

Constrained Flooding: A Robust and Efficient Routing Framework for Wireless Sensor Networks

Ying Zhang and Markus Fromherz
Palo Alto Research Center (PARC) Inc.
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
Emails: {yzhang, fromherz}@parc.com

Abstract

Flooding protocols for wireless networks in general have been shown to be very inefficient and therefore are mainly used in network initialization or route discovery and maintenance. In this paper, we propose a framework of constrained flooding protocols. The framework incorporates a reinforcement learning kernel, a differential delay mechanism, and a constrained and probabilistic retransmission policy. This type of protocol takes the advantages of robustness from flooding, but maintains energy efficiency by constraining retransmissions. Without the use of any control packets, such a protocol adapts to the specific routing requirements of the task and the dynamic changes of the network. We analyze this framework in simulation using a real-world application in sensor networks.

1. Introduction

Routing in sensor networks has very different characteristics than that in traditional communication networks. First of all, address-based destination specification is replaced by a more general feature-based specification [9], such as geographic location [11] and information gain [5], or a fixed but maybe mobile sink. Secondly, routing metrics are not just shortest delays, but also energy usage [21] and information density [5]. Thirdly, in addition to peer-to-peer communication, multicast (one-to-many) and converge-cast (many-to-one) are major traffic patterns in sensor networks. Even for peer-to-peer communication, routing is more likely to be source or sink driven than table-based [16], and source/sink pairs often are dynamic (changing from time to time) or mobile (moving during routing).

In the last several years, many routing protocols have been developed for wireless sensor networks, including Grid Routing [4], Directed Routing [13], Mint Routing [19],

Backbone Routing [7], and the Constraint-based Routing Framework [22, 25]. Most of these routing protocols have been implemented on Berkeley motes [1], a widely used sensor network platform. We may classify these routing mechanisms into two categories: *structure-based* protocols and *structure-less* protocols. Structure-based protocols generate a routing tree (e.g., Backbone [7], Mint [19], Grid [4], Adaptive Tree [25]), or one or more routing paths (e.g., AODV [14], DSR [10], Multipath routing [6]) before sending data. For dynamic networks, in addition to initialization, periodic or repair-based maintenance is required to keep the structure up to date. Structure-less protocols, instead of building and maintaining a routing structure, establish a “potential field”. Data from upstream flows downstream according to the potential field. Similar to structure-based protocols, however, the maintenance of the potential field is necessary for dynamic networks. Flooding (e.g., Directed Routing [13], Gradient Broadcast [20], Gradient Routing [15], Constrained Flooding [25]) and real-time search (e.g., Ant-Routing [3, 26], Q-Routing [2, 12], GEAR [21], LRTA* [23, 25]) protocols belong to this category.

Constrained flooding has been proposed as one of the meta-routing strategies [25] for constraint-based routing [22]. Constrained flooding consists of a real-time reinforcement kernel, shared by other meta-routing strategies [25] in constraint-based routing, and a constrained retransmission policy with differential delays. In this paper, we will present an in-depth study of the constrained flooding protocols, performance analysis and comparisons with other peer routing protocols for some real-world applications. We will show that constrained flooding protocols are not only robust with high success rates but also, unlike other flooding protocols, energy efficient.

A couple of flooding-based routing protocols have been proposed in the past, such as GRAD [15], GRAB [20] and DFRF [13]. GRAD is a basic gradient flooding protocol where all nodes with smaller cost to the sink retransmit

the data. GRAB adds credit control to the basic gradient routing which effectively reduces the degree of redundancy and energy consumption. DFRF incorporates priorities and multiple redundant transmissions to further increase success rates. DFRF has been implemented on Berkeley motes and applied to Shooter Localization [17], one of the most successful applications in sensor networks.

