Overcoming Untuned Radios in Wireless Networks with Network Coding

Dragan Petrović, Kannan Ramchandran, Jan Rabaey
{dragan, kannanr, jan}@eecs.Berkeley.edu
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

***Abstract*— The drive toward the implementation and massive deployment of wireless sensor networks calls for ultra-low-cost and low-power nodes. While the digital subsystems of the nodes are still riding Moore's Law, there is no such trend regarding the performance of analog components. This work presents a fully integrated architecture of both digital and analog components (including local oscillator) that offers significant reduction in cost, size and power consumption of the overall node. While such a radical architecture cannot offer the reliable tuning of standard designs, it is shown that by using random network coding, a dense network of such nodes can achieve throughput linear in the number of channels available for communication. Moreover, the ratio of the achievable throughput of the untuned network to the throughput of a tuned network with perfect coordination is shown to be close to $1/e$.**

This work makes use of known results from network coding theory that show that throughput equal to the max-flow in a graph is achievable. However, the challenge here is finding the max-flow of the random graph corresponding to the network.

***Index Terms*—Sensor Networks, Network Coding, Untuned Radios**

I. INTRODUCTION

THE emerging field of wireless sensor networks has become a very active area of both academic research [1]-[4] and industrial development [5]-[9]. The potential scenarios for use of sensor networks are far ranging. Some of the near-term applications include monitoring the structural integrity of buildings and bridges [10], environmental control within living and working spaces [11], habitat monitoring in animal sanctuaries [12], highway traffic control [13], and warehouse inventory tracking [14]. The potential applications that are being considered in the long-term include such science-fiction-like concepts as smart surfaces that can respond to contact or serve as a communication backplane and airplane wings that can provide real-time monitoring of and alerts regarding the state of every square centimeter of the wing surface. It may even be possible to eventually produce sensor nodes small enough to be inserted into the blood-stream to provide real-time diagnostics of factors such as blood pressure, blood flow, glucose and insulin levels, etc.

To make such deployments economically and technologically feasible, it is necessary to drastically reduce the cost, size and energy consumption of the nodes available today. Moore's Law still provides for exponential reduction of these metrics over time when it comes to the digital components that comprise the memory, computation and coding of the nodes. However, there is no equivalent trend to Moore's Law that applies to the analog components needed for the radios that enable the nodes to communicate with one other. This work introduces an architecture for the analog radios that can greatly reduce the cost, size (5x reduction) and energy consumption (10x reduction) of the nodes. In fact it is expected that the proposed architecture will allow the energy consumption of the nodes to be so low that they could be fully powered by energy scavenged from the environment [15]. The penalty for using such a radical architecture is that the radios become untuned and it is no longer possible to guarantee that any arbitrary pair of nodes will be able to communicate with each other. Instead, it becomes necessary to rely on the density of nodes to make the overall network capable of providing reliable communication.

Narrowband radios have shown to be the architecture of choice for low-power applications [6],[7],[16], as they are low in complexity and consume less power than spread spectrum or other wide-band techniques. One fundamental requirement of narrowband radios is that the transmitter's carrier frequency and the receiver's detection frequency must be well-matched. This is traditionally accomplished by employing a crystal at both the transmitter and receiver to provide the same low frequency reference. This reference frequency is multiplied via a phase-locked loop (PLL) to generate the carrier wave. However, the off-chip crystal contributes significantly to the cost, size, and power consumption of such transceivers. The cost is due to the external quartz crystals being more expensive than the silicon used for the baseband signal processing as well as the need to bond separate components. This problem is especially acute in the design of highly integrated transceivers for wireless sensor networks. The size of traditional low power transceivers is largely due to the external crystal reference and the interface between the crystal and the silicon integrated circuit (IC). Additionally, the power consumption of low power radios is dominated by the crystal referenced PLL. *Therefore, great savings in all three of these areas could be obtained by eliminating the off-chip crystal and PLL.*

Even when care is taken to ensure that all radios are tuned and are attempting to communicate on the same frequency, reliable communication is not guaranteed. Practical implementations of sensor networks are notorious for having unstable links because narrowband communication is susceptible to deep fades between nodes [17],[18]. Since it is not feasible to overcome these fades by transmitting with more power (due to power-constraints), it has been proposed that randomized algorithms be used to ensure reliable communication [19],[20]. Such algorithms propose to provide reliable multi-hop communication by exploiting the

This work was sponsored in part by NSF grants CCR-0325311 and CCR-0330514

broadcast nature of wireless transmissions. The key idea is for a transmitting node to send a beacon to many potential forwarding nodes and then select one node to be the next hop for the packet among those that respond to the beacon. However, collisions among the responses to the beacon as well as the time-varying quality of the communication channels (a channel may be good during the beaconing, but become bad during the response and/or data transmission) contribute significant overhead to such schemes.

