ processes (serves) the number $\mu_n(k_i) \geq 0$ of customers from its queue (or entire queue content, if it is less than $\mu_n(k_i)$), and these customers are independently routed to other processing nodes (including self, in general) with probabilities $p_{nj}(k_i)$, $j \in \mathcal{N}^p$, or leave the system with probability $1 - \sum_{j \in \mathcal{N}^p} p_{nj}(k_i)$.

The mapping of the network described above into the more general model of Section 4 is straightforward. A network mode m is a combination of modes m_i of all individual switches (both generating and processing), and a network control k is a combination of controls of all

individual switches. Naturally, if node n belongs to switch i , then $b_n(k) = b_n(k_i)$ for $n \in \mathcal{N}^u$, and $\mu_n(k) = \mu_n(k_i)$ and $p_{nj}(k) = p_{nj}(k_i)$ for $n \in \mathcal{N}^p$. For each processing node j

$$\lambda_j(k) = \sum_{i \in \mathcal{I}^u} \sum_{n \in \mathcal{N}_i} \hat{\lambda}_{nj}(k_i).$$

For this network, it is easy to see that the GPD algorithm specializes to the following algorithm in which all (generating and processing) switches choose their controls independently, according to the following rules. (We remind that $\beta > 0$ is a - typically small - parameter.)

Control rule for a generating switch $i \in \mathcal{I}^u$ (“Generalized Gradient rule”): At time t choose a control

$$k_i(t) \in \arg \max_{k_i \in K_i(m_i(t))} \sum_{n \in \mathcal{N}_i} \left[(\partial H^{(i)}(X^{(i)}(t)) / \partial x_n) b_n(k_i) - \sum_{j \in \mathcal{N}^p} \beta Q_j(t) \hat{\lambda}_{nj}(k_i) \right], \quad (101)$$

and running averages $X_n(t)$, $n \in \mathcal{N}_i$, are updated as follows:

$$X_n(t+1) = (1 - \beta)X_n(t) + \beta b_n(k_i(t)). \quad (102)$$

The initial values $X_n(0) \geq 0$, $n \in \mathcal{N}_i$, are arbitrary.

Control rule for a processing switch $i \in \mathcal{I}^p$ (“MaxWeight rule”): At time t choose a control

$$k_i(t) \in \arg \max_{k_i \in K_i(m_i(t))} \sum_{n \in \mathcal{N}_i} \left[Q_n(t) \mu_n(k_i) - \sum_{j \in \mathcal{N}^p} Q_j(t) p_{nj}(k_i) \mu_n(k_i) \right], \quad (103)$$

Note that each generating switch only needs to keep track of X_n for “its own” nodes. However, if a switch may generate or route data to processing nodes in other (processing) switches, then this switch needs to know queue lengths at those nodes.

In many cases of interest the above control rules reduce to very simple ones. For example, if a processing switch can never route served customers to other nodes, its control choice only depends on its own state. Or consider a simple but realistic case when a traffic source (utility node) n is independent of other sources (i.e., it in itself forms a switch with utility function $H_n(x_n)$), it is not time-varying (has only one mode), and in any time slot it has only two choices: “send” or “not send” a fixed amount $c > 0$ of data to a fixed processing node j . In this case the source chooses to send data in slot t if and only if $H'_n(X_n(t)) - \beta Q_j(t) \geq 0$.

We note that the non-degeneracy condition (55) for the specialized network described here has precisely same meaning as for the general network of Section 4: there exists at least one control policy which would be able to provide strictly negative long-term drift to all processing node queues, if they would have “infinite” supply of customers. (This condition is only slightly stronger than saying that there exists a control policy able to keep network stable.)

When this condition holds, the specialized GPD algorithm described here is asymptotically optimal.