The major elements of the constrained flooding (CF) framework are the following:

- **Real-time learning of potential/cost field:** other than most flooding protocols where the potential field has to be established and maintained, CF incorporates real-time reinforcement learning techniques to build and update the potential field. No extra control packets are necessary.
- **Temperature control for robustness:** unlike other gradient routing, where only nodes with smaller cost retransmit the packet, CF adds an extra variable, temperature, with the intent that only nodes with cost difference smaller than the temperature participate in the relay. The temperature will be cooled down over time to reduce the number of transmissions and increase energy efficiency.
- **Delayed transmission according to the gradient descent:** instead of adding random delays to retransmission packets to reduce broadcast collision, CF adds delays according to the cost difference. The smaller the cost, the sooner the retransmission.
- **Probabilistic retransmission based on data frequency:** the node will retransmit the delayed packet based on the number of times the same packet has been heard. The more times the packets has been heard, the less likely that it will be retransmitted by this node. This way, packets at nodes with larger cost (with larger delays) are likely to be suppressed, significantly reducing the number of transmissions.

Note that not all elements are novel and some of them (e.g. delayed, probabilistic and counter-based retransmission) have been proposed before [18]. CF integrates constraint-based message specification with reinforcement learning strategies and applies them effectively to control the broadcast storm. CF can also incorporate credit control from GRAB and priority management from DFRF in a layered routing architecture. Furthermore, CF is a meta-routing strategy that can be used with general routing specifications, of which geographical or energy-aware routings are canonical examples.

The remainder of the paper is organized as follows. Section 2 presents the constrained flooding framework and its main components. Section 3 analyzes the protocols in this framework via simulation and compares performances for a real-world application. Section 4 concludes the paper.

2. Constrained Flooding

In this section, we will first present the basic elements of constrained flooding protocols and then introduce a layered routing architecture with the components for implementing constrained flooding protocols.

2.1. Constrained Flooding Protocols

The basic elements of constrained flooding include (1) a real-time reinforcement learning kernel for establishing and maintaining the potential field, (2) a constrained propagation policy, (3) differential delays based on cost difference, and (4) probabilistic retransmission based on data frequency.

2.1.1. Potential field generation: The potential field of a network is the mapping from nodes to cost-to-go values. The definition of the cost-to-go depends on routing objectives. For example, if the shortest path is the routing objective, cost-to-go is the minimum number of hops from the current node to the sink node. In general, if the local cost function is o , the cost-to-go at node v^0 is $\min_p \sum_i o(v^i)$ where $p : v^0, v^1, \dots, v^{n-1}$ is a path from v^0 to the sink node.

Most flooding protocols generate the potential field by flooding from the sink node. Each node maintains the estimated cost-to-go c . The sink has zero cost-to-go and all the other nodes have infinite initial values. Starting from the sink node, a packet is sent out advertising its cost-to-go. Whenever a node v receives a packet from a neighbor with cost-to-go as c' , it calculates its new cost-to-go as $\min(c, o(v) + c')$ and sends a packet with its cost-to-go. After a complete flooding from the sink node, every node will have a cost-to-go value and the potential field is established. There are two problems with this method. First of all, if the network changes, one has to regenerate the potential field by flooding from the sink. For the sink to recognize the changes in the network, network performance metrics, e.g., packet loss rates, have to be maintained. If the network is highly dynamic, this method would have both high energy cost and high loss rates. Periodic flooding can be performed as well, again, with the extra energy cost. Secondly, if the sink is not fixed, for example, in the case of geographical routing, one has to locate the sink first through an initial flooding.

For real-time reinforcement learning, in particular Q-learning, each node maintains not only its estimated cost-to-go, but also cost-to-go values of its neighbors. Initially all the cost-to-go values are set to zero, if there is no knowledge about the sink. Every packet (not just the control packets) has a piggybacked cost-to-go field and all the nodes (excluding the ones in sleep state) are set to be in promiscuous listening mode. Whenever a non-sink node v over-

hears a packet, whether it is the designated receiver or not, it updates the corresponding cost-to-go value of the neighbor where the packet was sent from, and re-estimates its own cost-to-go using the formula