This paper proposes a fundamentally different way of designing and operating a transceiver. The quartz crystal is eliminated and replaced by an on-chip resonator such as an inductor-capacitor (LC)-circuit or a nano-electromechanical resonant structure. This makes it possible to economically produce millions of nodes and densely deploy them by weaving them into fabrics or mixing them with paint. The proposed architecture allows a sensor node to be developed entirely out of thin-film technologies (radio, digital, battery, energy scavenging, and sensing). However, the drawback of such architectures is that the variations in the manufacturing process are large, resulting in un-tuned radios. Therefore, two narrowband radios produced by such a process are not likely to be able to communicate with each other. To address this problem, a low-complexity communication protocol is proposed that makes use of the high density of nodes to ensure reliable communication using such un-tuned radios even without the need for handshaking protocols or re-transmission. By eliminating the need for this kind of coordination, the protocol is also made more robust to link failures, while the density that is made possible by such low cost designs makes the network robust to the failure of individual nodes.

Main Result: We consider using the nodes to form a communication backplane carrying data between a source and a destination. The data is transported in a multi-hop fashion by a network of nodes that employ untuned narrowband radios. Let N denote the number of unit-capacity channels available for communication. We show that by using random linear network coding, achievable throughput of the network is $O(N)$, same as in a fully-coordinated network of tuned radios. Moreover, the ratio of the achievable throughput of the untuned network to the achievable throughput of a tuned network is shown to be close to $1/e$. This work makes use of known results from network coding theory that show that throughput equal to the max-flow in a graph is achievable even when the connectivity of the network is not known a priori. *However, the challenge here is in finding the max-flow of the random graph corresponding to the network.*

II. UN-TUNED RADIOS

The drawback of using an on-chip resonator is that the variations in the manufacturing process are large meaning that the distribution of the resonant frequency of the manufactured oscillators will have large variance.

Figure 1 presents a qualitative illustration of the challenges presented by this approach. In a traditional narrow-band architecture with a quartz crystal reference, the signal bandwidth is typically orders of magnitude larger than the center frequency tolerance. When using un-tuned receivers, the situation is reversed and the carrier frequency variation is orders of magnitude greater than the signal bandwidth. To achieve frequency tolerance approaching a quartz crystal, prohibitively expensive trimming would have to take place. In addition, drift over time, temperature, and supply voltage would quickly render the trimming inaccurate. The other option is to leave the transmitters and receivers un-tuned, which means that two narrowband radios produced by such a process are not likely to be able to communicate with each other. If the input frequency range of a particular receiver did admit the transmitter's carrier frequency, the transmission would be received successfully. Otherwise, the result would be the same as if the transmitter were communicating on a channel orthogonal to the one that the receiver is monitoring. This means that the unreliable manufacturing process results in having multiple channels available for communication but neither the transmitting nor receiving nodes can select on which particular channel to communicate. Instead the channel on which a transmitter or receiver communicate are random.

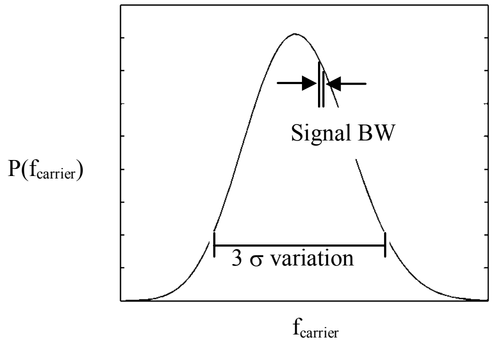

Figure 1. Signal bandwidth relative to process variation when using on-chip LC resonators to provide the carrier frequency for narrowband radio

Even though a particular transmit-receive pair may not be able to communicate because they would be effectively tuned to different channels, a sufficiently high density of nodes ensures a high probability that there are pairs of transmitters and receivers that can communicate with one another. The number of channels available in the system is determined by the ratio of the variations of the manufacturing process to the receiver bandwidth. As long as the bandwidth admitted by the receive filters is greater than the signal bandwidth, there will be a non-zero probability that two randomly selected nodes will be able to communicate with each other.

