

Accelerated Backpressure Algorithm

Michael Zargham[†], Alejandro Ribeiro[†], Ali Jadbabaie[†]

Abstract—We develop an Accelerated Back Pressure (ABP) algorithm using Accelerated Dual Descent (ADD), a distributed approximate Newton-like algorithm that only uses local information. Our construction is based on writing the backpressure algorithm as the solution to a network feasibility problem solved via stochastic dual subgradient descent. We apply stochastic ADD in place of the stochastic gradient descent algorithm. We prove that the ABP algorithm guarantees stable queues. Our numerical experiments demonstrate a significant improvement in convergence rate, especially when the packet arrival statistics vary over time.

I. INTRODUCTION

This paper considers the problem of joint routing and scheduling in packet networks. Packets are accepted from upper layers as they are generated and marked for delivery to intended destinations. To accomplish delivery of information nodes need to determine routes and schedules capable of accommodating the generated traffic. From a node-centric point of view, individual nodes handle packets that are generated locally as well as packets received from neighboring nodes. The goal of each node is to determine suitable next hops for each flow conducive to successful packet delivery.

A joint solution to this routing and scheduling problem is offered by the backpressure (BP) algorithm [1]. In BP nodes keep track of the number of packets in their local queues for each flow and share this information with neighboring agents. The differences in the number of queued packets at two neighboring terminals are computed for all flows and the transmission capacity of the link is assigned to the flow with the largest queue differential. We can interpret this algorithm by identifying queue differentials with pressure to send packets on a link.

A drawback of BP is the slow convergence rate of this iterative process for route discovery. This is better understood through an alternative interpretation of BP as a dual stochastic subgradient descent algorithm [2], [3]. Joint routing and scheduling can be formulated as the determination of per-flow routing variables that satisfy link capacity and flow conservation constraints. In this model the packet transmission rates are abstracted as continuous variables. The slow convergence rate of BP is thus expected because the convergence rate of subgradient descent algorithms is logarithmic in the number of iterations [4]. Simple modifications can speed up the convergence rate of BP by rendering it equivalent to stochastic gradient descent [5] – as opposed to stochastic subgradient descent. Nevertheless, the resulting linear convergence rate is still a limiting factor in practice.

To speed up the convergence rate of BP we need to incorporate information on the curvature of the dual function. This could

be achieved by using Newton’s method, but the determination of dual Newton steps requires coordination among all nodes in the network. To overcome this limitation, methods to determine approximate Newton steps in a distributed manner have been proposed. Early contributions on this regard are found in [6], [7]. Both of these methods, however, are not fully distributed because they require some level of global coordination. Efforts to overcome this shortcoming include approximating the Hessian inverse with the inverse of its diagonals [8] and the use of consensus iterations to approximate the Newton step [9], [10]. Our work makes use of the accelerated dual descent (ADD) family of algorithms that uses a Taylor’s expansion of the Hessian inverse to compute approximations to the Newton step [11]. Members of the ADD family are indexed by a parameter N indicating that local Newton step approximations are constructed at each node with information gleaned from agents no more than N hops away. Algorithms ADD- N were proven to achieve quadratic convergence rate and demonstrated to reduce the time to find optimal operating points – as measured by the number of communication instances – by two orders of magnitude with respect to regular subgradient descent. It has been shown that stochastic ADD converges for the Network flow optimization in [12]. The goal of this paper is to adapt the ADD family of algorithm to develop variations of BP that achieve quadratic convergence.

Section II introduces the problem of stabilizing queues in a communication network and reviews the backpressure algorithm, [1]. We proceed to demonstrate that the backpressure algorithm is equivalent to stochastic subgradient descent (Section II-A). Once in the optimization framework we review the generalization to soft backpressure (Section II-B). In Section III, we construct the accelerated backpressure algorithm, which is an approximation of Newton’s method using local information and limited communication with neighboring nodes. In Section IV, we prove that the accelerated backpressure algorithm stabilizes the queues. In Section V, we present numerical experiments which demonstrate the performance gains associated with implementing the accelerated backpressure algorithm as compared to the backpressure algorithm from [1] and soft backpressure algorithm from [5].

The main contribution of this work is the introduction of a locally computable second order method for solving the queue stabilization problem. This is particular relevant for cases where the packet arrival statistics vary with time. As shown in Figure 3, the accelerated backpressure algorithm can effectively stabilize queues in networks whose arrival rates vary at higher frequencies than its first order parent algorithms.

II. PRELIMINARIES

Consider a given a network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ where \mathcal{V} is the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of links between nodes. Denote as C_{ij} the capacity of link $(i, j) \in \mathcal{E}$ and define the

This research is supported by Army Research Lab MAST Collaborative Technology Alliance, AFOSR complex networks program, ARO P-57920-NS, NSF CAREER CCF-0952867, and NSF CCF-1017454, ONR MURI N000140810747, ONR Basic Research Challenge N000141210997 and NSF-ECS-0347285.

[†]Michael Zargham, Alejandro Ribeiro and Ali Jadbabaie are with the Department of Electrical and Systems Engineering, University of Pennsylvania.

neighborhood of i as the set $n_i = \{j \in \mathcal{V} | (i, j) \in \mathcal{E}\}$ of nodes j that can communicate directly with i . There is also a set of information flows \mathcal{K} with the destination of flow $k \in \mathcal{K}$ being the node $o_k \in \mathcal{V}$.

