

Exploiting Macrodiversity in Dense Multihop Networks and Relay Channels

Matthew C. Valenti

Lane Dept. Comp. Sci. & Elect. Eng.
West Virginia University
Morgantown, WV 25506-6109
mvalenti@wvu.edu

Neiyer Correal

Florida Commun. Research Labs
Motorola
Plantation, FL 33322
N.Correal@motorola.com

Abstract—Embedded networks of sensors and actuators must operate at extremely low power and use inexpensive single-antenna transceivers. The economics of such systems preclude the use of complex signal processing or antenna arrays at any one device. However, the same economics allows the coverage area to be blanketed with a high density of devices which results in a rich spatial diversity. This spatial diversity can be exploited by forming a virtual antenna array which combines observations made at multiple receivers. This paper illustrates the potential transmit energy savings that are possible by using such macrodiversity-combining approaches by the analysis and simulation of an idealized system. It concludes by discussing practical issues that must be considered before integrating macrodiversity-combining strategies into actual embedded networks.

I. INTRODUCTION

The combination of silicon scaling and energy-efficient multihop packet radio technology will soon allow low power devices to be embedded virtually everywhere, enabling a wide range of revolutionary applications (e.g. sensor networks) that will radically change the way that people and devices interact with their environments [1]. The devices that comprise these networks must operate at very low power and be inexpensive enough to allow high density deployments. Often, the battery must last as long as the device itself. Thus, the network must be designed to simultaneously use inexpensive devices and operate at very low power.

Recent advancements in communications research have enabled significant transmit power reductions, but typically these gains require sophisticated coded waveforms at the transmitter and complex detection algorithms at the receiver. Of particular relevance is the use of antenna arrays at both the transmitter and receiver to improve energy and/or bandwidth efficiency [2]. However, these approaches are not appropriate for use in embedded networks which require inexpensive devices with modest processing capabilities and a footprint that precludes the use of antenna arrays.

At first glance, it seems that the low-cost requirement of the devices that comprise an embedded network is a liability and that only simple, but energy inefficient transmission technologies can be permitted. But what appears to be a burden can be turned into an advantage because the low device cost allows the coverage area to be blanketed with a high density of devices. This high density creates an intrinsic spatial diversity that could, and should, be exploited in order to reduce the energy consumption of each device. However, the best way to exploit this rich, but indirect, spatial diversity is an open question and the focus

of this paper. A cross-layer approach to the problem is taken, and the proposed strategies involve nodes that work *collaboratively* by sharing observations taken at separate locations. In particular, a dense network could exploit its inherent rich spatial redundancy by macrodiversity-combining observations made at multiple locations, forming a *virtual* antenna array capable of massive energy savings in the presence of harsh fading and interference. The strategies discussed in this paper are similar to the concept of cooperative diversity [3], [4], but have been optimized for the extremely low power regime characteristic of sensor networks.

The purpose of this paper is to illustrate the potential gains that can be achieved by exploiting the macrodiversity present in dense ad-hoc networks and discuss some of the related research issues. Section II discusses the philosophy behind macrodiversity reception and presents a short literature survey showing how these concepts have previously been applied to ad-hoc networks. Section III gives a more detailed discussion of an idealized relay channel that uses macrodiversity-combining concepts, with simulation results following in Section IV. Section V discusses some practical considerations related to the integration of the proposed strategies in actual embedded networks, and Section VI concludes the paper.

II. MACRODIVERSITY IN AD-HOC NETWORKS

Macrodiversity combining is fairly straightforward if the receivers are connected over a reliable backbone, as in the uplink of a cellular network or infrastructure-based wireless local area network (LAN). In the uplink, observations from multiple access points could be combined on a packet- or symbol-level, thereby allowing each mobile station to reduce its transmit power in the presence of fading and interference [5]. However, the application of macrodiversity is less straightforward in ad-hoc systems comprised of packet radios performing multihop routing.