$$c \leftarrow (1 - \alpha)c + \alpha(o(v) + \min_w c'(w)) \quad (1)$$

where $0 < \alpha \leq 1$ is the learning rate, w is a neighbor of this node, and c' is the neighbor's cost-to-go. There are two advantages with this method. First of all, initial flooding from the sink is useful but unnecessary, since packets from any direction help establish the potential field. Similarly, when the network changes, no flooding from the sink is needed, ongoing packets in the network will update the potential field. Secondly, if one has knowledge about the sink, cost-to-go can be estimated initially rather than set to zero uniformly. For example, in the case of geographical routing with the shortest path objective, the cost-to-go at any node can be estimated by d/R , where d is the Euclidean distance between the node and the sink, and R is the maximum radio range. When there are no network holes, this estimation of the potential field can be accurate enough for efficient geographical flooding. When there are network holes, the potential field would be automatically updated during routing.

2.1.2. Constrained propagation: All flooding protocols piggyback the cost-to-go value in each data packet, and most of them use a simple rule to forward packets: if a node receives a packet with cost-to-go higher than its own, it retransmit the packet. This simple strategy works if all the cost-to-go values are accurate. If the cost-to-go values are not fully learned yet (e.g., there is a network hole in geographical routing) or if the network changes, this simple rule causes the loss of packets.

In constrained flooding, this simple rule is generalized using an extra variable, temperature T : if a node with cost-to-go c receives a packet with cost-to-go c' , and if $c - c' < T$, it retransmits the packet. This rule is identical to the simple rule when $T = 0$. However, in this generalized form, T can be changing during routing. For example, T can be set to large values initially and reduced gradually (cooling down) with the learning of the potential field.

The temperature cooling-down process may vary in this general framework. For example, one can reduce T after every packet received:

$$T \leftarrow k * T / (k + 1) \quad (2)$$

where $k = T_0$, the initial temperature.

2.1.3. Differential delays: Most flooding protocols add random delays to forwarding packets to avoid collision. In constrained flooding, delays are added to forwarding packets according to the cost difference $\Delta = c' - c$: the larger the difference, the smaller the delay. The delay function δ

has the property that if $\Delta_1 \geq \Delta_2$, $\delta(\Delta_1) \leq \delta(\Delta_2)$. For example, the delay function

$$\delta(\Delta) = D/e^\Delta \quad (3)$$

where D is a constant, works quite well.

2.1.4. Probabilistic retransmissions: For most flooding protocols, all the forwarding packets will be transmitted after random delays. For constrained flooding, after the designated delay time, a probabilistic policy will be used to decide whether or not to retransmit the packet: the more times the same packet has been heard, the less possible the packet will be retransmitted. Let p be a probability function for retransmission and C denote the number of times the packet has been heard in the past, then p should satisfy (1) $p(1) = 1$, i.e., if only one such a packet is heard, it should be transmitted, and (2) $p(C_1) \leq p(C_2)$ if $C_1 \geq C_2$. An example of such a function is

$$p(C) = 1/C^\gamma \quad (4)$$

where $\gamma \geq 0$. Here γ can be tuned to trade off between robustness and energy cost. The smaller γ , the more packets are to be transmitted and the higher the energy cost.

2.1.5. Constrained flooding protocols: Like most flooding protocols, in addition to a data field, constrained flooding protocols have the cost-to-go field and a unique identifier, which can be a sequence ID or a time stamp at the source. The cost-to-go is used to generate and update the potential field as well as to control the direction of flooding, and the unique identifier is used to check duplicates and to count the number of times such a packet has been heard. Furthermore, constrained flooding protocols have each sender's (including relay nodes) address in any packet, so that neighborhood structure can be maintained.

A flowchart for constrained flooding protocols is shown in Figure 1. Module 1 updates the temperature (e.g., using Eq. (2)) and its cost-to-go as per Eq. (1), module 2 determines if the packet shall be forwarded based on the cost difference and the current temperature at this node, module 3 adds delay to the forward packet (e.g., according to delay function Eq. (3)), and module 4 makes the final decision (e.g., according to probabilistic function Eq. (4)) of sending the packet after the designated delay time.