At time index t terminal $i \neq o_k$ generates a random number $a_i^k(t)$ of units of information to be delivered to o_k . The random variables $a_i^k(t) \geq 0$ are assumed independent and identically distributed across time with expected value $\mathbb{E}[a_i^k(t)] = a_i^k$. In the same time slot node i routes $r_{ij}^k(t) \geq 0$ units of information through neighboring node $j \in n_i$ and receives $r_{ji}^k(t) \geq 0$ packets from neighbor j . The difference between the total number of received packets $a_i^k(t) + \sum_{j \in n_i} r_{ji}^k(t)$ and the sum of transmitted packets $\sum_{j \in n_i} r_{ij}^k(t)$ is added to the local queue – or subtracted if this quantity is negative. Therefore, the number $q_i^k(t)$ of k -flow packets queued at node i evolves according to

$$q_i^k(t+1) = \left[q_i^k(t) + a_i^k(t) + \sum_{j \in n_i} r_{ji}^k(t) - r_{ij}^k(t) \right]^+, \quad (1)$$

where the projection $[\cdot]^+$ into the nonnegative reals is necessary because the number of packets in queue cannot become negative. We remark that (1) is stated for all nodes $i \neq o_k$ because packets routed to their destinations are removed from the system.

To ensure packet delivery it is sufficient to guarantee that all queues $q_i^k(t+1)$ remain stable. In turn, this can be guaranteed if the average rate at which packets exit queues do not exceed the rate at which packets are loaded into them. To state this formally observe that the time average limit of arrivals satisfies $\lim_{t \rightarrow \infty} a_i^k(t) = \mathbb{E}[a_i^k(t)] := a_i^k$ and define the ergodic limit $r_{ij}^k := \lim_{t \rightarrow \infty} r_{ij}^k(t)$. If the processes $r_{ij}^k(t)$ controlling the movement of information through the network are asymptotically stationary, queue stability follows if ¹

$$\sum_{j \in n_i} r_{ij}^k - r_{ji}^k \geq a_i^k \quad \forall k, i \neq o_k. \quad (2)$$

For future reference define the vector $r := \{r_{ij}^k\}_{k, i \neq o_k, j}$ grouping variables r_{ij}^k for all flows and links. Since at most C_{ij} packets can be transmitted in the link (i, j) the routing variables $r_{ij}^k(t)$ must satisfy

$$\sum_k r_{ij}^k(t) \leq C_{ij}. \quad (3)$$

The joint routing and scheduling problem can be now formally stated as the determination of nonnegative variables $r_{ij}^k(t) \geq 0$ that satisfy (3) for all times t and whose time average limits r_{ij}^k satisfy (2). The BP algorithm solves this problem by assigning all the capacity of the link (i, j) to the flow with the largest queue differential $q_i^k(t) - q_j^k(t)$. Specifically, for each link we determine the flow pressure

$$k_{ij}^* = \operatorname{argmax}_k [q_i^k(t) - q_j^k(t)]^+. \quad (4)$$

If the maximum pressure $\max_k [q_i^k(t) - q_j^k(t)]^+ > 0$ is strictly positive we set $r_{ij}^k(t) = C_{ij}$ for $k = k_{ij}^*$. Otherwise the link remains idle during the time frame. The backpressure algorithm

¹Stability is guaranteed only if the inequalities hold in a strict sense, i.e., $\sum_{j \in n_i} r_{ij}^k - r_{ji}^k > a_i^k$. Equality is allowed here to facilitate connections with optimization problems to be considered later on.

works by observing the queue differentials on each link and then assigning the capacity for each link to the data type with the largest positive queue differential, thus driving the time average of the queue differentials to zero– stabilizing the queues. For the generalizations developed in this paper it is necessary to reinterpret BP as a dual stochastic subgradient descent as we do in the following section.

A. Dual stochastic subgradient descent

Since the parameters that are important for queue stability are the time averages r_{ij}^k of the routing variables $r_{ij}^k(t)$ an alternative view of the joint routing and scheduling problem is the determination of variables r_{ij}^k satisfying (2) and $\sum_k r_{ij}^k \leq C_{ij}$. This can be formulated as the solution of an optimization problem. Let $f_{ij}^k(r_{ij}^k)$ be arbitrary concave functions on \mathbb{R}_+ and consider the optimization problem

$$\begin{aligned} r^* := \operatorname{argmax} \quad & \sum_{k, i \neq o_k, j} f_{ij}^k(r_{ij}^k) \\ \text{s.t.} \quad & \sum_{j \in n_i} r_{ij}^k - r_{ji}^k \geq a_i^k, \quad \forall k, i \neq o_k, \\ & \sum_{k \in \mathcal{K}} r_{ij}^k \leq C_{ij}, \quad \forall (i, j) \in \mathcal{E}. \end{aligned} \quad (5)$$

where the optimization is over nonnegative variables $r_{ij}^k \geq 0$. Since only feasibility is important for queue stability, solutions to (5) ensure stable queues irrespectively of the objective functions $f_{ij}^k(r_{ij}^k)$.

Since the problem in (5) is maximization of a concave function, it can be solved by descending on the dual domain. Start by associating multipliers λ_i^k with the constraint $\sum_{j \in n_i} r_{ij}^k - r_{ji}^k \geq a_i^k$ and keep the constraint $\sum_k r_{ij}^k \leq C_{ij}$ implicit. The corresponding Lagrangian associated with the optimization problem in (5) is

$$\mathcal{L}(r, \lambda) = \sum_{k, i \neq o_k, j} f_{ij}^k(r_{ij}^k) + \sum_{k, i \neq o_k} \lambda_i^k \left(\sum_{j \in n_i} r_{ij}^k - r_{ji}^k - a_i^k \right) \quad (6)$$

where we introduced the vector $\lambda := \{\lambda_i^k\}_{k, i \neq o_k}$ grouping variables λ_i^k for all flows and nodes. The corresponding dual function is defined as $h(\lambda) := \max_{\sum_k r_{ij}^k \leq C_{ij}} \mathcal{L}(r, \lambda)$.

