

Near-Optimal Data Dissemination Policies for Multi-Channel, Single Radio Wireless Sensor Networks

David Starobinski* Weiyao Xiao* Xiangping Qin† Ari Trachtenberg*

* Department of Electrical and Computer Engineering
Boston University, Boston, MA 02215
Email: {staro, weiyao, trachten}@bu.edu

† Samsung Information Systems America
75 West Plumeria Drive San Jose, CA 95134
Email: x.qin@samsung.com

Abstract

We analyze the performance limits of data dissemination with multi-channel, single radio sensors. We formulate the problem of minimizing the average delay of data dissemination as a stochastic shortest path problem and show that, for an arbitrary topology network, an optimal control policy can be found in a finite number of steps, using value iteration or Dijkstra's algorithm. However, the computational complexity of this solution is generally prohibitive. We thus focus on two special classes of network topologies of practical interest, namely single-hop clusters and multi-hop cluster trees. For these topologies, we derive the structure of policies that achieve an average delay within a factor $1 + \epsilon$ of the optimal average delay, in networks with large number of nodes. Through simulation, we show that these policies perform close to optimal even for networks with small and moderate numbers of nodes. Our analysis and simulations reveal that multi-channel data dissemination policies lead to a drastic reduction in the average delay, up to a factor as large as the total number of channels available, even though each node can communicate over only one channel at any point of time. Finally, we present the foundations of a methodology, based on extreme value theory, allowing the implementation of our near-optimal dissemination policies with minimal overhead.

A version of this paper appeared as:

- D. Starobinski, W. Xiao, X. Qin, and A. Trachtenberg, "Near-Optimal Data Dissemination Policies for Multi-Channel, Single Radio Wireless Sensor Networks", IEEE INFOCOM 2007.

I. INTRODUCTION

A wide variety of fundamental sensor networking services, such as routing and over-the-air programming, rely upon efficient data dissemination [1]–[3]. However, traditional sensor limitations (e.g., limited battery life and memory) joined with the complications of the wireless sensor broadcast channel (e.g., lossy channel, narrow bands, and energy-expensive communication) make it extremely difficult to analyze and implement efficient dissemination algorithms.

Even so, wireless sensor radios currently on the market do enjoy at least one currently under-utilized feature: they are able to communicate on any one of multiple (narrow) channels [4]–[7]. Thus, for example, MICA2 sensor motes operating in the 900 Mhz range can communicate on any one of more than 25 non-overlapping channels. The main challenge, however, is that the motes are equipped with a single radio interface (due, in large part, to considerations of cost and energy consumption) and, thus, can operate on only one of these channels at a time. The main effort of this work is to demonstrate, through theoretical analysis and simulation, that this multi-channel transceiving capability of sensor motes can be exploited for major efficiency gains.

Specifically, we propose a theoretical framework to evaluate the performance limits of data dissemination with multi-channel, single radio sensors, using expected delay as the primary optimization metric. Within this framework, we show how to model the problem of data dissemination as an instance of the stochastic shortest path problem [8]. This framework permits us to find an optimal dissemination policy for an arbitrary topology in a finite (though possibly prohibitive) amount of time using value iteration or Dijkstra’s algorithm.

The optimal solutions are typically very complicated and unintuitive. As such, we focus on two specific classes of topologies of practical interest: single-hop clusters and multi-hop cluster trees [9], [10]. For large size networks with these topologies, we are able to derive the structure of policies that exhibit nearly optimal expected delay (within a factor $(1+\epsilon)$, for any $\epsilon > 0$). These policies make use of a round-robin strategy applied both at the packet and channel levels. We, thus, refer to them as packet-channel round robin (PCRR) policies. One of our main theoretical contributions is to show that, with C channels available, the expected delay using PCRR can approach a value that is C times smaller than the optimal expected delay with a single channel (i.e., as in a multi-radio system). These results are validated by simulation, showing that PCRR policies are nearly optimal even for small and moderately-sized networks.

Our work provides a first step in rigorously characterizing the performance limits of multi-channel, single radio wireless sensor networks. As such, important practical considerations, such as control overhead, are not explicitly captured in the model. That said, we present a methodology, based on extreme value theory [11], that provides foundations for the practical implementation of our near-optimal dissemination policies. Specifically, given an upper bound estimate on the packet loss probability, we show that the PCRR policy can be implemented in such a way that, with very high probability, all the nodes in the network receive all the information without having to send any acknowledgement packets (a so-called ACK-less protocol). Our approach, thus, suggests a novel way to address the well-know “broadcast storm” problem [12] plaguing reliable data dissemination in wireless networks.

The rest of this paper is organized as follows. In Section II, we first present our model and then formalize our optimization problem, analyze its computational complexity, and show how it can be (theoretically) solved in finite time. In Section III, we analyze the problem of data dissemination in single cluster topologies, introduce the PCRR policy, and prove its near-optimality in networks with large number of nodes. In Section IV, we generalize our results to multi-hop cluster tree topologies. In Section V, we present our approach for developing ACK-less protocols based on extreme value theory. Simulation results are presented in Section VI. We provide concluding remarks in Section VII.

II. MODEL AND PROBLEM FORMULATION

A. Model

We consider the problem of disseminating a file consisting of M packets from a set of S sources (e.g., base stations) to N nodes in an arbitrary topology network, with C orthogonal channels available for communication. Each source has a copy of the entire file. The time axis is slotted and each packet transmission takes one time slot. Each node is equipped with a single, half-duplex radio. Thus, during a time slot, a node can either transmit or receive (but not both) on one of the C channels. To simplify exposition, we assume that packets do not need to be received in order at the various nodes for a file to be properly reconstructed although the results of Sections II, III and IV hold without this assumption. Note that several data dissemination protocols, such as Deluge [1], do not require that packets be received in order.

At each time slot, a control u specifies which nodes transmit and receive on each channel. Packets are not only transmitted by sources but possibly also by other nodes that have received some of the packets and serve as relays. Communications take place over a wireless broadcast channel, whose losses are independent and identically distributed at each time slot. As such, we can associate a probability $p_{ij}(u)$ of a packet transmission from node i to j being corrupted; note that this probability is a function of the control u because the packet loss is dependent on all simultaneous transmissions on the same channel.