Remark. In paper [6], which is an independent and contemporaneous work with present paper, authors study the following problem (in terms of the model of this section). It is a one-hop system, consisting of a single processing switch, and all processed customers leaving the system (that is, with all $p_{rj} \equiv 0$). Each processing queue j is “fed” by one independent traffic source $n = n(j)$, with strictly concave increasing utility function $H_n(x_n)$ of average traffic rate x_n . Each traffic source can generate traffic at *any* average rate in any time slot. (Note that in this case sources are never “bottlenecks” limiting system utility. If available traffic generation rates are constrained, as in our model, such additional constraints may effect optimal average traffic rate allocation and optimal utility.) The congestion control algorithm proposed in [6] is such that the control rule for the processing switch is exactly same as the MaxWeight rule (103) above (but with all $p_{nj} = 0$). However, the source traffic generation rule is different: the amount of traffic a source generates in time slot t is a static function of queue length $Q_j(t)$ (of the processing node it feeds). It is shown that this congestion control (asymptotically) maximizes aggregate utility $\sum_n H_n(x_n)$, while keeping queues stable. Thus the main difference of the algorithm in [6] from the corresponding special case of GPD algorithm is the “static” traffic generation rule, different from the “greedy” rule employed by GPD algorithm (see generating switch rule (101) above).

5.2 Optimal Congestion Control with Additional Constraints on the Average Traffic Rates and Average “Power Usage”

Let us consider the network described in the previous Section 5.1, but let us now introduce additional network control constraint that the average traffic rates generated by each source $n \in \mathcal{N}^u$ must be between certain lower and upper bounds, $0 \leq x_n^{min} < x_n^{max}$.

Furthermore, we will extend the network model itself to a seemingly more general model, in which switches as they operate consume certain resource, say “transmission power.” Namely, when switch i chooses control k_i , it uses power $w_i(k_i)$. We would like to impose a further constraint that the average power consumed by switch i does not exceed a fixed upper bound $w_i^{max} > 0$.

We now show that the model extension and all additional constraints can in fact be mapped into the framework of our general network (Section 4), so that, again, a special case of GPD algorithm provides (asymptotically) optimal control.

For each additional constraint we have, let us introduce an additional independent *virtual* processing node. Namely, for every (real) utility node $n \in \mathcal{N}^u$, we introduce independent virtual nodes $\eta(n)$ and $\zeta(n)$, which will correspond to the upper and lower bound constraints on the average rate of generated traffic. We put (with i being the switch containing n):

$$\lambda_{\eta(n)}(k) = \lambda_{\eta(n)}(k_i) = b_n(k_i),$$

$$\begin{aligned}
\mu_{\eta(n)}(k) &\equiv x_n^{max}, \\
\lambda_{\zeta(n)}(k) &\equiv x_n^{min}, \\
\mu_{\zeta(n)}(k) &= \mu_{\zeta(n)}(k_i) = b_n(k_i).
\end{aligned}$$

Note that since nodes $\eta(n)$ and $\zeta(n)$ are virtual, each switch i can maintain and update the queue lengths at all these virtual nodes corresponding to its own real nodes n .

The maximum average power constraints are dealt with similarly to the maximum average traffic rate constraints. Namely, for every switch $i \in \mathcal{I}^u \cup \mathcal{I}^p$, we introduce independent *virtual* node $\nu(i)$ and put:

$$\begin{aligned}
\lambda_{\nu(i)}(k) &= \lambda_{\nu(i)}(k_i) = w_i(k_i), \\
\mu_{\nu(i)}(k) &\equiv w_i^{max}.
\end{aligned}$$

Each switch i can maintain and update the queue length at the corresponding node $\nu(i)$.

The key observation is that *each virtual queue is stable if and only if its corresponding constraint is satisfied*.

The (asymptotically optimal) specialization of the GPD algorithm is as follows. All (generating and processing) switches choose their controls independently, according to the following rules.

Control rule for a generating switch $i \in \mathcal{I}^u$: *At time t choose a control*

$$\begin{aligned}
k_i(t) \in \arg \max_{k_i \in K_i(m_i(t))} \sum_{n \in \mathcal{N}_i} &\left[(\partial H^{(i)}(X^{(i)}(t)) / \partial x_n) b_n(k_i) - \sum_{j \in \mathcal{N}^p} \beta Q_j(t) \hat{\lambda}_{nj}(k_i) \right] \\
&+ \sum_{n \in \mathcal{N}_i} \beta [Q_{\zeta(n)}(t) - Q_{\eta(n)}(t)] b_n(k_i) - \beta w_i(k_i) Q_{\nu(i)}(t).
\end{aligned}$$