To compute a descent direction for $h(\lambda)$ define the primal Lagrangian maximizers for given λ as $r_{ij}^k(\lambda) := \operatorname{argmax}_{\sum_k r_{ij}^k \leq C_{ij}} \mathcal{L}(r, \lambda)$. A descent direction for the dual function is available in the form of the dual subgradient whose components $\bar{g}_i^k(\lambda)$ are obtained by evaluating the constraint slack associated with the Lagrangian maximizers

$$\bar{g}_i^k(\lambda) := \sum_{j \in n_i} r_{ij}^k(\lambda) - r_{ji}^k(\lambda) - a_i^k. \quad (7)$$

Since the Lagrangian $\mathcal{L}(r, \lambda)$ in (6) is linear in the routing variables r_{ij}^k the determination of the maximizers $r_{ij}^k(\lambda) := \operatorname{argmax}_{\sum_k r_{ij}^k \leq C_{ij}} \mathcal{L}(r, \lambda)$ can be decomposed into the maximization of separate summands. Considering the coupling constraints $\sum_k r_{ij}^k \leq C_{ij}$ it suffices to consider variables $\{r_{ij}^k\}_k$ for all flows across a given link. After reordering terms it follows

that we can compute $r_{ij}^k(\lambda)$ as

$$\begin{aligned} r_{ij}^k(\lambda) &= \operatorname{argmax} \sum_k f_{ij}^k(r_{ij}^k) + r_{ij}^k(\lambda_i^k - \lambda_j^k) \\ \text{s.t.} \quad &\sum_{k \in \mathcal{K}} r_{ij}^k \leq C_{ij}. \end{aligned} \quad (8)$$

Introducing a time index t , subgradients $\bar{g}_i^k(\lambda(t))$ could be computed using (7) with Lagrangian maximizers $r_{ij}^k(\lambda(t))$ given by (8). A subgradient descent iteration could then be defined to find the variables r^* that solve (5); see e.g., [13].

The problem in computing $\bar{g}_i^k(\lambda)$ is that we don't know the average arrival rates a_i^k . We do observe, however, the instantaneous rates $a_i^k(t)$ that are known to satisfy $\mathbb{E}[a_i^k(t)] = a_i^k$. Therefore,

$$g_i^k(\lambda) := \sum_{j \in n_i} r_{ij}^k(\lambda) - r_{ji}^k(\lambda) - a_i^k(t), \quad (9)$$

is a stochastic subgradient of the dual function in the sense that its expected value $\mathbb{E}[g_i^k(\lambda)] = \bar{g}_i^k(\lambda)$ is the subgradient defined in (7). We can then minimize the dual function using a stochastic subgradient descent algorithm. At time t we have multipliers $\lambda(t)$ and determine Lagrangian maximizers $r_{ij}^k(t) := r_{ij}^k(\lambda(t))$ as per (8). We then proceed to update multipliers along the stochastic subgradient direction according to

$$\lambda_i^k(t+1) = \left[\lambda_i^k(t) - \epsilon \left(\sum_{j \in n_i} r_{ij}^k(t) - r_{ji}^k(t) - a_i^k(t) \right) \right]^+, \quad (10)$$

where ϵ is a constant stepsize chosen small enough so as to ensure convergence; see e.g., [5].

Properties of the descent algorithm in (10) vary with the selection of the functions $f_{ij}^k(r_{ij}^k)$. Two cases of interest are when $f_{ij}^k(r_{ij}^k) = 0$ and when $f_{ij}^k(r_{ij}^k)$ are continuously differentiable, strongly convex, and monotone decreasing on \mathbb{R}_+ but otherwise arbitrary. The classical backpressure algorithm, [1] can be recovered by setting $f_{ij}^k(r_{ij}^k) = 0$ for all links flows k and links (i, j) and choosing the step size $\epsilon = 1$.

B. Soft backpressure

Assume now that the functions $f_{ij}^k(r_{ij}^k)$ are continuously differentiable, strongly convex, and monotone decreasing on \mathbb{R}_+ but otherwise arbitrary. In this case the derivatives $\partial f_{ij}^k(x)/\partial x$ of the functions $f_{ij}^k(x)$ are monotonically increasing and thus have inverse functions that we denote as

$$F_{ij}^k(x) := [\partial f_{ij}^k(x)/\partial x]^{-1}(x). \quad (11)$$

The Lagrangian maximizers in (8) can be explicitly written in terms of the derivative inverses $F_{ij}^k(x)$. Furthermore, the maximizers are unique for all λ implying that the dual function is differentiable. We detail these two statements in the following proposition.

Proposition 1 *If the functions $f_{ij}^k(r_{ij}^k)$ in (5) are continuously differentiable, strongly convex, and monotone decreasing on \mathbb{R}_+ , the dual function $h(\lambda) := \max_{\sum_k r_{ij}^k \leq C_{ij}} \mathcal{L}(r, \lambda)$ is differentiable for all λ . Furthermore, the gradient component along the λ_i^k direction is $g_i^k(\lambda)$ as defined in (9) with*

$$r_{ij}^k(\lambda) = F_{ij}^k \left(-[\lambda_i^k - \lambda_j^k - \mu_{ij}(\lambda)]^+ \right) \quad (12)$$

where $\mu_{ij}(\lambda)$ is 0 if $\sum_k F_{ij}^k(-[\lambda_i^k - \lambda_j^k]^+) \leq C_{ij}$ and in any other case is the solution to the equation

$$\sum_k F_{ij}^k \left(-[\lambda_i^k - \lambda_j^k - \mu_{ij}(\lambda)]^+ \right) = C_{ij}. \quad (13)$$

Proof: The dual problem $\min_{\lambda} h(\lambda)$ is strongly smooth because the primal (5) is maximization of a strongly concave function, [14]. Strong smoothness implies the dual function is everywhere differentiable and thus there must be unique primal maximizers. In our case, the dual gradient comes from differentiating (6) with respect to λ_i^k which yields (9). Substituting the unique maximizers we have the unique dual gradient.