In the rest of the paper, the abstraction of having multiple channels available for communication is made. Note that it is not necessary for the channels to be orthogonal (in fact, they are not!). What is important is the number of transmitter carrier frequencies that fall within the bandwidth being monitored by a particular receiver. The probability that any given transmitter falls within the receiver's range is dependent only on the ratio of the range of possible carrier frequencies to the receiver bandwidth. This ratio is equivalent to the number of channels in our analysis because it is equal to the maximum number of independent transmissions that can be made simultaneously without interfering with one another.

To maximize the throughput of the network, it is necessary to maximize the probability that during communication a channel is occupied by exactly one transmitter. This will maximize the number of channels that contain a decodable transmission. It can be shown that when there are N channels available for communication and each transmitter is independently and randomly assigned to a channel with the same probability distribution, the probability that a channel is occupied by exactly one transmitter is maximized when there are exactly N transmitters and each transmitter is equally likely to be assigned to any of the channels. In this case, the probability that a channel contains exactly one transmission is asymptotically, for large N, equal to $1/e$[1]. This implies that, in order to maximize the throughput, the network should be operated with N active transmitters within communication range of each other, in which case each transmission will experience a collision with probability $1-1/e$.

Having N active transmitters within communication range maximizes the probability that a receiver will have exactly one transmitter in the range of frequencies it monitors; however, it may still be possible that a transmitter occupies a unique channel, but no receiver is tuned to that channel. In order to increase the probability that a non-colliding transmission is heard as well as the probability that a receiving node hears at least one packet, each one is equipped with several receive radios, with each radio tuned to a different channel (by using a different LC-circuit as the local oscillator for each radio - this also allows the nodes to transmit on different random channels at different times by selecting any of its LC-circuits to provide the carrier frequency for the transmitter).

Denote the number of receive-radios on each node with L. In Section IV, we derive the relationship between the value of L and the throughput of the network. We show that it is possible to achieve throughput that is linear in the number of channels even with a constant value of L. It is important to show that this is achievable with a constant L because requiring L to grow with the number of channels would correspond to requiring more hardware on each node, and this is exactly what we are trying to avoid. It should be noted that the theoretical results use bounding techniques, so the constants of this linear throughput that are guaranteed by any particular value of L are pessimistic. To complement the theoretical result and give guidelines for practical deployments, simulations are used to estimate the throughput that can be achieved with different values of L.

III. MULTI-HOP COMMUNICATION

We consider using the nodes to form a communication backplane carrying data between a source and a destination. This scenario is shown in Figure 2. The source sends its data to the destination in a multi-hop fashion using geographic routing. The region between the source and the destination is divided into blocks such that any two nodes in adjacent blocks are within communication range of each other. This allows the nodes to make use of the broadcast nature of wireless communication by having multiple potential forwarders for each transmission. Usually, this is exploited by using beaconing to establish a connection between a transmitting node and a forwarder [19],[20], but in our scheme, the transmitters simply broadcast their data and assume that at least one node in the next block will receive and relay the packet.

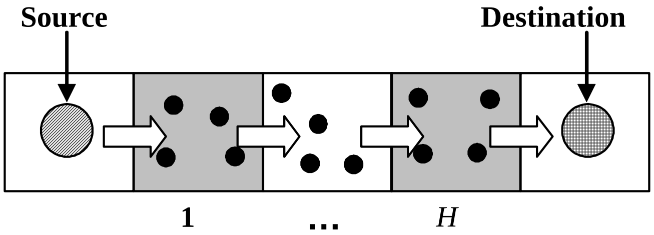

Figure 2. Proposed multi-hop communication method

In the proposed scheme, anytime a node transmits a packet, it includes in the packet's header the coordinates of the next block downstream as well as the coordinates of the destination so that only nodes in the next block will relay the information, again specifying the direction in which the data should propagate. It is assumed that the nodes know their position [22] to within the accuracy of one block, or that they know their position relative to each other and the destination, which can be achieved by sending out a one-time flood from the destination and having each node remember its hop-count to the destination. The source has to send its data at the same time and in waves. After sending one wave of data, it has to wait long enough for that wave to propagate four blocks so as to ensure that the next wave of data will not interfere with the retransmissions of the previous wave by the nodes in downstream blocks. By having the source send a wave of data simultaneously and the intermediate nodes relay wave as soon as they receive it, the network is made to operate as a time-slotted communication system.