The averages $X_n(t)$, $n \in \mathcal{N}_i$, are updated as in (102), with arbitrary initial values $X_n(0) \geq 0$, $n \in \mathcal{N}_i$. The virtual queue lengths, associated with this switch, are updated as follows:

$$Q_{\eta(n)}(t+1) = [Q_{\eta(n)}(t) - x_n^{max}]^+ + b_n(k_i), \quad n \in \mathcal{N}_i,$$

$$Q_{\zeta(n)}(t+1) = [Q_{\zeta(n)}(t) - b_n(k_i)]^+ + x_n^{min}, \quad n \in \mathcal{N}_i,$$

$$Q_{\nu(i)}(t+1) = [Q_{\nu(i)}(t) - w_i^{max}]^+ + w_i(k_i). \quad (104)$$

Control rule for a processing switch $i \in \mathcal{I}^p$: *At time t choose a control*

$$k_i(t) \in \arg \max_{k_i \in K_i(m_i(t))} \sum_{n \in \mathcal{N}_i} \left[Q_n(t) \mu_n(k_i) - \sum_{j \in \mathcal{N}^p} Q_j(t) p_{nj}(k_i) \mu_n(k_i) \right] - w_i(k_i) Q_{\nu(i)}(t).$$

The virtual queue length $Q_{\nu(i)}(t)$ is updated as in (104).

The non-degeneracy condition (55) for the network with additional average traffic rate and average “transmission power” constraints has the following meaning: there exists a control policy which would be able to provide strictly negative long-term drift to all processing node queues, while keeping average traffic generation rates strictly between their lower and upper bounds, and keeping average transmission powers strictly below their upper bounds. (This is only slightly stronger than requiring that it is feasible at all to keep network stable while keeping average traffic generation rates and average transmission power within the desired bounds.) When this condition holds, the specialized GPD algorithm for this network is asymptotically optimal.

We remark here that a special case of the above network, consisting of a single generating switch, with no average power constraints, corresponds to a problem of utility based scheduling in wireless systems, subject to average rate constraints [15, 2]. (In fact, this single-switch model is more general than those in [15, 2].) The specialized GPD algorithm for this case (which is the control rule for a generating switch, without the sums containing Q_j and $Q_{\nu(n)}$) provides an asymptotically optimal policy for the problem. This algorithm is different from that in [15], but is close to that in [2].

5.3 Managing Power Usage (and Other “Costs”) in Wireless Systems

Saving power is a very important consideration in wireless communication systems (cf. [24, 27, 12]). If a resource allocation algorithm (say, transmission rate allocation) has to comply with constraints on the average power usage, this seemingly presents a qualitatively more difficult problem. However, as the model in Section 5.2 demonstrates, the GPD algorithm allows one to accommodate average power constraints rather simply, by introducing virtual queues. In this section we would like to emphasize and extend this point, by looking at different examples and showing that GPD algorithm can handle a large variety of average power management problems by using virtual queues and virtual commodities.

Consider first a special case of the model in Section 5.2, consisting of a single generating switch, with no average rate constraints (but with average power constraint). We obtain a generalized version of the problem studied in [24] (where a linear utility function was used). Specialization of GPD algorithm to this case is a simple parsimonious dynamic strategy (different from that in [24]), not requiring any a priori knowledge of the stationary distribution of the switch mode m_i .

Next, consider another specialization of the network of Section 5.2, described as follows. Suppose, the processing switches are same as before. The generating switches, however, do *not* consume energy, do *not* have average rate constraints, and in each mode m_i have only one available control. Thus, this is an “open-loop” model, where the generating switches are “uncontrollable” - they simply generate random input traffic flows for the network of processing switches. Such model is a generalized version of the models considered in [27, 12], which address the question “How to control an open-loop network so that it remains stable,

without violating average power usage constraints by the processing switches?” The utility function does not matter, so we can set $H(\cdot) \equiv 0$. Then, we obtain a specialized version of the GPD algorithm (different from the algorithms of [27, 12]), which is a simple parsimonious dynamic policy. (Note, that parameter β of GPD algorithm is irrelevant here, because generating switches are not controlled at all.)