The key to achieving macrodiversity in a wireless multihop network is to broadcast each message to all nearby nodes rather than directing it to just a particular node. In a dense network, the message is likely to be received correctly by at least one node and thus the transmitter can afford to transmit at a lower power due to the *receiver* diversity effect. One or more of the nodes that received the broadcast may act as a relay, and in narrowband systems a negotiation is necessary to determine the order that these nodes attempt to forward the message. This process could be performed either explicitly through a control




Fig. 1. A dense double-hop network.

channel or implicitly within the transmission protocol. If direct-sequence spread-spectrum signaling is employed, *transmit* diversity can also be achieved by using a Rake receiver to diversity combine multiple relayed messages that are (perhaps intentionally) transmitted concurrently by multiple relays. Diversity routing is also possible by allowing the message to follow different paths through the network.

The benefits of applying macrodiversity concepts to ad-hoc networks was initially recognized in the 1970s by the DARPA-sponsored Packet Radio Network (PRNet) program [6]. In PRNet, the source first attempts to transmit a spread spectrum signal to a preselected relay node. If, after three attempts, the selected node fails to respond with a positive acknowledgment, the source switches to a broadcast-mode so that any node that receives (or had received) the message can forward it (provided that it has a clear channel to the destination).

More recently, Sendonaris et al [3] and Laneman et al [7], [4] have proposed a technique for exploiting macrodiversity in the wireless ad hoc network based upon the use of the classic relay channel [8]. A message is first broadcast by the source and is received simultaneously by both the destination and an intermediate relay node. The relay then forwards the message to the destination, where it is maximal ratio combined (MRC) with a stored copy of the direct message. A similar strategy termed *multihop diversity* is presented in [9]. Our work can be viewed as both a generalization of this related research that allows the use of *multiple* relays, as well as a specific instance of cooperative diversity that is tailored to the low cost requirements of embedded networks by using noncoherent reception and avoiding MRC combining.

III. A SIMPLE MACRODIVERSITY STRATEGY FOR THE RELAY CHANNEL

In the scenario shown in Fig. 1, a source must send a message to a destination but has insufficient power to transmit directly. However, there is a circular cloud of M randomly placed transceivers located halfway between the source and destination. Thus, the source could first send a message to one (or more) of these intermediate relay nodes, which could then forward the message to the destination. To keep the discussion simple, we assume that only the source and a *single* relay transmits. Thus, this is a *double-hop* network or equivalently, an instance of the classic *relay channel* [8]. In the classic relay channel, the relay could amplify and forward the message *without* decoding it. However, we do not allow this and require that the message must always be checked for errors before being forwarded. Furthermore, the classic relay channel allows the source to simultaneously broadcast to the relays and destination, but for simplicity we assume that the source does not transmit directly to the destination.

The source has the option of either transmitting to a *particular* relay or broadcasting it to *all* of them. The first option, called *receiver-directed* transmission [10], is the most common paradigm for transporting information in packet radio systems. It requires that the source select, prior to transmission and according to a routing algorithm, one of the M relays to serve as the forwarding relay. Once the source begins its transmission, only the selected relay participates in the process. Thus the spatial redundancy due to the presence of the other $M - 1$ relays is wasted.

Macrodiversity can be introduced into this scenario by allowing the source to *broadcast*, rather than constraining it to address a particular relay. Because of the assumed high network density, it is likely that at least one of the relays receives the message correctly and could act as the forwarding relay. Note that because broadcasting the message eliminates the need for the source to preselect a relay, this could result in a simpler routing algorithm. However, the situation is complicated by the likely event that more than one relay successfully receives the broadcast (we call the relays that receive the source transmission correctly *potential* relays). In this case, a negotiation must take place in order for the potential relays to determine which one will actually forward the message. This negotiation could be performed explicitly over a control channel or implicitly by being embedded into a combined MAC/network-layer protocol. Ideally, the potential relay that has the best instantaneous channel to the destination would be elected to forward the message. In practice, the instantaneous SNR might not be known and thus either average SNR or relative position could be used to determine the forwarding relay. If the initially selected relay fails to transmit successfully to the destination, then one of the other potential relays could be used to retransmit the packet; the benefit of this strategy would be that the extra energy required for retransmissions would be distributed across the cloud.