Sending the data in waves prevents collisions among broadcasts from different blocks, however the fact that the transmitters are untuned results in collisions by transmissions within the same block. If every node in a block transmits on a random frequency, it is likely that there will be transmissions at frequencies close to each other, thus any receiver tuned to those frequencies will detect a collision and will not be able to decode the individual transmissions. These collisions effectively erase some of the packets, making it seem as if nodes in neighboring blocks communicate with each other through an erasure channel. The question we are interested in is, given N unit-capacity communication channels, how much data can simultaneously be sent to the destination and have this data successfully received and decoded by the destination that is H hops away (for now, let us assume that H is a constant, though later it will be shown that H may be allowed to grow with N, as long as $H(N)=o(e^{N})$, without affecting the asymptotic throughput). We want to compare this to a fully-coordinated network employing tuned radios, in which case exactly N packets could be sent in each wave, provided that there are N nodes in each block and every node selects a unique frequency on which to communicate.

IV. THROUGHPUT

We will find the throughput of the untuned network by showing that the connectivity of the network can be modeled as a random graph and then applying known results from network coding literature. Namely, we make use of the result that for communication in a graph of unit capacity links for which the connectivity is not known a priori, a throughput equal to the

---

[1] The result and its derivation are similar to that for the throughput of the slotted ALOHA protocol [21]. The proof is omitted here due to space constraints.

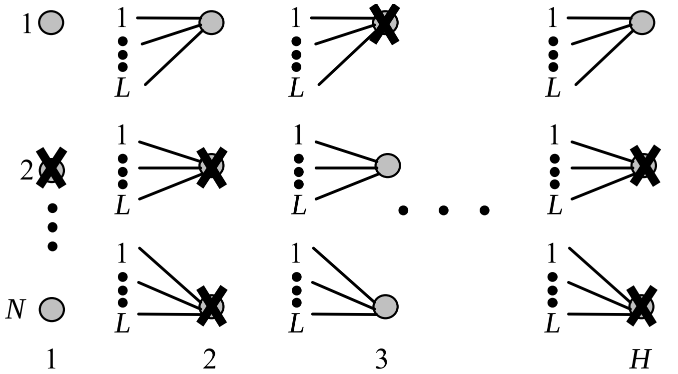

Figure 3. Random graph representing connectivity in the network of nodes with untuned radios. Each vertex (node) has L inputs, each one coming from a randomly and uniformly selected node in the previous column. Each node, along with its incoming and outgoing links, is deleted with probability 1-1/e.

max-flow between the source and the destination is achievable with arbitrarily high probability by using random linear network coding[2] over a high enough field size [23]. However, in order to make use of this result, we must find the max-flow of the graph. This is done by Result 1. Since the connectivity of this random graph is not static (i.e. each wave of data will encounter a different set of links), the packets have to carry the encoding vectors in their headers to provide the destination with just the right information needed to decode the source packets as in the scheme introduced in [24]. It can also be shown that simple routing (in which forwarding nodes are only allowed to forward one of the packets from each wave rather than combining all the packets they receive in each wave) achieves performance that is strictly sub-linear in N [25].

A. Random Graph Representation

We now create the random graph, shown in Figure 3, that models the connectivity of the network. The N vertices in each column correspond to the N nodes in each block during communication. The H columns correspond to attempting to communicate over H blocks. Each of the vertices has L incoming links corresponding to the L receivers on each node. Each link connects a vertex to a randomly, independently chosen vertex in the previous column. Since transmissions experience collisions with probability $1-1/e$, each of the vertices in the graph is deleted with probability $1-1/e$, in which case all of its incoming and outgoing links are also deleted[3]. This means that each of the links is deleted with probability $1-1/e$ because each receive-radio has probability $1-1/e$ of being tuned to a frequency range that does not contain a decodable transmission (i.e. either no transmission or more than one). The links that are not deleted are equally likely to connect to any of the vertices in the previous column that are not deleted because, given that a receive-radio has exactly one transmission in its receive frequency range, the source of that transmission is equally likely to be any of the transmitters that do not experience a collision.