In order to find the previously mentioned unique maximizers we consider the primal optimization (8). We dualize the additional constrain to get the extended Lagrangian

$$\bar{\mathcal{L}}(r, \mu) = \sum_k f(r_{ij}^k) + r_{ij}^k(\lambda_i^k - \lambda_j^k) + \mu_{ij}(C_{ij} - \sum_k r_{ij}^k). \quad (14)$$

Considering the KarushKuhnTucker (KKT) optimality conditions as defined in [15][Section 5.5] for (14) yields the equations

$$f'(r_{ij}^k) = -[\lambda_i^k - \lambda_j^k - \mu_{ij}]^+ \quad (15)$$

$$\sum_k r_{ij}^k \leq C_{ij} \quad (16)$$

for all $(i, j) \in \mathcal{E}$. Applying the definition of $F_{ij}^k(\cdot)$ from equation (11) to (15) we get the desired relation in (12). It remains to enforce (16) by selection of μ_{ij} which gives us condition (13). The assertion that $\mu_{ij} = 0$ when $\sum_k F_{ij}^k(-[\lambda_i^k - \lambda_j^k]^+) \leq C_{ij}$ hold by the principal of complementary slackness, detailed in [15][Section 5.5]. ■

While (13) does not have a closed for solution it can be computed quickly numerically using a binary search because it is a simple single variable root finding problem. Computation time cost remains small compared to communication time cost.

An important difference between backpressure and soft backpressure is that the former is equivalent to stochastic *subgradient* descent whereas the latter is tantamount to stochastic *gradient* descent because the dual function is differentiable. This improves the average convergence rate from logarithmic – expected distance to optimal variables decreasing like c/t for some constant c – to linear – expected distance to optimality proportional to c^t for some c . Linear convergence is still not satisfactory, however, motivating the accelerated backpressure algorithm that we introduce in the following section.

III. ACCELERATED BACKPRESSURE

Backpressure type algorithms exhibit slow convergence because they are fundamentally first order methods. In the centralized deterministic case we can speed up using the Newton method. In our case, the stochastic nature of the problem is not an issue because the arrival rates a_i^k do not appear in the Hessian

$$H = \begin{bmatrix} H_{11} & H_{12} & \cdots & H_{1n} \\ H_{21} & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ H_{n1} & \cdots & \cdots & H_{nn} \end{bmatrix} \quad (17)$$

where

$$[H_{ij}]_{ks} = \frac{\partial^2 \mathcal{L}(r(\lambda), \lambda)}{\partial \lambda_i^k \partial \lambda_j^s}. \quad (18)$$

Proposition 2 *The Dual (Generalized) Hessian, $H(\lambda) = \nabla^2 \mathcal{L}(r(\lambda), \lambda)$ is block sparse with respect to the graph \mathcal{G} :*

$$H_{ij} = \mathbf{0} \quad \forall i \neq j \text{ s.t. } (i, j), (j, i) \notin \mathcal{E} \quad (19)$$

and the diagonal blocks are H_{ii} are positive semidefinite.

Proof: The existence of the generalized Hessian is guaranteed by [16]. In order to compute the dual (generalized) Hessian, we begin by computing the optimal flow rates $r_{ij}^k(\lambda)$ from the optimal queue priorities as defined in (12). Substituting into

$$\frac{\partial \mathcal{L}(r(\lambda), \lambda)}{\partial \lambda_i^k} = \sum_{j \in n_i} r_{ji}^k(t) - r_{ij}^k(t) - a_i^k \quad (20)$$

and differentiating with respect to λ we construct the Hessian. Since a_i^k is a constant the key is differentiating $r_{ij}^k(\lambda)$ with respect to λ . We can differentiate (12) using the chain rule yielding

$$\frac{\partial r_{ij}^k(\lambda)}{\partial \lambda_i^k} = \frac{\partial F_{ij}^k(x)}{\partial x} \frac{\partial}{\partial \lambda_i^k} \left(-[\lambda_i^k - \lambda_j^k - \mu_{ij}]^+ \right) \quad (21)$$

The existence of $\partial F_{ij}^k(x)/\partial x$ is guaranteed by our assumptions on the edge costs f_{ij}^k . Differentiating $-[\lambda_i^k - \lambda_j^k - \mu_{ij}]^+$ is done by observing that this function becomes a saturated linear function of r_{ij}^k when r_{ij}^k is non-zero and exactly zero when $r_{ij}^k = 0$. The slope of the linear function is determined by whether the edge (i, j) is at capacity ($\mu_{ij} > 0$). Discontinuities in r_{ij}^k occur when $r_{ij}^k = 0$ and when the projection becomes active, resulting in points from the derivative. We define the derivatives to be zero out these points since they are adjacent to regions where the derivatives are defined and equal to zero. We express the derivative in terms of the active sets

$$\mathcal{A}_{ij} = \{k \in \mathcal{K} : r_{ij}^k(\lambda) > 0\} \quad (22)$$

for each edge (i, j) and the local variables μ_{ij} and $\{r_{ij}^k\}$. Using (21) the diagonal elements H_{ii} are defined

$$\begin{aligned} [H_{ii}]_{kk} &= \sum_{j \in n_i} F'(r_{ij}^k) \mathbf{1}(k \in \mathcal{A}_{ij}) \left(\frac{\mathbf{1}(\mu_{ij} > 0)}{|\mathcal{A}_{ij}|} - 1 \right) \\ &+ F'(r_{ji}^k) \mathbf{1}(k \in \mathcal{A}_{ji}) \left(\frac{\mathbf{1}(\mu_{ji} > 0)}{|\mathcal{A}_{ji}|} - 1 \right). \end{aligned} \quad (23)$$