Finally, we will denote by T the random variable representing the time (delay) until all nodes receive all packets. Our goal is thus to determine a control policy that minimizes the expected value of T (denoted \bar{T}).

B. Problem Formulation for General Networks

We next formalize our optimization problem and provide a computational methodology to solve it by casting it as a stochastic shortest path (SSP) problem [8]. In our specific case, this problem can then be solved deterministically in bounded time using value iteration or Dijkstra's algorithm [13].

a) SSP problem: The SSP problem is a generalization of the deterministic shortest path problem in a graph. Specifically, in the stochastic version, a path from a source to destination is determined probabilistically, meaning that one may transition from a vertex to any other vertex according to a given distribution (which, in turn, is determined by a chosen control). The shortest path, in this context, corresponds to the choice of controls at each vertex that minimize the expected cost to a given destination (or *termination state* in the literature). Clearly, the deterministic version of the problem thus corresponds to a case in which, controls from any given vertex assign a probability 1 for reaching some vertex and 0 for reaching all others.

b) Formulation: In our case, we build a graph of $|V| = 2^{NM}$ vertices, each of which correspond to an $N \times M$ binary matrix representing a possible configuration of the network in the middle of a data dissemination protocol. Specifically, the (n, m) -th entry of any such matrix is 1 if and only if node n has received packet m in the corresponding configuration. For simplicity, we order the states so that the initial state $i = 1$ and last state $i = |V|$ correspond to the all zero and all one matrices respectively. Note that this formulation implicitly requires global knowledge at each node about the data received at other nodes in the network; we will see how to remove this assumption in Section V for practical implementation.

To complete the model, we assume a set of possible controls $U(i)$ for each state i , and define a corresponding transition probability $q_{ij}(u)$ corresponding to the probability of reaching state j from state i if transmissions are enacted according to control $u \in U(i)$. Our goal is then to pick a control $\pi(i) \in U(i)$ at every state so as to minimize the expected delay from state 1 to the termination state $|V|$. This optimal control π will necessarily be stationary because the channel is assumed to be *i.i.d.*

c) *Solution*: Let denote by $T_{DP}^*(i)$ the time to reach the termination state starting from state i and $\bar{T}_{DP}^*(i)$ its expected value solved using dynamic programming. One of the main results for the SSP problem is that it has a unique solution satisfying Bellman's equations [8], [14]:

$$\bar{T}_{DP}^*(i) = \min_{u \in U(i)} \left[1 + \sum_{j=1}^{|V|} q_{ij}(u) \bar{T}_{DP}^*(j) \right], i = 1, \dots, |V|. \quad (1)$$

For each state i , the optimal policy $\pi(i)$ corresponds to the control that achieves the minimum in Eq. 1. To simplify notation, we will use \bar{T}^* to represent the optimal expected delay starting from initial state 1. Traditional approaches to this solution include *value iteration*, the most commonly used approach that generally requires an infinite number of iterations for convergence, and *policy iteration*, which is more computationally expensive at each step but terminates in finite time.

The special structure of our problem allows for an especially efficient value iteration solution satisfying (1). Specifically, our graph has acyclic transition probabilities, in that a path can never visit the same state twice, except for self-transitions (which can be eliminated [8, Vol. 2: p. 97]); this is because nodes cannot lose packets that they correctly received and decoded. As such, value iteration is guaranteed to converge within $|V|$ iterations, where V is the state space. Since each iteration involves $|V||U|$ operations, where $|U|$ is the size of the control space, the optimal solution can be computed with at most $|V|^2|U|$ operations. The following theorem summarizes this result.

Theorem 1: The optimal dissemination of an M -packet file to a network with N nodes, S sources, and C channels can be computed using value iteration with at most $|V|^2|U|$ operations, where $|V| = 2^{NM}$ and $|U| = 2^{(N+S)C}$.

Note that, in practice, we can even further reduce the computational complexity of this problem by noting that its optimal policy must be *consistently improving*, meaning that

$$q_{i,j}(\pi(i)) > 0 \quad \Rightarrow \quad \bar{T}_{DP}^*(i) > \bar{T}_{DP}^*(j).$$

We can then use Dijkstra's algorithm to compute the optimal policy and optimal expected delays [8, Vol. 2: p. 135].

C. Problem Formulation for Single Clusters and Cluster Trees

In general, the derivation of an optimal policy for our problem is computationally prohibitive for large numbers of nodes, packets, or channels. As such, the remainder of this work focuses on analytical approaches for some specific topologies of practical interest, such as single-hop clusters and (multi-hop) cluster trees. Even in such cases, the optimal policies can be quite involved, but we are able to derive near-optimal policies, as $N \rightarrow \infty$.

1) *Single Clusters*: In this model, the network consists of a single-hop topology (shown in Figure 1). On each channel, only one node can transmit (broadcast) in each time slot (to avoid packets collision). For each receiver, the packet loss probability is p , independently of any other events.

2) *Cluster Trees*: This model corresponds to a collection of clusters organized along a tree (see Figure 2). The height of the tree is D . The number of clusters at each depth d is denoted by W_d , where $d = 0, 1, \dots, (D - 1)$. The number of nodes in the cluster w at depth d is $N_{dw} = \alpha_{dw}N$, where $\sum_{d=0}^{D-1} \sum_{w=0}^{W_d-1} \alpha_{dw} = 1$.

We assume that the root cluster contains all the sources S . A node in a cluster can communicate directly with all the other nodes in the same cluster as well as with nodes belonging to a direct parent or child cluster. We assume a 2-hop interference model. Thus if a node transmits on a certain channel, then all the nodes belonging to any ancestor or descendant clusters within a distance of two hops and operating on the same channel must remain quiet to avoid packet collision. As before, the packet loss probability is p for each pair of source/destination.




Fig. 1. Single cluster topology




Fig. 2. Cluster tree topology

III. ANALYSIS OF SINGLE CLUSTERS

In this section, we analyze the single cluster model described in Section II-C. We first consider the single channel case. We determine the optimal policy and derive tight bounds on the expected delay. We then consider the multi-channel case. We derive lower and upper bounds on the expected delay and determine a policy that is asymptotically optimal as $N \rightarrow \infty$. These results are first proven for the case when the number of sources is larger or equal to the number of channels, i.e., $S \geq C$, and then extended to the case $S < C$. In the following, we denote by $T_N^*(C)$ the random variable representing the completion time using the optimal policy in a cluster of N nodes with C channels available for communication.