Finally, let us further extend, rather than specialize, the model of Section 5.2 as follows. Suppose, for each switch i , in addition to the power $w_i(k_i)$, a control k_i incurs some “cost” $c_i(k_i)$. The value of $c_i(k_i)$ may be the power itself, $c_i(k_i) = w_i(k_i)$, or some function of it, or something completely unrelated to it. Suppose, we would like to maximize the utility function

$$\sum_{i \in \mathcal{I}^u} H^{(i)}(x^{(i)}) - \sum_{i \in \mathcal{I}^u \cup \mathcal{I}^p} y^{(i)}$$

where $y^{(i)}$ is the average cost of controls chosen by switch i . As reader probably already guessed by now, the GPD algorithm easily extends to this case as well, by simply treating costs c_i for each i as another (“virtual”) commodity type. It is straightforward to derive the corresponding form of GPD algorithm from its general form. Note that the resulting control policy is such that each switch i will not need to maintain even its own cost rate estimate, because the utility function is linear in $y^{(i)}$. Thus, the computational complexity of the algorithm is essentially same as that of the algorithm without costs. We remark that, in particular, this algorithm solves the problem of “keeping queues stable at the minimum possible average cost” for the open-loop network described earlier in this section.

6 Appendix

6.1 Proof of Lemma 1

We first show that conditions (10)-(15) imply (16). Let us call a time point $t > 0$ *regular* if proper derivatives $x'(t)$, $q'(t)$, and $f'(t)$ exist, and conditions (11)-(13) hold for this t . Almost all $t \geq 0$ are regular. It will suffice to show that (16) holds at each regular point t . Indeed, suppose t is regular. Then, first, (16) holds trivially when $q_n(t) > 0$. Second, when $q_n(t) = 0$ we must have $q'_n(t) = 0$ (by regularity), and then $v_n(t) \leq 0$ because otherwise (15) would imply $(d^+/dt)q_n(t) = v_n(t) > 0$.

Now, suppose (10), (11), (13), (14) and (16) hold for (x, q) . Let us define

$$f_n(t) = q_n(0) + \int_0^t v_n(\xi) d\xi,$$

so that (12) and the first part of (15) hold. To prove the second part of (15), we further define

$$h_n(t) = \int_0^t I\{q_n(\xi) = 0\}([v_n(\xi)]^+ - v_n(\xi)) d\xi$$

and observe that $q_n(t) = f_n(t) + h_n(t) \geq 0$ for all $t \geq 0$, function $f_n(t)$ is continuous, function $h_n(t)$ is continuous non-decreasing non-negative with $h_n(0) = 0$, and, finally, time t cannot be a point of increase of $h_n(\cdot)$ when $q_n(t) > 0$. This implies (15) (cf. [7], Proposition 2.2.3). Function f satisfying (11), (12) and the first part of (15), is unique because (according to these relations) it is uniquely determined by x and the initial condition $f(0) = q(0)$.

6.2 Estimate of the Speed of Convergence to a Convex Set

The following lemma is a slight modification of Lemma 3 in [21], and is presented here with a proof for completeness. Note that in the differential inclusion (105) in Lemma 20 it is only required that $v(t) \in V$. (There are *no* other conditions on $v(t)$.)

Lemma 20 *Suppose V is a convex compact subset of a finite-dimensional space R^N . Suppose a vector-function $(x(t), t \geq 0)$, taking values in R^N , is absolutely continuous, satisfying the following differential inclusion for almost all $t \geq 0$:*

$$x'(t) = v(t) - x(t), \quad v(t) \in V . \quad (105)$$

Then

- (i) *Both $x(t)$ and the distance $\rho(x(t), V)$ are Lipschitz continuous in $[0, \infty)$;*
- (ii) *The distance $\rho(x(t), V)$ is non-increasing, and moreover, for almost all $t \geq 0$,*

$$\frac{d}{dt}\rho(x(t), V) \leq -\rho(x(t), V) , \quad (106)$$

which implies that

$$\rho(x(t), V) \leq \rho(x(0), V)e^{-t} .$$

- (iii) *The entire trajectory $(x(t), t \geq 0)$ is contained within the convex hull \bar{V} of $V \cup \{x(0)\}$.*