Using same method the off diagonal elements of H_{ii} are computed

$$\begin{aligned} [H_{ii}]_{ks} &= \sum_{j \in n_i} F'(r_{ij}^k) \mathbf{1}(k, s \in \mathcal{A}_{ij}) \left(\frac{\mathbf{1}(\mu_{ij} > 0)}{|\mathcal{A}_{ij}|} \right) \\ &+ F'(r_{ji}^s) \mathbf{1}(k, s \in \mathcal{A}_{ji}) \left(\frac{\mathbf{1}(\mu_{ji} > 0)}{|\mathcal{A}_{ji}|} \right). \end{aligned} \quad (24)$$

We observe that $\sum_{j \in n_i} [H_{ij}]_{kk} = [H_{ii}]_{kk}$ from the definitions in (23) and (24). Also, $F'(r_{ij}^k) \leq 0$ from (11) and the our assumption that $f_{ij}^k(\cdot)$ is monotone decreasing, therefore the diagonal elements are positive and according to [17][Section 6.2] H_{ii} is positive semidefinite.

We can also compute the off diagonal blocks by differentiating using (21). The diagonal elements of the the off-diagonal blocks H_{ij} are given by

$$\begin{aligned} [H_{ij}]_{kk} &= F'(r_{ij}^k) \mathbf{1}(k \in \mathcal{A}_{ij}) \left(1 - \frac{\mathbf{1}(\mu_{ij} > 0)}{|\mathcal{A}_{ij}|} \right) \\ &+ F'(r_{ji}^k) \mathbf{1}(k \in \mathcal{A}_{ji}) \left(1 - \frac{\mathbf{1}(\mu_{ji} > 0)}{|\mathcal{A}_{ji}|} \right) \end{aligned} \quad (25)$$

and the off-diagonal elements of the off-diagonal blocks are

$$\begin{aligned} [H_{ij}]_{ks} &= -F'(r_{ij}^k) \mathbf{1}(k, s \in \mathcal{A}_{ij}) \left(\frac{\mathbf{1}(\mu_{ij} > 0)}{|\mathcal{A}_{ij}|} \right) \\ &- F'(r_{ji}^s) \mathbf{1}(k, s \in \mathcal{A}_{ji}) \left(\frac{\mathbf{1}(\mu_{ji} > 0)}{|\mathcal{A}_{ji}|} \right). \end{aligned} \quad (26)$$

Observe that the action set \mathcal{A}_{ij} is empty by definition if there is no link $(i, j) \in \mathcal{E}$, therefore k cannot be an element of \mathcal{A}_{ij} . Plugging this fact into the definition of the off diagonal blocks H_{ij} found in equations (25) and (26) we conclude that $H_{ij} = \mathbf{0}$ whenever $i \neq j$ and $(i, j), (j, i) \notin \mathcal{E}$. ■

Proposition 2 guarantees that all elements of the Hessian can be computed using local information. Elements of the Hessian require knowledge of the local active sets \mathcal{A}_{ij} from (3), which are computed using the local flow values $\{r_{ij}^k\}_k$. Furthermore, Proposition 2 gives us positive semi-definiteness of the diagonal blocks H_{ii} indicating that the nodes depend positively on their own queues. This fits our intuition because we expect penalties on a specific queue to become larger when those queues become larger. Later we will use this fact to show that our splitting matrix \bar{D} is invertible.

Newtons method is centralized because it requires the inverse of the Hessian matrix. At each time instance nodes send their multipliers $\lambda_i^k(t)$ to their neighbors. After receiving multiplier information from its neighbors, each node can compute the multiplier differentials $\lambda_i^k(t) - \lambda_j^k(t)$ for each edge. The nodes then solve for μ_{ij} on each of its outgoing edges by using a rootfinder to solve the local constraint in (13). The capacity of each edge is then allocated to the commodities via reverse waterfilling as defined in (12). Once the transmission rates are set each node can observe its net packet gain which is equivalent to the stochastic gradient as defined in (9). The nodes must also compute the matrix of second derivatives, H_{ij} for each of its outgoing edges according to (23)–(26) which are functions of local transmission rates $\{r_{ij}^k\}_{k, j \in n(i)}$. In order to compute the Newton direction nodes must then send their sub gradients and g_i and all of their Hessian blocks H_{ij} to every node in the Network. With this information each node can invert the global Hessian matrix and compute its Newton direction $d_i = -\sum_j [H^{-1}]_{ij} g_j$. Finally, each node updates its multipliers by adding d_i^k to its current multipliers. In this algorithm the multipliers are no longer equivalent to the queue lengths but can be thought of as "queue priorities" which take into account the queue lengths and the networks ability to handle this queues.