A. Single Channel

With a single channel, only one source can transmit in each time slot. Denote by T_n^m the number of slots needed for user n to receive packet m . Under the assumption of a packet dropping probability p , T_n^m has a geometric distribution with mean $1/(1-p)$. Thus, because of the broadcast nature of the channel, the number of slots needed for all the users to receive packet m is $T^m = \max_{n=1, \dots, N} T_n^m$. The minimum number of time slots to complete the transmission of all the packets is thus $T_N^*(1) = \sum_{m=1}^M T^m$. We note that the order in which the source transmits the packets is arbitrary. Thus, any transmission policy in which, in each time slot, the source transmits a packet needed by at least one of the N nodes in the network is optimal.

We next provide lower and upper bounds on the optimal expected delay.

Proposition 1: Consider a single cluster topology with $C = 1$ channel. Then,

$$M\lambda^{-1}(\log(N) + \gamma) \leq \bar{T}_N^*(1) \leq M(\lambda^{-1}(\log(N) + \frac{1}{2N} + \gamma) + 1),$$

where $\lambda = \log(\frac{1}{p})$ and $\gamma \approx 0.577215$ is the Euler's constant.

proof: For the random variables T_n^m , we have $\Pr\{T_n^m = i\} = p^{i-1}(1-p)$ and $\Pr\{T_n^m \geq i\} = \sum_{j=i}^{\infty} p^{j-1}(1-p) = p^{i-1}$. Consider an equivalent continuous random variable X_n^m , where X_n^m has the pdf $f(x) = \sum_{i=1}^{\infty} p^{i-1}(1-p)\delta(x-i)$, where $\delta(x)$ is the Dirac's delta function. The ccdf of X_n^m is $\bar{F}(x) = p^{\lceil x \rceil - 1}$, for $x \geq 0$.

We now define the random variables Y_n^m and $Z_n^m = Y_n^m + 1$ where $\bar{F}_{Y_n^m}(x) = p^x$ and $\bar{F}_{Z_n^m}(x) = \min(1, p^{x-1})$. We note that Y_n^m are independent exponential random variables with parameter $\lambda = \log(\frac{1}{p})$. Clearly, we have $\bar{F}_{Y_n^m}(x) \leq \bar{F}_{X_n^m}(x) \leq \bar{F}_{Z_n^m}(x)$. Therefore, $Y_n^m \leq_{st} X_n^m \leq_{st} Z_n^m$, where the notation $Y \leq_{st} X$ means that the random variable Y is stochastically smaller than the random variable X [15]. Using properties of stochastic ordering, we then obtain

$$\mathbb{E} \max_{i=1, \dots, N} Y_i^m \leq \mathbb{E} \max_{i=1, \dots, N} X_i^m \leq \mathbb{E} \max_{i=1, \dots, N} Z_i^m.$$

Since Y_n^m are independent exponential random variables, from [16, p. 73], we have

$$\mathbb{E} \max_{i=1, \dots, N} Y_i^m = \sum_{n=1}^N \frac{1}{i\lambda}.$$

Using known bounds on the harmonic sum [17],

$$0 < \frac{1}{2(N+1)} \leq \sum_{n=1}^N \frac{1}{i} - \log(N) - \gamma \leq \frac{1}{2N},$$

where $\gamma \approx 0.577215$ is the Euler's constant. Therefore, we have

$$\frac{\log(N) + \gamma}{\log(\frac{1}{p})} \leq \mathbb{E} \max_{n=1, \dots, N} X_n^m \leq \frac{\log(N) + \frac{1}{2N} + \gamma}{\log(\frac{1}{p})} + 1. \quad (2)$$

The proof of the theorem follows by noting that $\mathbb{E}T^m = \mathbb{E} \max_{n=1, \dots, N} X_n^m$ and $\bar{T}_N^*(1) = M\mathbb{E}T^m$. ■

From the proposition, we observe that $\bar{T}_N^*(1) = M \log_{1/p}(N) + o(\log(N))$.

B. Multiple Channels

We now address the case $C > 1$. In general, the structure of the optimal policy appears to be quite intricate. One exception is the unconstrained channel case, that is, $C \geq M$, in which a distinct channel can be dedicated to the transmission of each packet. Note that $T_N^*(C) = T_N^*(M)$, when $C \geq M$.

We will exploit the results obtained for the single channel and unconstrained channel cases to provide bounds on the optimal expected delay for the case $1 < C < M$. Using these bounds, we show that, as $N \rightarrow \infty$, a simple scheduling policy, called Packet-Channel Round Robin, achieves an expected delay that is within a multiplicative factor of $(1 + \epsilon)$ of the optimal expected delay, where ϵ is an arbitrary small positive constant. The results in this section are first derived for the case when the number of sources is larger or equal to the number of channels, i.e., $S \geq C$ (or $S \geq M$ if $C \geq M$). At the end of the section, we show that they can be easily extended to the case $S < C$.

a) Unconstrained channels: Suppose $C \geq M$. Then, a different packet can be continuously transmitted on each channel, e.g., packet 1 on channel 1, packet 2 on channel 2, etc. As in the single channel case, the number of slots needed for user n to receive packet m , denoted by T_n^m , is geometrically distributed with parameter p . However, this time, as soon as user n receives packet m , it can switch to another channel to receive the next packet it needs, and so on until all the packets are received. The time needed for user n to receive all the packets is, thus, $T_n = \sum_{m=1}^M T_n^m$ and the completion time for all the nodes is $T_N^*(M) = \max_{n=1, \dots, N} T_n$. Any control policy is optimal as long as nodes tune to channels on which they can receive a new packet.

We next provide bounds on the optimal expected delay. We note that, being the sum of independent geometric random variables, T_n is a Pascal (or negative binomial) random variable with parameters p and M . We can then use a procedure similar to the proof of Theorem 1. We replace T_n by an equivalent continuous random variable and stochastically bound it.

Proposition 2: Consider a single cluster topology with $C \geq M$ channels and $S \geq M$ sources. Then,

$$\begin{aligned}\bar{T}_N^*(M) &\geq \lambda^{-1}(\log(N) + \log(M) + \gamma) \quad \text{and,} \\ \bar{T}_N^*(M) &\leq (\lambda^{-1} + \epsilon)(\log(N) + \frac{1}{2N} + \gamma) + f(\epsilon),\end{aligned}$$

where $\lambda = \log(1/p)$, $\gamma \approx 0.577215$ is the Euler's constant, and for any $\epsilon > 0$, $f(\epsilon)$ is a finite function, which is independent of N .