We label the resulting random graph as $G_{L,1-1/e}$ and show that the max-flow of $G_{L,1-1/e}$ is close to $1/e$ if L is large enough.

B. Max-flow of Random Graph

Result 1: For any constant β such that $\beta < 1/e$ there exists a *constant* number of inputs/node L such that the max-flow of $G_{L,1-1/e}$ is greater than $\beta \cdot N$ with high probability as N goes to infinity.

We will prove this result by applying a modified version of a technique used in Percolation Theory. The first step is to relate the likelihood of having many disjoint end-to-end (E2E) paths to the likelihood of having even a single path E2E. Let us define the following notation:

Let A be the event that there exists *a path* E2E and let A_r be the event having the following property: starting with any graph in A_r, the deletion of *any* r vertices will still result in A. This is equivalent to saying that any graph in A_r has at least $r+1$ vertex-disjoint paths E2E.

Lemma 1: Let r be a positive integer. Then

$$1 - P_{p_2}(A_r) \le \left(\frac{q_2}{q_2 - q_1}\right)^r \left\{1 - P_{p_1}(A)\right\} \tag{1}$$

whenever $0 \le p_2 \le p_1 \le 1$. Here, $q_1 = 1 - p_1$ and $q_2 = 1 - p_2$, and the notation $P_p(\cdot)$ represents the probability that the event in parentheses occurs when vertices in the graph are deleted with probability p.

Proof of Lemma 1: This proof is based on the proof in [26] of a similar result from Percolation Theory. Let $X_{i,j}\ \forall\ i \in \{1,...,N\}$ and $j \in \{1,...,H\}$ be i.i.d. random variables uniformly distributed in the interval $[0,1]$, and to each vertex in row i and column j of the grid assign the value $X_{i,j}$. To create graphs G_{L,p_1} and G_{L,p_2} that have vertices deleted with probability p_1 and p_2 respectively, do the following: first assign the values $X_{i,j}$ to each vertex in the grid. Then assign L links from each node in columns 2 through H to a randomly selected node in the previous column. Finally, to create graph G_{L,p_1}, for each vertex i,j, delete it iff $X_{i,j} \le p_1$. To create graph G_{L,p_2}, for each vertex i,j, delete it iff $X_{i,j} \le p_2$.

We are interested in relating the likelihood that A_r occurs in G_{L,p_2} to the likelihood that A occurs in G_{L,p_1}. Note that if A_r does not occur in G_{L,p_2}, then there must be a set of vertices, B, such that:

a) All of the vertices in the set B are not deleted in G_{L,p_2}

b) $|B| \le r$

c) The graph $\overline{G}_{L,p_2}$ obtained by deleting from G_{L,p_2} the vertices in B satisfies $\overline{G}_{L,p_2} \notin A$.

There may exist many such sets B, in which case it is sufficient to pick any such set. Suppose that $G_{L,p_2} \notin A_r$, and that

---

[2] In random linear network coding, forwarding nodes send on each outgoing link a random linear combination of the packets it receives on the input links. Each input packet is multiplied by a randomly chosen element from some Galois Field and these products are added together to form the outgoing packet [23].

[3] In the random graph, the vertices are deleted independently of one another. This is not the case in the network since the collisions are not independent; however, this approximation becomes accurate as N tends to infinity. The independence assumption allows for analytical tractability in what follows.

every vertex i,j in the set B satisfies $p_2 < X_{i,j} \le p_1$. It then follows from c) that $G_{L,p_1} \notin A$. Conditional on B, there is a $[(p_1 - p_2)/(1 - p_2)]^{|B|} = [(q_2 - q_1)/q_2]^{|B|}$ probability that $p_2 < X_{i,j} \le p_1$ for all vertices in B; therefore,

$$P\left(G_{L,p_1} \notin A \mid G_{L,p_2} \notin A_r\right) \ge \left(\frac{q_2 - q_1}{q_2}\right)^r \tag{2}$$

Applying Bayes's theorem and the fact that $P\left(G_{L,p_1} \notin A \cap G_{L,p_2} \notin A_r\right) \le P_{p_2}\left(G_{L,p_1} \notin A_r\right)$ gives the result of Lemma 1. □

This result of Lemma 1 is particularly useful if we can show that the probability that that $G_{L,p_1} \notin A$ decays exponentially (with N) to zero for some p_1. In other words, if we can show that $\{1 - P_{p_1}(A)\} \le e^{-N\alpha(p_1,L)}$, then we have

$$1 - P_{p_2}(A_r) \le \left(\frac{q_2}{q_2 - q_1}\right)^r e^{-N\alpha(p_1,L)}, \tag{3}$$

and applying $r = \beta \cdot N$ tells us that the probability of not having $\beta \cdot N$ (actually $\beta \cdot N + 1$) paths decays to zero exponentially as long as

$$\beta < \alpha(p_1, L)/\log\left(\frac{q_2}{q_2 - q_1}\right). \tag{4}$$

The problem now becomes finding an appropriate bound on $P\left(G_{L,p_1} \notin A\right)$.