Proof. Function $x(t)$ is Lipschitz continuous on any finite interval because $x(t)$ (and therefore the derivative $x'(t) = v(t) - x(t)$) is bounded (on such interval). This easily implies that $\rho(x(t), V)$ is Lipschitz continuous on any finite interval. This in particular means that, both derivatives $x'(t)$ and $(d/dt)\rho(x(t), V)$ exist at almost all $t \geq 0$. Then, to prove statements (i) and (ii), it suffices to show that (106) holds for any point $t > 0$ such that derivatives $x'(t)$ and $(d/dt)\rho(x(t), V)$ exist, (105) holds, and $\rho(x(t), V) > 0$. (For this will immediately imply that $x(\cdot)$ is bounded on its entire domain $[0, \infty)$.) Let t is such a point. Consider the point ξ which is the (unique) point of V closest to $x(t)$, i.e., $\rho(x(t), \xi) = \rho(x(t), V)$. (See Figure 1.) Consider vectors

$$\alpha_1 = \xi - x(t), \quad \alpha_2 = v(t) - \xi,$$

and

$$\alpha = v(t) - x(t) = \alpha_1 + \alpha_2 .$$

Consider auxiliary linear function

$$h(\theta) = \xi + \alpha_2(\theta - t), \quad \theta \geq t,$$




Figure 1: Proof of Lemma 20.

and note that $h(\theta) \in V$ for all sufficiently small increments $(\theta - t) \geq 0$. Then, we can write

$$\begin{aligned} \frac{d}{dt} \rho(x(t), V) &\leq \frac{d^+}{d\theta} \rho(x(\theta), h(\theta))|_{\theta=t} = \\ &= -\|\alpha_1\| = -\rho(x(t), V). \end{aligned}$$

The proof of (i) and (ii) is complete. Statement (iii) is a corollary of (ii), because $v(t) \in \bar{V}$ for almost all t and $x(0) \in \bar{V}$. \blacksquare

6.3 Proof of (32)

If v^* lies in the interior of V , the statement is obvious. Suppose v^* lies on the boundary of V . Let $C^\circ(v^*)$ denote the cone polar to cone $C^*(v^*)$. (In other words, $C^\circ(v^*)$ is the cone normal to $C^*(v^*)$ at point 0.)

Denote by Ξ the set of all unit length vectors ξ lying on the boundary of $C^\circ(v^*)$. Let $u^*(x)$ denote the normal projection of x on V , that is the point of V closest to x . Since cone $C^\circ(v^*)$ is polar to the normal cone $C^*(v^*)$ (and V is compact), the following holds:

$$\frac{1}{\delta} \sup_{\xi \in \Xi} \|u^*(v^* + \delta\xi) - (v^* + \delta\xi)\| \downarrow 0 \text{ as } \delta \downarrow 0.$$

(Otherwise we would be able to construct $\xi \in \Xi$ and a convex cone A , containing ξ in the interior, such that the set $v^* + A$ intersects with V at a single point v^* . This is impossible since ξ is an element of the polar cone $C^\circ(v^*)$.) This in particular implies that

$$\sup_{\xi \in \Xi} \alpha(u^*(v^* + \delta\xi), \xi) \downarrow 0 \text{ as } \delta \downarrow 0, \quad (107)$$

where $\alpha(\cdot, \cdot)$ denotes the angle between two vectors, and

$$\inf_{\xi \in \Xi} \|u^*(v^* + \delta\xi) - v^*\|/\delta \rightarrow 1 \text{ as } \delta \downarrow 0. \quad (108)$$

Now we can finish the proof of (32). Since vector $\nabla H(x(t)) - q(t)$ is uniformly bounded, at any time t when its distance to $C^*(v^*)$ is away from 0 - namely, it is greater or equal to some $\epsilon_1 > 0$ - the angle between $\nabla H(x(t)) - q(t)$ and $C^*(v^*)$ is bounded away from 0 as well. This implies that the angle between $\nabla H(x(t)) - q(t)$ and the cone $C^\circ(v^*)$ is less than and bounded away from the right angle $\pi/2$. To summarize, if $\xi(t)$ is the vector in Ξ forming the smallest angle with $\nabla H(x(t)) - q(t)$, then $\rho(\nabla H(x(t)) - q(t), C^*(v^*)) \geq \epsilon_1$ implies $\alpha(\nabla H(x(t)) - q(t), \xi(t)) \leq \pi/2 - \epsilon_3$ for some $\epsilon_3 > 0$ (depending on ϵ_1). According to (107) and (108), we can always choose $\delta > 0$ small enough so that $\alpha(u^*(\delta\xi(t)) - v^*, \delta\xi(t)) < \epsilon_3/2$ and $\|u^*(\delta\xi(t))\| \geq \delta/2$. Finally, when $\rho(\nabla H(x(t)) - q(t), C^*(v^*)) \geq \epsilon_1$ we have $\alpha(\nabla H(x(t)) - q(t), u^*(\delta\xi(t)) - v^*) \leq \pi/2 - \epsilon_3/2$ (and, obviously, $\|\nabla H(x(t)) - q(t)\| \geq \epsilon_1$), and therefore