A. A distributed approximation of the Newton direction

In order to accelerate backpressure and retain its distributed nature we compute the dual update direction based on the ADD-N algorithm defined in [18] which leverages the sparsity of

Algorithm 1: Local Approximation of Newton's Direction

1 **for** all neighbors $j \in n(i)$ **do**
 2 Observe local flows $\{r_{ij}^k\}_k$ and $\{r_{ji}^k\}_k$
 3 Compute active sets

$$\mathcal{A}_{ij} = \{k \in \mathcal{K} : r_{ij}^k(\lambda) > 0\}$$

 4 Compute H_{ij} according to (25) and (26).
 5 Set $\bar{B}_{ij} = -H_{ij}$.
 6 **end**
 7 Compute H_{ii} according to (23) and (24).
 8 Compute $\bar{D}_{ii}^{-1} = (H_{ii} + I)^{-1}$
 9 Set $\bar{B}_{ii} = I$
 10 Compute $g_i = \{g_i^k\}_k$ according to

$$g_i^k = \sum_{j \in n_i} r_{ji}^k(t) - r_{ij}^k(t) - a_i^k(t)$$

11 **for** all neighbors $j \in n(i)$ **do**
 12 Communicate \bar{D}_{ii}^{-1} and g_i
 13 Receive \bar{D}_{jj}^{-1} and g_j
 14 **end**
 15 Compute $d_i = \{d_i^k\}_k$ according to

$$d_i = -\bar{D}_{ii}^{-1} g_i - \sum_{j \in n_i} \bar{D}_{ii}^{-1} \bar{B}_{ij} \bar{D}_{jj}^{-1} g_j$$

the Hessian to approximate its inverse. The ADD-N algorithm requires a splitting of the Hessian. We define a block splitting $H = (D + I) - (B + I)$ where D is block diagonal defined by

$$D_{ii} = H_{ii} \quad (27)$$

and B is block sparse defined by

$$B = D - H. \quad (28)$$

The ADD-N method gives us the approximate Hessian inverse

$$H^{-1} \approx \bar{H}^{(N)} = \sum_{\tau=0}^N ((D + I)^{-1} (B + I))^\tau (D + I)^{-1}. \quad (29)$$

The diagonal blocks $[D + I]_{ii}$ are positive definite for all i because D_{ii} is positive semi-definite from Proposition 2. We define simplified notation for our splitting

$$H = \bar{D} - \bar{B} \quad (30)$$

where $\bar{D} = D + I$ and $\bar{B} = B + I$ giving us the expression

$$\bar{H}^{(N)} = \sum_{\tau=0}^N (\bar{D}^{-1} \bar{B})^\tau \bar{D}^{-1}. \quad (31)$$

By construction, the matrix $\bar{H}^{(N)}$ is block sparse such that $\bar{H}_{ij}^{(N)} > \mathbf{0}$ only if node i and j are N or fewer hops apart in \mathcal{G} . We are interested in the fully distributed case where only one hop neighbor information is required. Selecting $N = 1$, the dual update direction is computed as

$$d_i(t) = -\bar{D}_{ii}^{-1} g_i(t) - \sum_{j \in n_i} \bar{D}_{ii}^{-1} \bar{B}_{ij} \bar{D}_{jj}^{-1} g_j(t). \quad (32)$$

The update direction d_i^k can be computed using only local information acquired by a single exchange of information with

Algorithm 2: Accelerated Backpressure for node i

1 Observe $q_i^k(0)$.
 2 Initialize $\lambda_i^k(0) = q_i^k(0)$ for all k and $i \neq \text{dest}(k)$
 3 **for** $t = 0, 1, 2, \dots$ **do**
 4 **for** all neighbors $j \in n(i)$ **do**
 5 Send multipliers $\{\lambda_i^k(t)\}_k$ – Receive multipliers $\{\lambda_j^k(t)\}_k$
 6 Compute μ_{ij} such that

$$\sum_k [q_i^k(t) - q_j^k(t) + \beta_{ij}^k - \mu_{ij}]^+$$

 7 Transmit packets at rate

$$r_{ij}^k(t) = F(-[q_i^k(t) - q_j^k(t) + \beta_{ij}^k - \mu_{ij}]^+)$$

 8 **end**
 9 Send variables $\{r_{ij}^k(t)\}_{kj}$ – Receive variables $\{r_{ji}^k(t)\}_{kj}$
 10 Compute stochastic gradient $\{g_i^k(t)\}_k$

$$g_i^k(t) = \sum_{j \in n_i} r_{ji}^k(t) - r_{ij}^k(t) - a_i^k(t)$$

 11 Compute d_i^k by executing **Algorithm 1**
 12 Update the dual variables

$$\lambda_i^k(t+1) = \lambda_i^k(t) + \epsilon d_i^k(t)$$

 13 **end**

neighboring nodes as described in Algorithm 1. Each node uses knowledge of the flows entering $\{r_{ij}^k\}_k$ and leaving $\{r_{ji}^k\}_k$ to compute the active sets \mathcal{A}_{ij} . Using the active sets and the flows node i can compute the Hessian submatrix H_{ij} for each of its neighbors $j \in n(i)$. The sub matrix $H_{ii} = -\sum_{j \in n(i)} H_{ij}$. From the Hessian submatrices node i can compute the submatrices of our splitting $\bar{B}_{ij} = -H_{ij}$, $\bar{B}_{ii} = I$, $\bar{D}_{ii} = I + H_{ii}$. The gradient $\{g_i^k\}_k$ is compute from the flows $\{r_{ij}^k\}_k$ and $\{r_{ji}^k\}_k$. Node i can invert \bar{D} locally because it is block diagonal. Node i exchanges its local variables \bar{D}_{ii}^{-1} and $g_i = \{g_i^k\}_k$ with its neighbors $j \in n(i)$ allowing the computation of the approximate newton direction at node i , d_i given by equation (32).