Proof: We first prove the first inequality. Since $T_N^*(M) = \max_{n=1, \dots, N} \sum_{m=1}^M T_n^m$, we have

$$T_N^*(M) \geq \max_{n=1, \dots, N} \max_{m=1, \dots, M} T_n^m. \quad (3)$$

The expression in the rhs of Eq. (3) corresponds to the maximum of NM independent random variables geometrically distributed with mean $1/(1-p)$. Therefore, the first inequality follows from Eq. (2) (replacing N by NM).

For the second inequality, similar to the proof in Proposition 1, we use the random variables X_n^m , Y_n^m and Z_n^m . Let $X_n = \sum_{m=1}^M X_n^m$, $Y_n = \sum_{m=1}^M Y_n^m$ and $Z_n = \sum_{m=1}^M Z_n^m$. Using stochastic ordering properties, we have

$$\bar{F}_{Y_n}(x) \leq \bar{F}_{X_n}(x) \leq \bar{F}_{Z_n}(x).$$

Now, we note that Y_n is an Erlang random variable, that is,

$$\bar{F}_{Y_n}(x) = \sum_{i=0}^{M-1} \frac{(\lambda x)^i}{i!} e^{-\lambda x}.$$

Now, for any $\epsilon' > 0$, $e^{\epsilon'x}$ increases at a faster rate than $\sum_{i=0}^{M-1} \frac{(\lambda x)^i}{i!}$ as $x \rightarrow \infty$. Thus there exists a constant $x_{\epsilon'}$, such that for all $x \geq x_{\epsilon'}$, $e^{\epsilon'x} \geq \sum_{i=0}^{M-1} \frac{(\lambda x)^i}{i!}$, namely $\bar{F}_{Y_n}(x) \leq e^{-(\lambda - \epsilon')x}$ for all $x \geq x_{\epsilon'}$. Construct a random variable $Y'_n = x_{\epsilon'} + Y_n^e$, where Y_n^e is exponentially distributed with rate $(\lambda - \epsilon')$. We have $\bar{F}_{Y_n}(x) \leq \bar{F}_{Y'_n}(x)$ for all x . Therefore, $Y_n \leq_{st} Y'_n$ and

$$\mathbb{E} \max_{i=1, \dots, N} Y_i \leq \mathbb{E} \max_{i=1, \dots, N} Y'_i.$$

Since $\mathbb{E} \max_{i=1, \dots, N} Y'_i = x_{\epsilon'} + \mathbb{E} \max_{i=1, \dots, N} Y_i^e$ and Y_i^e are independent exponential random variables with rate $(\lambda - \epsilon')$, we again have from Eq. (2):

$$\begin{aligned}\mathbb{E} \max_{i=1, \dots, N} Y'_i &\leq x_{\epsilon'} + \frac{\log(N) + \frac{1}{2N} + \gamma}{\lambda - \epsilon'} + 1 \\ &= \left[\frac{1}{\lambda} + \frac{\epsilon'}{\lambda(\lambda - \epsilon')} \right] [\log(N) + \frac{1}{2N} + \gamma] + x_{\epsilon'} + 1.\end{aligned}$$

Let $\epsilon = \frac{\epsilon'}{\lambda(\lambda - \epsilon')}$. Then $\epsilon' = \frac{\epsilon \lambda^2}{1 + \epsilon \lambda}$ and let $f(\epsilon) = x_{\epsilon'} + M + 1 = x_{\frac{\epsilon \lambda^2}{1 + \epsilon \lambda}} + M + 1$. We then have, $\mathbb{E} \max_{i=1, \dots, N} Y_i \leq (\lambda^{-1} + \epsilon)(\log(N) + \frac{1}{2N} + \gamma) + f(\epsilon) - M$. Since $\mathbb{E} \max_{i=1, \dots, N} Z_i = \mathbb{E} \max_{i=1, \dots, N} Y_i + M$, we get $\mathbb{E} \max_{i=1, \dots, N} Z_i \leq (\lambda^{-1} + \epsilon)(\log(N) + \frac{1}{2N} + \gamma) + f(\epsilon)$.




Fig. 3. Example: PCRR policy for $M = 3$, $C = 2$

Therefore,

$$\begin{aligned} \bar{T}_N^*(M) &= \mathbb{E} \max_{n=1, \dots, N} T_n \leq \mathbb{E} \max_{i=1, \dots, N} Z_i \\ &\leq (\lambda^{-1} + \epsilon)(\log(N) + \frac{1}{2N} + \gamma) + f(\epsilon). \end{aligned}$$

The result follows. ■

From the proposition, we conclude that $\log_{1/p}(N) + o(\log(N)) \leq \bar{T}_N^*(M) \leq (1 + \epsilon) \log_{1/p}(N) + o(\log(N))$, where ϵ is an arbitrary constant such that $\epsilon > 0$.

b) Constrained channels: Assume now $C < M$. We next show that, for the optimal control policies, the completion time with C channels is always stochastically larger than the completion time in a single channel system running C times faster.

Theorem 2: Consider a single hop cluster with $C \leq M$ channels. Then $T_N^*(C) \geq_{st} \frac{1}{C} T_N^*(1)$.