2.2. Constrained Flooding Components

The constrained flooding framework can also incorporate mechanisms proposed by other flooding protocols, such as credit control [20], priority management for duplicated transmissions [13], and packet aggregations [13]. On the other hand, components that are used in constrained flooding, such as delayed and probabilistic retransmission, can be applied to any flooding-based routing. In this section,

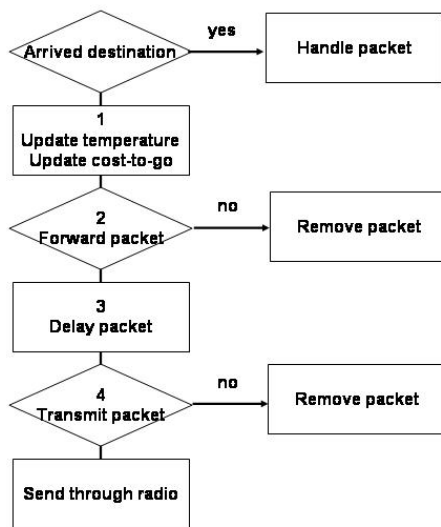


Figure 1. Flowchart for constrained flooding protocols

we advocate a layered routing architecture, then describe the essential components for constrained flooding as well as other components that can be used within the framework.

2.2.1. Layered routing architecture: A routing component is a module which handles events (e.g., *Packet_Received*, *Packet_Sent*, *Clock_Tick*) and executes commands (e.g., *Send_Packet*). A routing component A is wired to another routing component B , denoted $A \rightarrow B$, if A passes commands to B and B passes events to A . The minimum configuration of layers includes the MAC layer, a routing layer, and the application layer, with the MAC layer at the bottom and the application layer at the top. The MAC layer performs the low-level communication protocol (i.e., *GenericComm* in *TinyOS*). The routing layer forwards the received packet (via broadcast or selectively to the next hop) or passes the received packet up if it is the destination of that packet. The application layer generates the application scenarios. Furthermore, one can have a set of *control layers* between the routing layer and the application layer, and a set of *support layers* between the MAC layer and the routing layer. The control layers are for network initialization and maintenance, e.g., building and maintaining a routing table or a spanning tree. The support layers are for adding features such as transmitting and receiving queues, data aggregation/fragmentation, neighborhood management, confirming/delaying transmissions, and checking duplications. It is the algorithm designer's choice to put individual functions into different layers so that common functions can be

shared by different algorithms. For example, many algorithms need neighborhood management and transmission queues, and flooding-based routing may add transmission delays to forwarding packets to reduce collisions. For a given application scenario, one should select the right algorithm with a set of control and support routing components.

2.2.2. Constrained flooding layers: The basic constrained flooding protocols have the following layers: *application* \rightarrow *constrained_flood* \rightarrow *delay_transmit* \rightarrow *neighborhood* \rightarrow *MAC*, where *constrained_flood* is the routing layer, *delay_transmit* and *neighborhood* are support layers.

- *constrained_flood*: This layer implements the basic functions of modules 1, 2 and 3 in Figure 1. It maintains the cost-to-go of this node as well as the cost-to-go values of its neighbors. The value of cost-to-go at this node will be added to each packet passing this layer, which will be transmitted with the packet. Delays are calculated for forwarding packets.
- *delay_transmit*: This layer implements the function of module 4 in Figure 1. It maintains a timeout queue and the counts for all the packets it has heard. After the timeout for each packet, the probabilistic transmission rule is applied.
- *neighborhood*: Each packet sent from this layer will piggyback the node's address, which will be retrieved from the packet received at this layer. A list of neighbors is maintained.

Using the layered architecture, one can incorporate other routing components into the constrained flooding framework. These components include *initialization*, *credit_control*, *duplicate_transmission*, and *aggregate_queue*.

- *initialization*: Initial flooding from the sink node generates an initial potential field effectively.
- *credit_control*: Credit control for flooding has been used in GRAB [20] to control the width of flooding, trading off energy usage and success rates.
- *duplicate_transmission*: Duplicated transmission using priority management has been implemented in DFRF [13] to increase packet success rates.
- *aggregate_queue*: Packet aggregation assembles a number of packets into one packet before sending and disassembles a packet into individual packets after receiving. In the shooter localization application [17], maximum four packets are aggregated into one packet.