The dual update for the accelerated backpressure algorithm is given by

$$\lambda_i(t+1) = \lambda_i(t) - \sum_{j \in n_i} \bar{H}_{ij} g_j \quad (33)$$

where $[\lambda_i]_k = \lambda_i^k$ is the local vector of duals at node i and \bar{H}_{ij} is the i, j block of the matrix \bar{H} . The dual updates for variables belonging to node i depend only on the variables r_{ij}^k and g_i^k belonging to nodes $j \in n_i$. The accelerated backpressure algorithm is given by

Algorithm 2 works like soft backpressure but the set of queue priorities $\lambda_i^k(t)$ are not equivalent to the queue lengths q_i^k . The queue priorities are updated using information about not only node i 's queue but also the queues at neighboring nodes $j \in n_i$. The additional information is weighted according the approximate inverse Hessian. The cost of this change is that rather than simply observe the queues as they evolve we must observe the realized flows r_{ij}^k and update the queue priorities according. This requires one additional exchange of information with neighboring nodes, which appears in Algorithm 1. However the slightly increased communication overhead of the accelerated

backpressure algorithm is offset by significant improvement in convergence rate which we demonstrate in the following sections.

IV. STABILITY ANALYSIS

In order to claim the accelerated backpressure algorithm is an alternative to the backpressure and soft backpressure algorithms we need to guarantee that it achieves queue stability. In this section we leverage results from the the stability analysis of the backpressure algorithm detailed extensively in [2]. We begin by proving the stability of the dual variables because they are analogous to queue lengths in the backpressure analysis.

Proposition 3 *Given a Network \mathcal{G} with packet arrivals $\mathbb{E}[a_i^k(t)] = a_i^k$ and edge capacities C_{ij} , routing packets according to Algorithm 2, we have stable dual variables*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \sum_{i,k} \mathbb{E} \lambda_i^k(\tau) \leq \frac{B}{2\delta} \quad (34)$$

for each commodity k and node $i \neq \text{dest}(k)$, where $B, \delta > 0$.

Corollary 1 *Stability of the dual variables $\lambda_i^k(t)$ implies stability of the queues $q_i^k(t)$.*

Proposition 3 guarantees the stability of the dual variables $\lambda_i^k(t)$ (queue priorities) which take the role the queues have in BP and SBP for determining routing in the ABP algorithm. Demonstrating that the $\lambda_i^k(t)$ sequence is stable tells us that the algorithm itself is stable in the sense that the routing assignments stabilize (because they are an explicit function of $\lambda_i^k(t)$). It is Corollary 1 which guarantees that the queues themselves remain stable. This is done by observing that the evolution of $\lambda_i^k(t)$ and $q_i^k(t)$ both evolve based on the sequence of dual gradients $g_i^k(t)$. Since the ABP algorithm solves the optimization (5), the dual gradient $g_i^k(t)$ tends to zero on average so we are not surprised to see that both $\lambda_i^k(t)$ and $q_i^k(t)$ are stable.

V. NUMERICAL RESULTS

Our choice of objective function is

$$f_{ij}^k(r_{ij}^k) = -\frac{1}{2} (r_{ij}^k)^2 + \beta_{ij}^k r_{ij}^k \quad (35)$$

with 5 data types. The quadratic term captures an increasing cost of routing larger quantities of packets across a link and help to eliminate myopic routing choices that lead to sending packets in cycles. The linear term β is introduced to reward sending packets to their destinations. In our simulations, $\beta_{ij}^k = 10$ for all edges routing to their respective data type destinations $j = \text{dest}(k)$ and all other i, j, k , $\beta_{ij}^k = 0$. The Link capacities are select uniformly randomly $[0, 100]$.

We demonstrate that the ABP algorithm solved the deterministic networking problem (2). In figure 2, the ABP algorithm is applying to case with fixed arrival rates and the queues stabilize after only about 30 iterations with under 1000 total queued packets across all queues. Soft backpressure requires over 100 iterations to stabilize and the queues total around 8000 queued packets and backpressure appears to be stable after about 100 iterations and over 10000 packets remain queued. An important observation is that in the deterministic case the accelerated




Fig. 1. The numerical experiments for the ABP algorithm presented in this section are performed on this simple 10 node network with 5 data types. The destinations are unique for each data type and are chosen randomly.

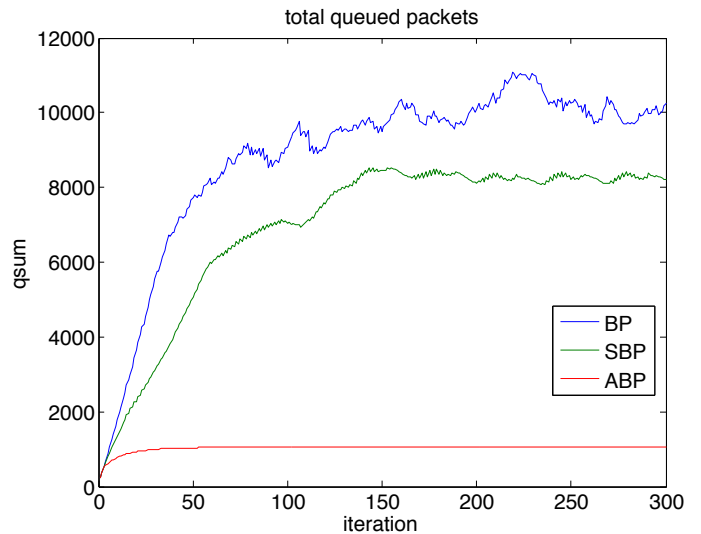


Fig. 2. The Accelerated Backpressure algorithm stabilizes the queues faster and with smaller queue totals than the Soft Backpressure or the traditional Backpressure algorithms in the case where the arrivals rates are deterministic.

backpressure algorithm converges exactly to the optimal dual variables (causing the queue totals to have no variance) while the backpressure and softbackpressure converge on average and retain some variance.