Lemma2: If $L(1 - p_1) > 1$, then $P\left(G_{L,p_1} \notin A\right) \le H \cdot P\left(Y/N < 1 - Z^*\right)$ where Y is a random variable drawn from the Binomial$\left(N,\ 1 - p_1 + (1 - p_1)Z^{*L}\right)$ distribution and $Z^* = [L(1 - p_1)]^{1/(1-L)}$.

Lemma 2 allows us to relate the probability that no E2E path exists in G_{L,p_1} to the probability that the mean of N Bernoulli random variables deviates from its expected value by some amount. Since this probability decays to zero exponentially in N, this result, along with Lemma 1 will allow us to prove that the number of vertex-disjoint paths in G_{L,p_2} will be linear in N with high probability. The constant, β, of this linear relationship will depend on the value of L. Note that β also depends on the value p_1; however, the value p_1 is not fundamental to the graph G_{L,p_2}, and we are allowed to assign any value to p_1, as long as it is larger than p_2, so as to maximize the bound on β guaranteed by Lemma 1 and Lemma 2.

Also note that the condition that $L(1 - p_1) > 1$ is imposed to ensure that $1 - Z^* > 0$. It can be shown that if $L(1 - p_1) > 1$, then $G_{L,p_1} \in A$ with high probability, otherwise $G_{L,p_1} \notin A$ with high probability. However, in our case it is not enough to show that $G_{L,p_1} \in A$ with high probability for appropriate values of L and p_1. We must also bound the rate of this convergence (using Lemma 2) in order to apply Lemma 1 to our original problem.

Proof of Lemma 2: Consider the number of vertices in each column of G_{L,p_1} that were not deleted and have a path back to column 1. Let us call these vertices “good” and all the others “bad.” Conditioned on the number of bad (good) vertices in column j, vertices in column $j+1$ are themselves good or bad independently of each other and with equal probability. Let the number of bad vertices in column j be $Z \cdot N$ for some Z that satisfies $0 \le Z \le 1$. Then vertices in column $j+1$ are themselves bad with probability $p_1 + (1 - p_1)Z^L$.

Consider the probability that the number of bad nodes in column $j+1$ is greater than $Z \cdot N$, given that the number of bad nodes in column j is $Z \cdot N$. If $p_1 + (1 - p_1)Z^L < Z$, this probability should be exponentially small in N. This probability is minimized when the difference $Z - p_1 + (1 - p_1)Z^L$ is greatest. Setting the first derivative of this function, which is convex in the interval $[0,1)$, to zero shows that the expression is maximized when $Z = [L(1 - p_1)]^{1/(1-L)}$.

Let R_j represent the ratio of bad nodes to the total number of nodes in column j (i.e. the total number of bad nodes in column j is $R_j \cdot N$). $G_{L,p_1} \notin A$ is equivalent to $R_H = 1$. Note that if $R_j = 1$ for some $j \in [1, H)$, then $R_k = 1\ \forall\ k \in [j+1, H]$ because if column j has no connectivity back to column 1, then none of the columns after j will have connectivity to column 1 either. We prove Lemma 2 by arguing that $P\left(G_{L,p_1} \notin A\right)$ is upper bounded by the probability that there exist a $j \in [1, H]$ for which $R_j = 1$, and this is upper bounded by the probability that there exist a $j \in [1, H]$ for which $R_j > Z^*$. Mathematically,