3. Performance Evaluations

In this section, we evaluate the routing performance of the constrained flooding framework for a real-world example: shooter localization [17], a successful sensor network application. For this application, all the acoustic data that are above certain threshold have to be sent to the base station, so that the source of the sound can be located. The communication problem in shooter localization is a typical convergecast [8]. In this experiment, the sensor distribution, the sink location and the event trace data are obtained from the actual experiments on the hardware platform. The network is about 10 hops (Figure 2).




Figure 2. Node distribution in the shooter localization experiment

Three algorithms are compared for this scenario:

- *Directed Routing (DFRF)* [13] – this algorithm is actually used in the shooter localization application, implemented in TinyOS/NesC; DFRF integrates both duplicate transmissions (up to 3 times) and packet aggregations (up to 4 packets).
- *Backbone tree (Backbone)* [7] – this algorithm uses directed diffusion [9] to create a backbone tree in the initialization phase and passes packets to parents during routing. To handle the problem of asymmetric links, it establishes a symmetric link neighborhood in the initialization;
- *Constrained Flooding (CF)* – basic CF with initialization;
- *Constrained Flooding with Aggregation and Duplication (CF+Agg+Dup)* – basic CF with initialization, aggregation and duplication components.

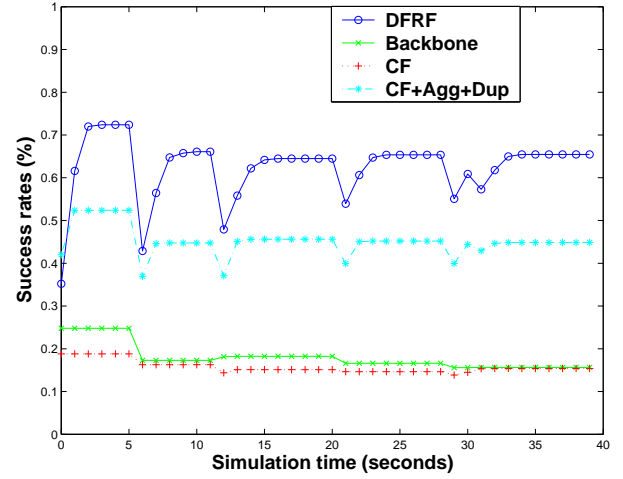


Figure 3. Shooter Localization: Success rates




Figure 4. Shooter Localization: Energy efficiency

Two performance metrics are used for comparisons: *success rate* measures the total number of packets received at the sink vs. the total number of packets sent from the sources. Energy efficiency measures the total number of packets received at the sink vs. the total number of packets sent in the air. *Rmase* [24], a sensor network simulator, is used for simulation. Results are based on the average of 10 random runs. Figures 3 and 4 show success rates and energy efficiency, respectively. We see that DFRF has the highest success rates. CF+Agg+Dup increases the success rates and has the highest energy efficiency. We have also tried the augmentation of aggregation and duplication to the Backbone layers but it does not seem to increase success rates either.

4. Conclusion

We have presented in this paper the framework of constrained flooding, its basic elements with augmented utilities in layered routing architecture. We have shown in our experiments that CF can be both robust and energy efficient, with the added utility layers, such as initialization, credit control, aggregation and duplication. Lots of research still need to be done on the selection of parameters and routing components and understanding the relationship between different components.

Acknowledgement: This work is funded in part by the Defense Advanced Research Project Agency contract # F33615-01-C-1904. Thanks to anonymous referees for valuable comments.