In order to demonstrate the effect of having a faster convergence rate we consider the case where the nodes are divided into two equal sized sets. At time t one of the two sets will be actively receiving packets from outside the network. The frequency at which we switch between active sets captures how fast the system is varying. The example shown in Figures 4 and 3 switches active sets every 10 iterations. The key observation in 4 is that the dual variables are highly volatile for backpressure and soft backpressure. Recalling that in the case of backpressure and soft backpressure the duals are equivalent to the queue lengths, volatility in the duals is volatility in the queues. However, accelerated backpressure converges quickly enough that the variations in the underlying statistics at a rate of once every 10 iterations has no significant effect. From Figure 3 it is clear that by having much more stable dual variables we are able to compute more effective transmission rate assignments and keep the queues from growing.




Fig. 3. The variation in the underlying statistics for the packet arrival rates causes the backpressure and soft backpressure algorithm to stabilize much more slowly and in some cases not at all.

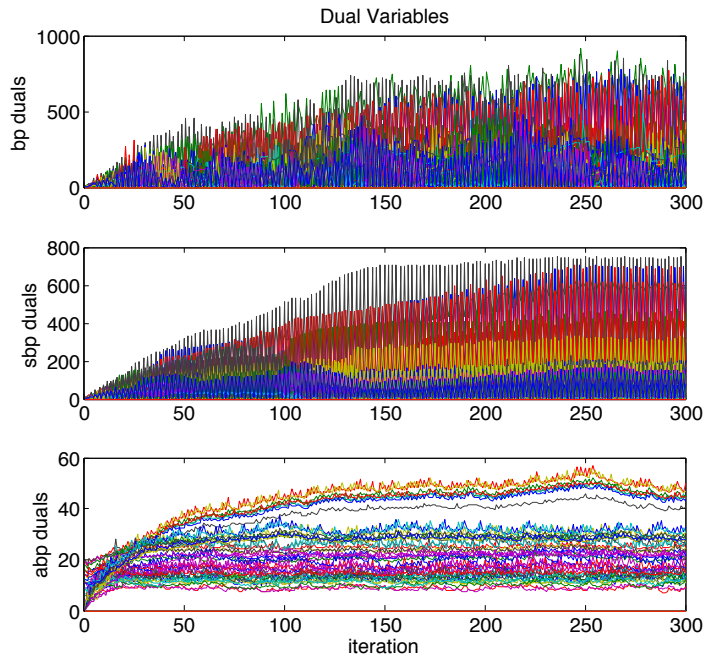


Fig. 4. The Accelerated Backpressure algorithm converges to the optimal dual variables at a rate significantly greater than the backpressure or soft backpressure algorithm, allowing the algorithm to track the optimal duals even when the underlying arrival rates vary.

VI. CONCLUSION

In this work we have presented a novel method for computing packet routes in networks based on applying an approximate Newton's method to the backpressure routing problem. This approach retains the distributed information structure necessary for implementing the algorithm efficiently at the node level. We presented node level protocols and proved that our algorithm stabilizes queues provided the arrival rates and capacities are chosen such that it is possible to stabilize the queues. In numerical experiments, we demonstrate significant improvements in

convergence rate leading to much smaller queues and the ability to stabilize queues even when the arrival rate statistic vary. Our forthcoming work focuses on extending the analysis of the ABP algorithm to explicitly consider the convergence rates.

REFERENCES

- [1] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37:1936–1948, 1992.
- [2] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, 1:1–144, 2006.
- [3] M. J. Neely, E. Modiano, and C. E. Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Journal on Selected Areas in Communications, Special Issue on Wireless Ad-Hoc Networks*, 23:89–103, 2005.
- [4] A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, forthcoming, 2008.
- [5] A. Ribeiro. Stochastic soft backpressure algorithms for routing and scheduling in wireless ad-hoc networks. In *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2009.
- [6] Bertsekas and Gafni. Projected newton methods and optimization of multi-commodity flow. *IEEE Transactions on Automatic Control*, 28:1090–1096, 1983.
- [7] J. G. Klinckewicz. A newton method for convex separable network flow problems. *Bell Laboratories*, 1983.
- [8] S. Athuraliyan and S. Low. Optimization flow control with newton-like algorithm. *Telecommunications Systems*, 15:345–358, 2000.
- [9] A. Jadbabaie, A. Ozdaglar, and M. Zargham. A distributed newton method for network optimization. In *Proceedings of IEEE CDC*, 2009.
- [10] E. Wei, A. Ozdaglar, and A. Jadbabaie. A distributed newton method for network utility maximization. *IEEE Transactions on Automatic Control (Accepted)*, 2011.
- [11] M. Zargham, A. Ribeiro, A. Ozdaglar, and A. Jadbabaie. Accelerated dual descent for network optimization. *IEEE Transactions on Automatic Control*, submitted.
- [12] M. Zargham, A. Ribeiro, and A. Jadbabaie. Network optimization under uncertainty. In *Proceedings of IEEE CDC*, 2012.
- [13] A. Ribeiro and G. B. Giannakis. Separation theorems of wireless networking. *IEEE Transactions on Information Theory*, 2007.
- [14] Sham M. Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Toyota Technological Institute, Chicago, USA*, 2009.
- [15] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- [16] M. Zargham, A. Ribeiro, and A. Jadbabaie. Accelerated dual descent for constrained convex network flow optimization. In *Proceedings of IEEE CDC*, 2013.
- [17] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York, 1979.
- [18] M. Zargham, A. Ribeiro, A. Ozdaglar, and A. Jadbabaie. Accelerated dual descent for network optimization. In *Proceedings of IEEE ACC*, 2011.