$$\begin{aligned} B_3(t) &= \max_{v \in V} [\nabla H(x(t)) - q(t)] \cdot (v - v^*) \geq [\nabla H(x(t)) - q(t)] \cdot (u^*(\delta\xi(t)) - v^*) \geq \\ &\quad \epsilon_1(\delta/2) \sin(\pi/2 - \epsilon_3/2) > 0. \end{aligned}$$

6.4 Some Elementary Facts

Proposition 1 *Suppose $H(v)$ is a continuously differentiable concave function with convex open domain $\tilde{V} \subseteq R^N$, $N \geq 1$. Suppose $H(v)$ has same value along the line segment L connecting two points $v^{(1)}, v^{(2)} \in \tilde{V}$. Then, $\nabla H(v^{(1)}) = \nabla H(v^{(2)})$. (And then $\nabla H(v)$ is same for all $v \in L$.)*

Proof is elementary - we present it for completeness. Suppose $\nabla H(v^{(1)}) \neq \nabla H(v^{(2)})$. Let us pick a non-zero vector $\zeta \in R^N$ such that

$$\nabla H(v^{(1)}) \cdot \zeta > \nabla H(v^{(2)}) \cdot \zeta. \quad (109)$$

Consider function

$$A(\xi) = \frac{1}{2}[H(v^{(1)} + \zeta\xi) + H(v^{(2)} - \zeta\xi)],$$

which is well defined for all sufficiently small real ξ . By concavity of H , for all (small) ξ , $A(\xi) \leq A(0) = H(v^{(mid)})$, where $v^{(mid)} = [v^{(1)} + v^{(2)}]/2$. On the other hand, by (109), $A'(\xi) > 0$. Contradiction completes the proof. \blacksquare

The following fact is standard and is easy to verify directly.

Proposition 2 *Suppose $V \subset R^N$, $N \geq 1$, is a compact set. The multivalued operator which takes a vector $\zeta \in R^N$ into the set $V_\zeta \doteq \arg \max_{v \in V} \zeta \cdot v$ is closed. Namely, if sequences of*

$\zeta^{(i)} \in R^N$ and $v^{(i)} \in V$, on $i = 1, 2, \dots$ are such that $\zeta^{(i)} \rightarrow \zeta$, $v^{(i)} \in V_{\zeta^{(i)}}$, and $v^{(i)} \rightarrow \hat{v}$, then $\hat{v} \in V_{\zeta}$.

Since V is compact, closedness of the operator implies its upper semicontinuity (cf. [9], Lemma XVI.5.1). Namely, for any $\zeta \in R^N$ and $\epsilon > 0$, there exists $\delta > 0$ such that $\|\zeta' - \zeta\| < \delta$ implies $V_{\zeta'} \subseteq \{v \in V \mid \rho(v, V_{\zeta}) < \epsilon\}$.

Acknowledgement. I would like to thank the referees for very useful comments, which helped in improving presentation in the paper.