Proof: We prove this result using a sample path argument. Consider an arbitrary time slot in the system with C channels, denoted by SYS_C , assuming packet m_i is transmitted in channel i , $i = 1, \dots, C$. Denote an equivalent single channel system by SYS_1 evolving over C time slots. Each event in SYS_C is mapped to SYS_1 by having packet m_i being sent in time slot i , $i = 1, \dots, C$. Since we have single radio, a node in SYS_C can listen on one channel only, say channel j . Thus, in SYS_C , the same node listens only during slot j but is forced not to listen during all the other slots. Clearly, given the same starting state, SYS_1 using C time slots is equivalent to SYS_C using one time slot. The optimal policy using a single channel will always perform at least as well as SYS_1 , since, in practice, nodes are allowed to listen to packet transmissions in each time slot. Therefore, $C T_N^*(C) \geq_{st} T_N^*(1)$, and the theorem follows. ■

We now introduce a simple control policy called Packet-Channel Round Robin (PCRR). We will show in the sequel that this policy is near optimal (within a multiplicative factor of $1 + \epsilon$), as $N \rightarrow \infty$. To explain the policy, we introduce a few notations. We use the variable c to index channels, i.e., $c = 1, 2, \dots, C$, the variable m to index packets, i.e., $m = 1, \dots, M$, and the variable t to index time slots, i.e., $t = 1, 2, \dots, T_{PCRR}$. T_{PCRR} is a variable and represents the final time slot, that is, the time slot at which all the nodes have completed receiving all the packets. Then, the PCRR policy states that at time slot t , packet $[(C(t-1) + c - 1) \bmod (M) + 1]$ should be transmitted on channel c . An illustration of the policy for the case $C = 2$ and $M = 3$ is given in Fig. 3, the number in the block is the packet to send.

We next provide a stochastic relation between the completion time of the PCRR policy using C channels, denoted by $T_{PCRR}(C)$, and the completion time in the unconstrained channel case using the optimal policy, $T_N^*(M)$. The proof, based on a sample path argument, is similar to that of Theorem 2.

Theorem 3: Consider a single hop cluster with $C \leq M$ channels. Then $T_{PCRR}(C) \leq_{st} \frac{M}{C} T_N^*(M) + 1$.

We now prove the main result of this section, namely that the PCRR policy is $(1 + \epsilon)$ -optimal, as $N \rightarrow \infty$.

Theorem 4: Consider a single hop cluster with $C \leq M$ channels. Then

$$1 \leq \lim_{N \rightarrow \infty} \frac{\bar{T}_{PCRR}(C)}{\bar{T}_N^*(C)} \leq 1 + \epsilon,$$

for any, arbitrarily small, positive constant ϵ .

Proof: From Theorems 2 and 3, we have

$$\frac{1}{C}T_N^*(1) \leq_{st} T_N^*(C) \leq_{st} T_{PCRR}(C) \leq_{st} \frac{M}{C}T_N^*(M) + 1.$$

From stochastic ordering properties, similar inequalities apply for the expectations of the random variables. Using the results of Propositions 1 and 2, we thus have

$$\begin{aligned} \lambda^{-1} \frac{M}{C} \log(N) + o(\log(N)) &\leq \bar{T}_N^*(C) \leq \bar{T}_{PCRR}(C) \\ &\leq (\lambda^{-1} + \epsilon) \frac{M}{C} \log(N) + o(\log(N)), \end{aligned}$$

and the theorem follows. ■

From Theorem 4, we observe that, as $N \rightarrow \infty$, the expected delay using the PCRR policy with $C \leq M$ channels approach a value that is C times smaller than the optimal expected delay with a single channel. It is important to emphasize that this result holds even though each node is only equipped with a single radio.

We also note that the performance of the PCRR policy can be improved by letting it skip the transmissions of packets already received by all the nodes. We can then easily prove that PCRR is optimal in the single channel and the unconstrained channel cases. Note, however, that all the results obtained for the PCRR policy in this section hold without this requirement.

c) The case $S < C$: We now show that the results presented in the previous section can be extended to the case where $S < C$. In this case, let the control policy consist of 2 stages. During stage 1, packets will be sent out in a round-robin fashion over a single channel. Once at least $(C - S)$ nodes have received all M packets, then stage 2 starts, during which PCRR is employed. Let the time the system spends in stage 1 be $T_{s1}(N)$.

We now show that $\bar{T}_{s1}(N) = O(1) = o(\log(N))$. Consider a policy that selects a priori $(C - S)$ out of N nodes, then use a single channel policy to transmit the M packets to each of these nodes. The expected time for these nodes to receive all packets is $\bar{T}_{(C-S)}^*(1) = O(1)$. In practice, obviously $\bar{T}_{s1}(N) \leq \bar{T}_{(C-S)}^*(1)$, thanks to the broadcast property of the wireless channel. Thus, $\bar{T}_{s1}(N) = o(\log(N))$.

IV. ANALYSIS OF MULTI-HOP CLUSTER TREES

We next analyze the cluster tree model described in Section II-C. We first consider the single channel case and then extend our results to the multi-channel case. In each case, we compute lower and upper bounds on the optimal expected time and derive (near) optimal policies, as $N \rightarrow \infty$. We will denote by $\bar{T}_{CT}^*(C)$, the minimum expected time to completion of the optimal policy for a cluster tree with C channels available.

A. Single Channel

We consider the case $C = 1$ and provide a lower bound on the minimum expected delay. As before, we use the notation $\bar{T}_N^*(1)$ to represent the minimum expected delay to disseminate M packets in a single hop cluster of N nodes. We also define $\alpha = \max_{w=0,1,\dots,W_{D-1}-1}(\alpha_{(D-1)w})$. Thus, αN represents the number of nodes in the cluster of largest size at depth $D - 1$ in the tree.

Theorem 5: Consider a cluster tree topology with $C = 1$ channel. Then

$$\begin{aligned} \bar{T}_{CT}^*(1) &\geq (D - 1) + (M - 1) + \bar{T}_{\alpha N}^*(1) \\ &= M \log_{\frac{1}{\alpha}}(N) + o(\log(N)). \end{aligned}$$

proof: We analyze the expected time to complete sending all the packets to all the nodes belonging to the cluster of largest size at depth $D - 1$, denoted as cluster v . This represents a lower bound on the expected

time to transmit all the packets to all the nodes in the network. First, we note that it will take at least $D - 1$ time slots for the first packet to arrive to any cluster at depth $D - 1$. Additionally, there are at least $M - 1$ time slots during which nodes in clusters at depth $D - 2$ receive the remaining $M - 1$ packets. During these time slots, nodes belonging to children clusters at depth $D - 1$ will be unable to either transmit or receive any packet. Consider now time slots during which packets are transmitted to nodes in cluster v . The minimum expected number of such slots is $\bar{T}_{\alpha N}^*(1)$, and the result follows. ■

We will next propose a simple, asymptotically optimal policy called Multi-hop Packet-Channel Round Robin (MPCRR) that proceeds in two stages. In stage 1, we make sure that at least one node in each cluster gets all the M packets. In stage 2, the node serving as a source implements the PCRR policy described in Section III-B. There are three rules in stage 1: (i) a packet is sent out from a parent cluster if and only if there is at least one child cluster with none of the nodes having the packet; (ii) when initiating transmissions, transmissions by descendant clusters are given higher priority over those by ancestor clusters if there is a channel contention; (iii) once transmission starts, it stops only when at least one node in each child cluster has received the packet.

