

Backpressure Routing Made Practical

Scott Moeller*, Avinash Sridharan*, Bhaskar Krishnamachari*

*Viterbi School of Engineering

University of Southern California, Los Angeles, CA 90007

Email: {smoeller,asridhar,bkrishna}@usc.edu

Omprakash Gnawali†

†School of Engineering

Stanford University, Stanford, CA 94305

Email: gnawali@cs.stanford.edu

Abstract—Current data collection protocols for wireless sensor networks are mostly based on quasi-static minimum-cost routing trees. We consider an alternative, highly-agile approach called backpressure routing, in which routing and forwarding decisions are made on a per-packet basis. Although there is a considerable theoretical literature on backpressure routing, it has not been implemented on practical systems to date due to concerns about packet looping and large packet delays. Addressing these concerns, we present the Backpressure Collection Protocol (BCP) for sensor networks, the first ever implementation of dynamic backpressure routing in wireless networks. In particular, we demonstrate for the first time that replacing the traditional FIFO queue service in backpressure routing with LIFO queues reduces the average end-to-end packet delays for delivered packets drastically (98% under low load). Under static network settings, BCP shows a more than 60% improvement in max-min rate over the state of the art Collection Tree Protocol (CTP). We also empirically demonstrate the superior delivery performance of BCP in network settings of extreme external interference.

We’ve implemented the Backpressure Collection Protocol (BCP), the first dynamic backpressure routing protocol grounded in the emerging theoretical works of Performance Optimal Lyapunov Networking ([3],[4]), a utility optimization enhancement to the landmark work by Tassiulas and Ephremides [6]. Though there have been numerous theoretical contributions to this field in recent years, and though efforts have been made to apply the framework to MAC design ([7],[1]) or at the transport layer [5], there has been no tangible system development of dynamic backpressure routing. In writing BCP, we targeted many-to-one collection scenarios, common in wireless sensor networks.

The crux of this protocol is the following weight computation at a node i for each of its neighbors:

$$w_{i,j} = (\Delta Q_{i,j} - V \cdot ETX_{i \rightarrow j}) \cdot \bar{R}_{i \rightarrow j}$$

Here, $\Delta Q_{i,j} = Q_i - Q_j$ is the queue differential (backpressure), with Q_i and Q_j representing the backlog of nodes i and j respectively, $\bar{R}_{i \rightarrow j}$ is the estimated link rate, and $ETX_{i \rightarrow j}$ is the expected total number of transmissions on the given link to successfully send a unique packet¹. The V parameter is a constant that trades *system queue occupancy* for ETX minimization. Once the weight is computed, each node uses it

as the basis for making independent routing (who to try and send the packet to) and forwarding (whether to transmit the packet) decisions as follows: **Routing decision:** Node i identifies the link (i, j^*) with the highest value of the backpressure weight as the next hop for the packet. **Forwarding decision:** if $w_{i,j^*} > 0$, the packet is forwarded (i.e. sent to the link layer for transmission to the designated neighbor), else the packet is held until the metric is recomputed.

In implementing BCP dynamic backpressure routing for wireless sensor networks, we generated solutions to three substantial challenges that arose:

- ETX minimization combats link layer losses that occur at low source rates due to looping.
- Floating Queues prevent queue overflow, allowing scalability with bounded probability of packet drops.
- LIFO service priority addresses high latency of packet delivery, closing the delay gap between backpressure routing and tree algorithms.

We implemented BCP in TinyOS 2.x, a lightweight operating system for wireless sensor networks, and benchmarked it against the Collection Tree Protocol (CTP) [2]. Using quasi-static routing trees generated through low overhead exponential back-off beaconing mechanisms, CTP has become the standard against which new TinyOS protocols are compared. The code footprint of BCP (CTP) is 23 KB (27 KB), indirectly indicating the relative simplicity of the link estimation and routing techniques employed in backpressure algorithms. Experiments were conducted on a 40 node subsection of the Tutonet, an indoor wireless sensor network testbed consisting of IEEE 802.15.4-based Tmote Sky devices. All experiments were carried out on 802.15.4 channel 26, having low uncontrolled external interference from 802.11 sources within Tutonet.

In static network tests over source rates of 0.5 packets per second per source through 1.66 packets per second per source, measured sink goodput for 35 minute experiments are given in Figure 1. At per-source rates at or above 1.0 packets per second per source, CTP begins to tail drop packets near the sink. This results in a loss of delivery ratio, particularly for motes near the rear of the network. In BCP, we observe that the dynamic backpressure routing avoids the formation of hot spots near the sink, supporting stable source rates up to 1.66 packets per second per source, an increase of more than 60%. Per node forwarding Queue sizes were maintained at 12 packets for both

¹This ETX term in the link weight computation is an innovation on our part. Before we included the ETX minimization, in our testbed experiments, we found that link-layer losses can cause substantial packet drops in loop formations of even minor duration. While backpressure should stop looping behaviors once backlogs build up, this was short-circuited by the presence of these losses.

CTP and BCP.

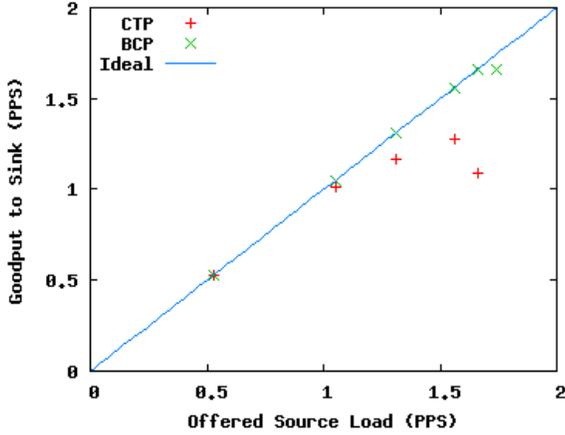


Fig. 1. Goodput in static tests over source rate for BCP and CTP.