References

- [1] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, P. Whiting. Providing Quality of Service over a Shared Wireless Link. *IEEE Communications Magazine*, 2001, Vol.39, No.2, pp.150-154.
- [2] M. Andrews, L. Qian, A. L. Stolyar. Optimal Utility Based Multi-User Throughput Allocation subject to Throughput Constraints. *Proceeding of INFOCOM'2005*, Miami, March 13-17, 2005.
- [3] R. Agrawal, V. Subramanian. Optimality of Certain Channel Aware Scheduling Policies. In *Proc. of the 40th Annual Allerton Conference on Communication, Control, and Computing. Monticello, Illinois, USA, October 2002*.
- [4] J.G.Dai and W.Lin. Maximum Pressure Policies in Stochastic Processing Networks. *Operations Research*, Vol. 53, (2005), pp. 197-218.
- [5] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. John Wiley and Sons, New York, 1986.
- [6] A.Eryilmaz, R.Srikant. Fair Resource Allocation in Wireless Networks using Queue-length-based Scheduling and Congestion Control. *Proceeding of INFOCOM'2005*, Miami, March 13-17, 2005.
- [7] J.M. Harrison. *Brownian Motion and Stochastic Flow Systems*. Wiley, 1985.
- [8] A.Jalali, R.Padovani, and R.Pankaj. Data Throughput of CDMA-HDR, a High Efficiency - High Data Rate Personal Communication Wireless System. In *Proc. of the IEEE Semiannual Vehicular Technology Conference, VTC2000-Spring*, Tokyo, Japan, May 2000.
- [9] L. V. Kantorovich and G. P. Akilov. *Functional Analysis*, 2nd ed. Pergamon Press, New York, 1982.
- [10] F.P.Kelly, A.K.Maullo, D.K.H.Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, Vol. 49, 1998, pp. 237-252.

- [11] F.P.Kelly. Fairness and stability of end-to-end congestion control. *European Journal of Control*, Vol. 9, 2003, pp. 159-176.
- [12] T. Klein and H.Viswanathan. Centralized power control for multi-hop wireless networks. 2003, submitted.
- [13] S. Karlin. *Mathematical Methods in Games, Programming, and Economics*. Dover, New York, 1992.
- [14] S.Liu, T.Basar, R.Srikant. Controlling the Internet: a survey and some new results. 2003, preprint.
- [15] X.Liu, E.K.P.Chong, N.B.Shroff. A Framework for Opportunistic Scheduling in Wireless Networks. *Computer Networks*, Vol. 41, 2003, pp. 451-474.
- [16] S. Low. A Duality Model of TCP and Queue Management Algorithms. *IEEE/ACM Transactions on Networking*, Vol. 11, (2003), No.4, pp. 525-536.
- [17] R.T.Rockafellar. *Convex Analysis*. Princeton Univ. Press, 1970.
- [18] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [19] A.L. Stolyar. On the Stability of Multiclass Queueing Networks: A Relaxed Sufficient Condition via Limiting Fluid Processes. *Markov Processes and Related Fields*, Vol. 1(4), 1995, pp.491-512.
- [20] A. L. Stolyar. MaxWeight Scheduling in a Generalized Switch: State Space Collapse and Workload Minimization in Heavy Traffic. *Annals of Applied Probability*, 2004, Vol.14, No.1, pp.1-53.
- [21] A.L. Stolyar. On the Asymptotic Optimality of the Gradient Scheduling Algorithm for Multi-User Throughput Allocation. *Operations Research*, 2005, Vol. 53, No.1, pp. 12-25.
- [22] L.Tassiulas, A.Ephremides. Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks. *IEEE Transactions on Automatic Control*, Vol. 37, 1992, pp. 1936-1948.
- [23] L.Tassiulas, P.P.Bhattacharya. Allocation of Interdependent Resources for Maximal Throughput. *Commun.Statist. - Stochastic Models*, Vol. 16, (2000), pp. 27-48.
- [24] D.Tse and S. Hanly. Multi-access fading channels: Part I: Polymatroid Structure, Optimal Resource Allocation and Throughput Capacities. *IEEE Transactions on Information Theory*, Vol. 44, (1998), no. 7, pp. 2796-2815, 1998.
- [25] P. Viswanath, D. Tse, and R. Laroia. Opportunistic Beamforming using Dumb Antennas. *IEEE Transactions on Information Theory*, Vol. 48(6), 2002, pp. 1277-1294.
- [26] J. T. Wen and M. Arcak. A Unifying Passivity Framework for Network Flow Control. *IEEE Transactions on Automatic Control*, Vol. 49, (2004), No.2, pp. 162-174.

- [27] E. M. Yeh, A. S. Cohen. Throughput and Delay Optimal Resource Allocation in Multiaccess Fading Channels. *Proceedings of ISIT*, Yokohama, Japan, June-July 2003.