Denote the expected time to completion using MPCRR policy with single channel by $\bar{T}_{MPCRR}(1)$. The following theorem provides bounds on $\bar{T}_{MPCRR}(1)$.

Theorem 6: Consider cluster tree with $C = 1$ channel. Then $\bar{T}_{CT}^*(1) \leq \bar{T}_{MPCRR}(1) \leq M \log_{1/p}(N) + o(\log(N))$.

proof: In stage 1, following an analysis similar to that of Section III-B, denote the expected time of stage 1 by $\bar{T}_{s1}(N)$, it can be shown that

$$\bar{T}_{s1}(N) \leq \frac{3(M-1) + D}{\prod_{d=0}^{D-1} \prod_{w=0}^{W_d-1} (1 - p^{\alpha_{dw}N})},$$

which can be bounded by a constant K , $K = \frac{3(M-1)+D}{(1-p)^{DW}}$, where $W = \max_{d=0..D-1} W_d$.

In the second stage, we let the sources implement the PCRR described in Section III-B. We note that sources should be at least 3-hops apart to avoid interferences. However, we remind that transmissions by a source in a certain cluster can be received not only by other nodes belonging to the the same cluster but also nodes belonging to a direct parent or child cluster. Thus, in each time slot, this policy performs at least as well as the optimal policy for a single cluster of N nodes. Denote the expected completion time in stage 2 by $\bar{T}_{s2}(N)$.

Therefore,

$$\begin{aligned} \bar{T}_{CT}^*(1) &\leq \bar{T}_{MPCRR} = \bar{T}_{s1}(N) + \bar{T}_{s2}(N) \\ &\leq K + \bar{T}_N^*(1) \leq M \log_{1/p}(N) + o(\log(N)). \end{aligned}$$

■

The following result directly follows from Theorems 5 and 6.

Theorem 7: Consider a cluster tree with $C = 1$ channel. Then

$$\lim_{N \rightarrow \infty} \frac{\bar{T}_{MPCRR}(1)}{\bar{T}_{CT}^*(1)} = 1.$$

B. Multiple Channels

We next study the case of $C > 1$ channels available in a multihop tree cluster network. The results are similar to the single cluster scenario. First, we provide a lower bound:

Theorem 8: Consider a cluster tree with $C > 1$ channels. Then

$$\begin{aligned} \bar{T}_{CT}^*(C) &\geq (D-1) + \bar{T}_{\alpha N}^*(C) \\ &= \frac{M}{C} \log_{1/p}(N) + o(\log(N)). \end{aligned}$$

Next, we give an upper bound:

Theorem 9: Consider a cluster tree with $C > 1$ channels. Then $\bar{T}_{\text{CT}}^*(C) \leq \bar{T}_{\text{MPCRR}}(C) \leq (\lambda^{-1} + \epsilon) \frac{M}{C} \log(N) + o(\log(N))$.

The techniques to prove Theorem 8 and 9 are the same as those used to prove Theorems 5 and 6, respectively. The following theorem states the fact that MPCRR is $(1 + \epsilon)$ -optimal, as $N \rightarrow \infty$.

Theorem 10: Consider a cluster tree with C channels. Then

$$1 \leq \lim_{N \rightarrow \infty} \frac{\bar{T}_{\text{MPCRR}}(C)}{T_{\text{CT}}^*(C)} \leq 1 + \epsilon,$$

for any, arbitrarily small, positive constant ϵ .

V. TOWARDS AN ACK-LESS PROTOCOL FOR DATA DISSEMINATION

In this section, we develop the foundations for an ACK-less data dissemination protocol based on the PCRR policy. In order to provide for reliable communication, use of acknowledgements (ACKs) is a common way to inform a source that a packet has been received (otherwise the source must retransmit the packet). The transmission of ACK packets is susceptible to cause a large overhead. Although some mitigation approaches exist [1], this overhead can easily become intolerable in dense networks. Hence, we propose a strategy for designing an ACK-less protocol. The main idea is to have the source compute in advance the number of slots needed so that, with very high probability, all the nodes receive all the packets. Our strategy only requires the source to know the number of intended recipients and have a conservative estimate on the packet loss probability. To simplify exposition, we will discuss the single cluster case and assume $S \geq C$. However, our results can be extended to cluster trees and to the case $S < C$, using approaches similar to those presented in the previous sections.

First consider the unconstrained case, i.e., $C \geq M$. Recall T_n^m represents the time for the user n to receive the packet m successfully, and $\Pr\{T_n^m = t\} = p^{t-1}(1-p)$. $T_n = \sum_{m=1}^M T_n^m$ is the time for user n to successfully receive M packets. The completion time, using an optimal policy, for all of N nodes to receive M packets is $T_N^*(M) = \max_{n=1, \dots, N} T_n$. We next prove a result based on techniques from extreme value theory, showing that with appropriate normalization, $T_N^*(M)$ can be bounded by random variables converging to a Gumbel distribution, as $N \rightarrow \infty$.

Theorem 11: There exist \tilde{T}_l^* and $\tilde{T}_u^* = \tilde{T}_l^* + 1$, satisfying

$$\tilde{T}_l^* \leq_{st} T_N^* \leq_{st} \tilde{T}_u^*,$$

$$\lim_{N \rightarrow \infty} \Pr\{(\tilde{T}_l^* - b_N(M))/a_N(M) \leq x\} = G(x),$$

where $G(x)$ is the cdf of the Gumbel distribution, i.e., $G(x) = \exp(-\exp(-x))$ and $a_N(M) = 1/\log(\frac{1}{p})$ and $b_N(M)$ is a constant whose expression is given in the proof of the theorem below.

proof: T_n is negative binomially distributed with distribution function $\Pr\{T_n = t\} = D(t) = I(1-p; M, t-M+1)$ [18], where $I(z; a, b)$ is regularized beta function, and from [19, p. 516], we have $I(z; a, b) = 1 - \frac{(1-z)^b}{B(a, b)} \sum_{i=0}^{a-1} (-1)^i \binom{a-1}{i} \frac{(1-z)^i}{b+i}$ when a and b are integers and $B(a, b)$ is the complete beta function with $B(a, b) = \frac{(a-1)!(b-1)!}{(a+b-1)!}$ [19, p. 594, 597].