$$\begin{aligned}
&P\left(G_{L,p_1} \notin A\right) < P_{p_1}\left(\exists\ j \in [1,H] \quad s.t. \quad R_j > Z^*\right) \\
&\le P_{p_1}\left(R_1 > Z^*\right) + \sum_{j=2}^{H} P_{p_1}\left(R_j > Z^* \mid R_{j-1} \le Z^*\right) \\
&< P_{p_1}\left(R_1 > Z^*\right) + \sum_{j=2}^{H} P_{p_1}\left(R_j > Z^* \mid R_{j-1} = Z^*\right) \\
&= P_{p_1}\left(R_1 > Z^*\right) + (H - 1) \cdot P\left(Y/N < 1 - Z^*\right) \\
&< H \cdot P\left(Y/N < 1 - Z^*\right)
\end{aligned} \tag{5}$$

where the second inequality is by the union bound, the last inequality is because $p_1 < p_1 + (1 - p_1)Z^{*L}$, and Y is a random variable drawn from the Binomial$\left(N,\ 1 - p_1 + (1 - p_1)Z^{*L}\right)$ distribution. □

Now, we must find the rate at which $P\left(Y/N < 1 - Z^*\right)$ decays to zero and apply this to Lemma 1 to show that the number of vertex-disjoint paths in G_{L,p_2} grows linearly with N. Fortunately, this rate is well known [27]. We use the following notation: $q = 1 - [p_1 + (1 - p_1)Z^{*L}]$ and $\varepsilon = q - 1 + Z^*$ to write [27]:

$$P\left(Y/N < 1 - Z^*\right) = P\left(q - Y/N > \varepsilon\right) \le \left(\frac{q-\varepsilon}{q}\right)^{-N(q-\varepsilon)} \left(\frac{1-q+\varepsilon}{1-q}\right)^{-N(1-q+\varepsilon)} \tag{6}$$

The right-hand side of the equation can also be expressed as

$$\exp\left(\log\left[\left(\frac{q-\varepsilon}{q}\right)^{-N(q-\varepsilon)} \left(\frac{1-q+\varepsilon}{1-q}\right)^{-N(1-q+\varepsilon)}\right]\right) \tag{7}$$

giving us

$$\alpha(p_1, L) = \log\left[\left(\frac{q-\varepsilon}{q}\right)^{(q-\varepsilon)} \left(\frac{1-q+\varepsilon}{1-q}\right)^{(1-q+\varepsilon)}\right] \tag{8}$$

as the $\alpha(p_1, L)$ we need to plug into (4). Evaluating (4) with a large enough *but constant* value for L and appropriately chosen value for p_1 provides the guarantee that he max-flow of $G_{L,1-1/e}$ is at least $\beta \cdot N$ for any $\beta < 1/e$, proving Result 1. ■

The throughput is not dependent on H because we held H constant while letting N go to infinity. However, (5) implies that H may grow with N without affecting the throughput result as long as $\log(H(N))/N$ goes to zero as N grows.

Note that the result is proven using bounding techniques, so it is not tight. For example, for $L = 10$, the result proves that the throughput is guaranteed to be at least 2%, whereas simulations show that $L = 10$ is good enough to give throughput of 30%. Figure 4 shows the result of simulations for various values of L at $N = 1000$.

This max-flow result, together with the random network coding result of [23] tells us that each wave of data can deliver nearly N/e packets from the sources to the destination, compared to N packets that could be transported in each wave if the network were composed of nodes with tuned radios and perfect coordination.

In practical deployments, having about 10 radios per node would be realistic. In this case, the theoretical results tell us that a throughput of only $0.02 \cdot N$ can be guaranteed. However, simulations show that a throughput of $0.3 \cdot N$ can be expected. This is a good trade-off for many applications in which the demand on bandwidth is not as strict as the demand for low-cost nodes that can operate at power levels comparable to the power levels that can be supplied by the energy scavenging mechanisms.

V. CONCLUSION

This paper shows how network coding can be used to achieve high throughput in an ad-hoc wireless network of nodes with untuned radios. This makes it possible to build ultra low cost and low power devices, and deploy a high-performance network by utilizing randomized algorithms and high density of such nodes.

ACKNOWLEDGMENT

The authors would like to thank Brian Otis, Massimo Franceschetti, Alistair Sinclair, and Vinod Prabhakaran for sharing their expertise and insights during the course of this work.

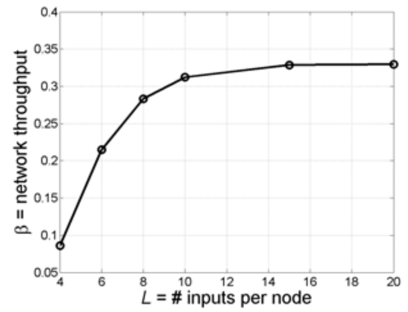

Figure 4. Simulation results showing the throughput achievable as a function of the number of inputs per node for 1000 channels and 1000 nodes per block.