References

- [1] U. Berkeley. Berkeley wireless embedded systems. <http://webs.cs.berkeley.edu/>.
- [2] J. A. Boyan and M. L. Littman. Packet routing in dynamically changing networks: A reinforcement learning approach. In J. A. J. D. Crowan, G. Tesauro, editor, *Advances in Neural Information Processing Systems*, volume 6, pages 671–678. Morgan Kaufmann Publishers, Inc., 1994.
- [3] G. D. Caro and M. Dorigo. Antnet: Distributed stigmergetic control for communication networks. *Journal of Artificial Research*, 9:317 – 365, 1998.
- [4] Y. Choi, M. G. Gouda, H. Zhang, and A. Arora. Routing on a Logical Grid in Sensor Networks. Technical Report TR04-49, Department of Computer Sciences, The University of Texas at Austin, 2004.
- [5] M. Chu, H. Haussecker, and F. Zhao. Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. *Int. Journal on High Performance Computing Applications*, June 2002.
- [6] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient energy-efficient multipath routing in wireless sensor networks. *Mobile Computing and Communications Review*, 1(2), 2002.
- [7] T. He, S. Krishnamurthy, J. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh. An energy-efficient surveillance system using wireless sensor networks. In *Proc. 2nd International Conference on Mobile Systems, Applications and Services (MobiSys04)*, 2004.
- [8] Q. Huang and Y. Zhang. Radial coordination for convergecast in wireless sensor networks. In *Proc. IEEE 1st workshop on Embedded Networked Sensors (EmNeTS-I)*, 2004.
- [9] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proc. 6th Int'l Conf. on Mobile Computing and Networks (ACM Mobicom)*, Boston, MA, 2000.
- [10] D. B. Johnson and D. B. Maltz. Dynamic source routing in ad hoc wireless networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*, pages 153–181. Kluwer Academic, 1996.
- [11] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. 6th Int'l Conf. on Mobile Computing and Networks (ACM Mobicom)*, Boston, MA, 2000.
- [12] S. Kumar and R. Miiikkulainen. Confidence-based q-routing: An on-line adaptive network routing algorithm. In *Proceedings of Artificial Neural Networks in Engineering*. 1998.
- [13] M. Maroti. Directed Flood-Routing Framework. Technical Report ISIS-04-502, Institute for Software Integrated Systems, Vanderbilt University, 2004.
- [14] C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *Proc. 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, February 1999.
- [15] R. Poo. Gradient routing in ad hoc networks. <http://www.media.mit.edu/pia/Research/ESP/texts/poorieeepaper.pdf>.
- [16] E. Royer and C. Toh. A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications*, April 1999.
- [17] G. Simon, M. Maroti, A. Ledeczi, G. Balogh, B. Kusy, and A. Nadas. Sensor network-based countersniper system. In *Proc. 2nd ACM International Conference on Embedded Networked Sensor Systems (SenSys04)*, 2004.
- [18] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. *ACM Wireless Networks*, 8(2), March 2002.
- [19] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proc. 1st ACM International Conference on Embedded Networked Sensor Systems (SenSys03)*, 2003.
- [20] F. Ye, G. Zhong, S. Lu, and L. Zhang. A robust data delivery protocol for large scale sensor networks. In *Information Processing in Sensor Networks, Lecture Notes in Computer Science*, number 2634, page 658, 2003.
- [21] Y. Yu, R. Govindan, and D. Estrin. Geographical and energy aware routing: a recursive data dissemination protocol for wireless sensor networks. Technical report ucla/csd-tr-01-0023, UCLA Computer Science Department, May 2001.
- [22] Y. Zhang and M. Fromherz. Message-initiated constraint-based routing for wireless ad-hoc sensor networks. In *Proc. IEEE Consumer Communication and Networking Conference*, 2004.
- [23] Y. Zhang and M. Fromherz. Search-based adaptive routing strategies for sensor networks. In *AAAI04 workshop on Sensor Networks*, 2004.
- [24] Y. Zhang, M. Fromherz, and L. Kuhn. Rmase: Routing modeling application simulation environment. <http://www.parc.com/era/nest/Rmase>.
- [25] Y. Zhang, M. Fromherz, and L. Kuhn. Smart routing with learning-based qos-aware meta-strategies. In *Proc. Quality of Service in the Emerging Networking*, Lecture Notes in Computer Science 3266, 2004.
- [26] Y. Zhang, L. Kuhn, and M. Fromherz. Improvements on ant-routing for sensor networks. In *Fourth International Workshop on Ant Colony Optimization and Swarm Intelligence*, Lecture Notes in Computer Science 3172, 2004.