Create a continuous R.V. \tilde{T}_l^i with c.d.f. $F(x) = I(1 - p; M, x - M)$, where $x \in R$. Then

$$\begin{aligned}\bar{F}(x) &= 1 - I(1 - p; M, x - M) \\ &= \frac{p^{x-M}}{B(M, x - M)} \sum_{i=0}^{M-1} (-1)^i \binom{M-1}{i} \frac{p^i}{x - M + i} \\ &= \frac{\prod_{j=1}^M (x - j)}{(M-1)!} p^{x-M} \sum_{i=0}^{M-1} \binom{M-1}{i} \frac{(-1)^i p^i}{x - M + i}.\end{aligned}$$

Since $F(x)$ is an increasing function of x , we have

$$F(x) \leq D(\lceil x \rceil) \leq F(x + 1).$$

Let $T_l^* = \max_{i=1..N} \tilde{T}_l^i$, and $T_u^* = \max_{i=1..N} (\tilde{T}_l^i + 1) = T_l^* + 1$, then we have

$$T_l^* \leq_{st} T_N^*(M) \leq_{st} T_u^*.$$

Let $a_N(M) = 1/\log(\frac{1}{p})$, $b_N(M) = \log_{\frac{1}{p}}(N) + (M-1)\log_{\frac{1}{p}}\left(\log_{\frac{1}{p}}(N) + (M-1)\log_{\frac{1}{p}}(\log_{\frac{1}{p}}(N))\right) + (M-1)\left(\log_{\frac{1}{p}}(\frac{1-p}{p})\right) - \log_{\frac{1}{p}}[(M-1)!] + 1$, it can be shown,

$$\lim_{N \rightarrow \infty} N\bar{F}(a_N(M)x + b_N(M)) = -\log G(x), \quad x \in R.$$

Thus, F is in the domain of attraction of G with normalizing constants $a_N(M)$, $b_N(M)$ [16, p. 209], namely,

$$\lim_{N \rightarrow \infty} \Pr\{(T_l^* - b_N(M))/a_N(M) \leq x\} = G(x).$$

■

Using Theorem 11, we can approximate the probability that all the nodes in the network have received all the packets after t time slots, as follows,

$$\Pr\{(T_N^* - b_N(M))/a_N(M) \leq t\} \approx G(t).$$

By setting t large enough, one can guarantee that with, high probability, all the nodes have received successfully all the packets.

Now we extend the analysis to the constrained channel case, $C < M$ and assume for simplicity that M is divisible by C . Consider the PCRR policy, which transmits the same C packets every $\frac{M}{C}$ time slots. Let T be the total time to transmit all packets to all nodes. Let T_i where $i = 1, \dots, \frac{M}{C}$ be the number of time slots used to transmit packets $(i-1)C + 1, (i-1)C + 2, \dots, (i-1)C + C$. Clearly, we have $T_i = T/\frac{M}{C}$ for all i . For every C packets, C channels are available. Thus Theorem 11 can be applied, and we have

$$\Pr\{(T_i - b_N(C))/a_N(C) \leq t\} \approx G(t).$$

Since $T_i = T/\frac{M}{C}$, we have

$$\Pr\{T \leq \frac{M}{C}t\} = \prod_{i=1}^{\frac{M}{C}} \Pr\{T_i \leq t\}.$$

Thus,

$$\Pr\{T \leq \frac{M}{C}(a_N(C)t + b_N(C))\} \approx G(t)^{\frac{M}{C}}. \quad (4)$$

If M is not divisible by C , we can employ a slightly conservative approach, whereby we replace $\frac{M}{C}$ in Eq. 4 by $\lceil \frac{M}{C} \rceil$.




Fig. 4. Performance of PCRR algorithm with different number of channels for $S \geq C$

VI. NUMERICAL RESULTS

We next present numerical results to evaluate the following aspects: (i) the gain achieved using multiple channels with single radios; (ii) the performance of PCRR compared to the optimal policy; and (iii) the design accuracy of the proposed ACK-less protocol.

A. Multi-channel Gain

1) *Set Up*: We consider the case of disseminating a file consisting of $M = 20$ packets employing the PCRR transmission policy, on a single cluster topology. The packet loss probability p is set to 0.3. The results are obtained by taking average on 1000 identical simulations. We present results first for the case $S \geq C$ and then for the case $S = 1$.

2) *Simulation Results*: Figure 4 shows the gain of multiple channels compared to single channel for a single cluster topology and $S \geq C$. As expected, the average completion time in each case increases logarithmically with N (note that the x-axis in the figure is on a logarithmic scale). The figure shows that significant reduction in the expected completion time can be achieved using only two channels. Furthermore, as predicted by the asymptotic analysis, we observe, using linear regression, that the slopes of the curves corresponding to C channels are about C times smaller than those corresponding to a single channel, where $C = 2, 5, 10$.

Figure 5 illustrates the performance of the PCRR policy for the case $S = 1$, i.e., when there is initially only one source possessing the entire file. In this implementation, as soon as a node receives all the M packets, it can serve as an additional source. Obviously, the maximum number of sources cannot exceed C . As expected, the performance of PCRR in this case is similar to the unconstrained channel case. Interestingly, when $C = 10$, we find that the expected completion time decreases initially as the network size increases. This fact can be explained by noting that, as the network size increases, it takes less time to find nodes that have received all the packets and serve as sources on different channels.