REFERENCES

[1] http://www.tinyos.net/
[2] http://webs.cs.berkeley.edu/nest-index.html
[3] http://bwrc.eecs.berkeley.edu/Research/Pico_Radio/Default.htm
[4] http://www.cens.ucla.edu/
[5] http://www.intel-research.net/berkeley/index.asp
[6] http://www.zigbee.org
[7] http://www.xbow.com
[8] http://www.ember.com
[9] http://www.dust-inc.com/flash-index.shtml
[10] http://www.cs.berkeley.edu/~binetude/ggb/
[11] http://www.cbe.berkeley.edu/research/briefs-Wireless.htm
[12] A. Mainwaring, J. Polastre, R. Szewczyk, and D. Culler, "Wireless Sensor Networks for Habitat Monitoring," Intel Research, IRB-TR-02-006, Jun. 19, 2002.
[13] S. Coleri, S. Cheung, and P. Varaiya, "Sensor Networks for Monitoring Traffic," 42nd Annual Allerton Conference on Commuinication, Control, and Computing, September 2004.
[14] http://www.wherenet.com/
[15] S. Roundy, B. Otis, Y.H. Chee, J. Rabaey, P. Wright, "A 1.9GHz RF Transmit Beacon using Environmentally Scavenged Energy," Dig. IEEE Int. Symposium on Low Power Elec. and Devices, Seoul, Korea, 2003.
[16] B. Otis, Y.H. Chee, R. Lu, N.Pletcher, J. Rabaey, "An Ultra-Low Power MEMS-Based Two-Channel Transceiver for Wireless Sensor Networks", IEEE Symp. on VLSI Circuits, Honolulu, HI, June 2004.
[17] A. Willig et al., "Measurements of a Wireless Link in an Industrial Environment using an IEEE 802.11-Compliant Physical Layer," IEEE Trans. on Industrial Electronics, vol. 43, Dec. 2002.
[18] N. Patwari, Y. Wang, R. O'Dea, "The Importance of the Multipoint-to-Multipoint Indoor Radio Channel in Ad Hoc Networks," IEEE Wireless Communication and Networking Conference (WCNC), Orlando FL, March 2002.
[19] M. Zorzi, R.R. Rao, ``Geographic Random Forwarding (GeRaF) for ad hoc and sensor networks: energy and latency performance," IEEE Trans. on Mobile Computing, vol. 2, Oct.-Dec. 2003.
[20] J. Van Greuen, D. Petrović, A. Bonivento, J. Rabaey, K. Ramchandran, A. Sangiovanni-Vincentelli, "Adaptive Sleep Discipline for Energy Conservation and Robustness in Dense Sensor Networks", ICC 2004.
[21] M. Ferguson, "On the Control, Stability, and Waiting Time in a Slotted ALOHA Random-Access System," IEEE Transactions on Communications, Vol. 23, Issue 11, Nov 1975 pp. 1306 – 1311.
[22] L. Doherty, L. El Ghaoui, K. Pister, "Convex Position Estimation in Wireless Sensor Networks", Infocom 2001, Anchorage, AK, April 2001.
[23] T. Ho, M. Medard, J. Shi, M. Effros, and D. Krager, "On randomized network coding," Proceedings of 41st Annual Allerton Conference on Communication, Control, and Computing, October 2003.
[24] P. Chou, Y. Wu, and K. Jain, "Practical network coding," Allerton Conference on Communication, Control, and Computing, Monticello, IL, October 2003. Invited paper.
[25] D. Petrović, J. Rabaey, K. Ramchandran, "Wireless Sensor Networks Using Untuned Narrowband Radios," UC Berkeley Technical Report, to be published February 2005.
[26] M. Aizenman, J. T. Chayes, L. Chayes, J. Fröhlich, L. Russo, "On a sharp transition from area law to perimeter law in a system of random surfaces," Communications in Mathematical Physics vol. 92, pp. 19-69, 1983.
[27] H. Chernoff, "A measure of asymptotic efficiency for test of a hypothesis based on the sum of observations," Annals of Mathematical Statistics, 23:493-507, 1952.