Packet delivery latency is a well known challenge to backpressure routing algorithms. The exploration of routing alternatives and persistent queue backlogs both result in delivery delays that are both highly variable and substantially greater on average than delays achieved by tree-based algorithms. Early experimentation with BCP on Tutonet brought about novel insights into the stability of the queues in backpressure networks, supporting the usage of LIFO service priority and novel Floating Queues. The delay reductions achieved through LIFO queues can be seen in the packet delivery CDF of figure 2, provided for static network tests sourcing 0.25 packets per second for a duration of 35 minutes.

The average packet delay is nearly two orders of magnitude lower under LIFO within figure 2. We find analytically that the average delivered packet delay under FIFO service priority \bar{W}^{FIFO} is related to the average serviced packet delay under LIFO priority \bar{W}^{LIFO} like $\bar{W}^{FIFO} = \bar{W}^{LIFO} + \frac{b^{min}}{\lambda}$ for queue arrival rate λ and minimum infinitely reachable queue backlog b^{min} . In the Tutonet topology, nodes near the rear of the network running BCP have larger long term backlog $\equiv b^{min}$. Therefore, as the queue arrival rate tends to zero the delay advantage of LIFO queues becomes arbitrarily large.

Under non-static experiments, backpressure routing should be expected to outperform quasi-static tree routing algorithms. We therefore engineered an external interference experiment in which, after five minutes of source packets at 0.25 packets per second, two 802.11 devices began synchronized broadcasts at 200 packets per second with a duty cycle of 10 seconds on, 20 seconds off. As expected, the tree routing algorithm shows signs of substantial stress in this experiment, while BCP routed around the period bursts of external interference and in doing so preserved good delivery ratios and average packet transmissions throughout the test.

REFERENCES

[1] U Akyol, M Andrews, P Gupta, J Hobby, I Saniee, and A Stolyar. Joint scheduling and congestion control in mobile ad-hoc networks. *INFOCOM*, Jan 2008.

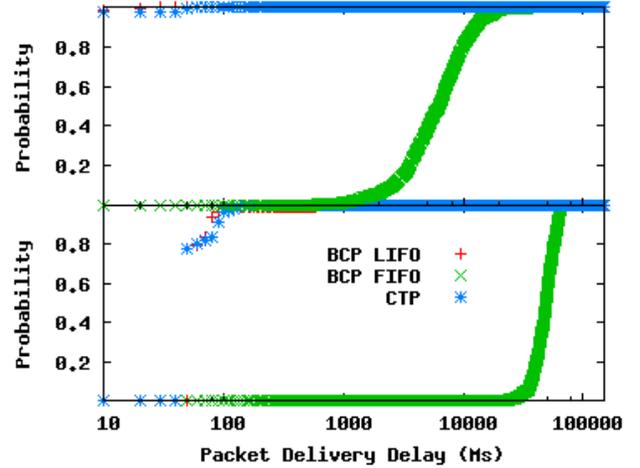


Fig. 2. Source to sink delay CDF at 0.25 PPS for motes 4 (top) and 40 (bottom) under CTP, BCP-FIFO and BCP-LIFO.

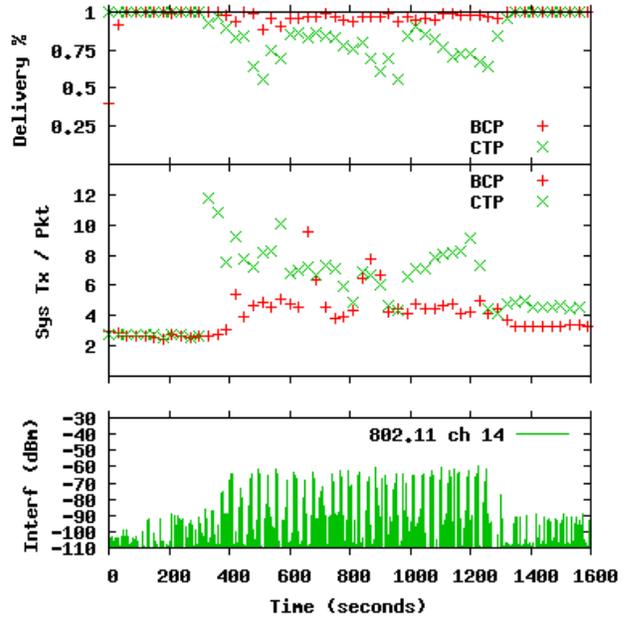


Fig. 3. Thirty second windowed average sourced packet delivery ratio (top) and system transmissions per packet (middle). Spectrum analyzer results are plotted at bottom for the colliding 802.11 channel 14 traffic.

[2] O Gnawali, R Fonseca, K Jamieson, D Moss, and P Levis. Collection tree protocol. *Sensys*, Apr 2009.

[3] M J Neely. Dynamic power allocation and routing for satellite and wireless networks with time varying channels. *PhD Thesis, Massachusetts Institute of Technolog*, 2003.

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

[5] A Sridharan, S Moeller, and B Krishnamachari. Implementing backpressure-based rate control in wireless networks. *ITA Workshop*, Sep 2008.

[6] L Tassiulas and A Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 1992.

[7] A Warrier, S Janakiraman, S Ha, and I Rhee. Diffq: Practical differential backlog congestion control for wireless networks. *INFOCOM*, Jan 2009.