B. Comparison with Optimal Policy

We next compare the performance of the PCRR policy with that of the optimal policy obtained by solving Eq. 1. As discussed in section II-B, solving the optimal policy is computationally involved. We thus consider a small example where $M = 3, C = 2, N = 1..20$. Even then, the state and control spaces in the dynamic programming problem are huge. Therefore, we consider the case for which in-order packet delivery is required, which substantially reduces the size of the state and control spaces. Figure 6 shows

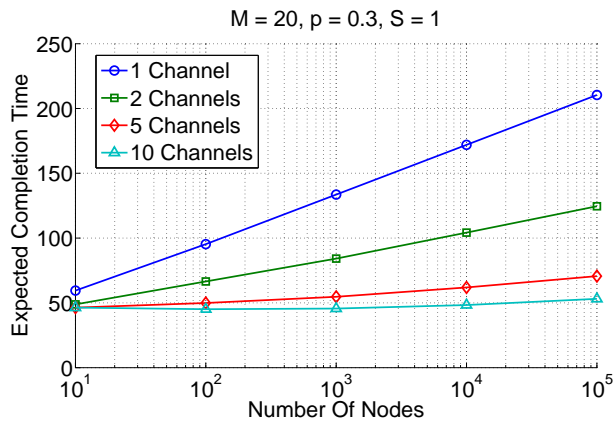


Fig. 5. Performance of PCRR algorithm with different number of channels for $S = 1$

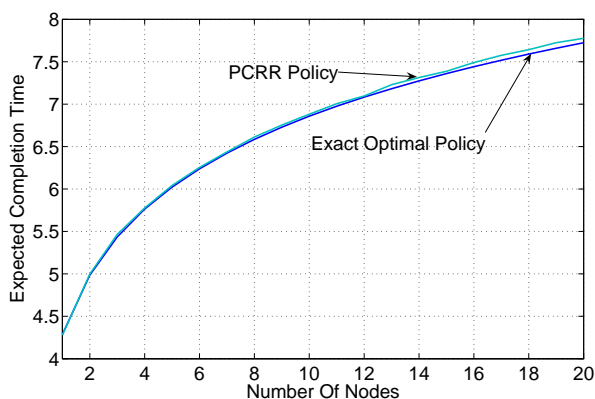


Fig. 6. Comparison of PCRR and optimal policy for single cluster with in-order packet delivery

that the performance of PCRR is very close to that of the optimal policy. Although the PCRR policy is proved to be near-optimal only when N is very large, it performs very well compared to the exact optimal policy, even in small networks.

C. Performance of the ACK-less Protocol

In this simulation, we evaluate our methodology based on extreme value theory to approximate the probability that all the nodes have successfully received all the packets after t time slots.

1) *Set up*: We consider a single cluster topology consisting of $N = 100$ nodes. The other parameters are set as follows: $M = 20$, $C = 10$, and $p = 0.1$. The results are obtained by taking average over 10000 identical simulations.

2) *Simulation Results*: Figure 7 depicts the probability that all the nodes in the network have received all the packets after t time slots, as a function of t . The figure compares the results obtained with our approximate analysis (see Eq. 4) and simulation. Although the two curves are known to converge only when $N \rightarrow \infty$, the result shows that they are already quite close for $N = 100$ nodes. This result illustrates the potential of our method towards the design of efficient, ACK-less data dissemination schemes.




Fig. 7. Probability of completion versus number of time slots

VII. CONCLUDING REMARKS

In this paper, we have shown that the multichannel transceiving capability of sensors can be exploited to achieve significance reduction in the delay of data dissemination. In particular, judicious variations of round robin strategies can achieve near-optimal performance in important, practical topologies. Surprisingly, the presence of a separate radio (interface) for each channel is not needed to achieve substantial performance gain, proportional to the number of channels. Finally, we have shown that extreme value theory could prove very useful in designing reliable data dissemination protocols with minimal control overhead.

ACKNOWLEDGEMENT

The work of the third author was conducted at Boston University. This work was supported in part by NSF CAREER grants ANI-0132802 and CCR-0133521 and NSF grants ANI-0240333 and CNS-0435312.

REFERENCES

- [1] J. Hui and D. Culler, "The dynamic behavior of a data dissemination protocol for network programming at scale." in *SenSys'04*, Baltimore, Maryland, USA, Nov. 2004.
- [2] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: a scalable and robust communication paradigm for sensor networks," *Mobicom*, 2000.
- [3] S. Kulkarni and L. Wang, "MNP: Multihop network reprogramming service for sensor networks," in *25th IEEE International Conference on Distributed Computing Systems*, 2005, pp. 7–16.
- [4] W. Xiao and D. Starobinski, "Poster abstract: Exploiting multi-channel diversity to speed up over-the-air programming of wireless sensor networks," in *SenSys'05*, San Diego, California, USA, Nov. 2005.
- [5] R. Simon, L. Huang, E. Farrugia, and S. Setia, "Using multiple communication channels for efficient data dissemination in wireless sensor networks," in *Mobile Adhoc and Sensor Systems Conference*, Washington, DC, USA, Nov. 2005.
- [6] G. Zhou, C. Huang, T. Yan, T. He, J. A. Stankovic, and T. F. Abdelzaher, "MMSN: Multi-frequency media access control for wireless sensor networks," *INFOCOM*, 2006.
- [7] M. J. Miller and N. H. Vaidya, "Leveraging channel diversity for key establishment in wireless sensor networks," in *INFOCOM*, 2006.
- [8] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.
- [9] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *33rd Annual Hawaii International Conference on System Sciences(HICSS)*, Jan 2000.
- [10] S. Lindsey and C. S. Raghavendra, "Pegasis: Power efficient gathering in sensor information systems," *IEEE Aerospace Conference*, Mar. 2002.
- [11] J. Galambos, *The Asymptotic Theory of Extreme Order Statistics*. Robert Krieger Publishing Company, 1987.
- [12] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," *Wirel. Netw.*, vol. 8, no. 2/3, pp. 153–167, 2002.
- [13] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms*. MIT Press, 2001.
- [14] D. Bertsekas and J. N. Tsitsiklis, "An analysis of stochastic shortest path problems," *Mathematics of Operations Research*, vol. 16, pp. 580–595, 1991.
- [15] S. Ross, *Stochastic Processes*, 1996, p. 405.
- [16] B. C. Arnold, N. Balakrishnan, and H. N. Nagaraja, *A First Course in Order Statistics*. John Wiley & Sons, Inc, 1992.
- [17] R. M. Young, "Euler's constant," *Math. Gazette*, vol. 75, pp. 187–190, 1991.
- [18] W. Feller, *An Introduction to Probability Theory and Its Applications*. John Wiley & Sons, Inc., 1957.
- [19] M. Fogiel and J. R. Ogden, *Handbook of Mathematical, Scientific, and Engineering Formulas, Tables, Functions, Graphs, Transforms*. Research & Education Assoc., 1984.