IV. SIMULATION RESULTS

A set of simulations was performed to compare the potential performance of the proposed strategy of broadcasting to all relays with that of a more conventional receiver-directed strategy. The system topology is as shown in Fig. 1, with the source and destination separated by 10 m and the M relays placed randomly inside a 7 m wide circle whose center is halfway between source and destination. After each attempted packet transmission, the relays were repositioned.

With the receiver-directed system, the node closest to the center of the circle is always preselected to act as the forwarding relay and the other $M - 1$ relays simply ignore the transmission. For the broadcast-oriented system, the message is broadcast to all M nodes, some (or none) of which receive the message correctly (i.e. the *potential* relays). Two ideal forwarding scenarios were considered. In the first scenario, the potential relay with the best instantaneous channel (i.e. highest instantaneous relay-destination SNR) forwards the message, while in the second scenario the one closest to the destination (i.e. highest average relay-destination SNR) forwards it. If the forwarding relay's message fails to reach the destination correctly, no other relays transmit and an end-to-end error is logged. In practice, performance could be further improved by allowing the other

potential relays to successively attempt a transmission (ordered by decreasing instantaneous or average SNR) until all fail.

The simulated system uses packets of length $N = 80$ bits transmitted at a rate of 1 Mbaud over a quasi-static Rayleigh fading channel using noncoherently detected frequency shift keying (NFSK)¹. It is assumed that the transceivers are sufficiently separated so that the channels from source to M relays were statistically independent (likewise for the channels from relays to destination). The noise spectral density is $N_o = 10^{-18}$ W/Hz and the average received power at the m^{th} transceiver is $P_r = K_o (d_m/d_o)^{-n} P_t$, where $d_o = 1$ m is a reference distance, d_m is the transmitter-receiver separation, n is the path loss exponent, K_o is the channel power gain at the reference distance, and P_t is the transmitted power. Corresponding to a typical indoor channel [12], $n = 3$, and for a 2.4 GHz operating frequency, $K_o = (c/4\pi d_o f_c)^2 \approx 10^{-4}$.

No FEC encoding is used, and thus the packet will only be received correctly only if all N bits are correct (neglecting the overhead for synchronization). Thus, the instantaneous packet error probability at the m^{th} transceiver is

$$\begin{aligned} p_e(\gamma_m) &= 1 - (1 - p_b(\gamma_m))^N = 1 - \left(1 - \frac{1}{2}e^{-\gamma_m/2}\right)^N \\ &= -\sum_{k=1}^N \binom{N}{k} \left(-\frac{1}{2}\right)^k e^{-k\gamma_m/2}, \end{aligned} \quad (1)$$

where γ_m is the instantaneous SNR of the channel. The average packet error rate in fading is found by taking the expected value of $p_e(\gamma_m)$ with respect to the SNR [11]:

$$\begin{aligned} \bar{p}_e(\Gamma_m) &= \int_0^\infty \frac{1}{\Gamma_m} e^{-\gamma_m/\Gamma_m} p_e(\gamma_m) d\gamma_m \\ &= \sum_{k=1}^N \binom{N}{k} \left(-\frac{1}{2}\right)^{k-1} \left(\frac{1}{2 + k\Gamma_m}\right), \end{aligned} \quad (2)$$

where Γ_m is the average SNR of the channel.

The transmit power of the source (P_1) and relay (P_2) were varied independently in 0.5 dB increments, and enough transmissions were attempted for each (P_1, P_2) pair to log 1000 independent errors. The target end-to-end error probability was 10^{-2} and it was noted which (P_1, P_2) pairs satisfied this requirement. Fig. 2 shows a set of contours describing the results of the simulation. The contour plot shows (for each of the three types of systems) which combinations of (P_1, P_2) achieve the desired end-to-end error probability for several values of M . (P_1, P_2) pairs that lie to the upper-right of a contour will satisfy the error requirement, while those to the lower-left will violate it. The least power efficient scenario occurs when there is just a single relay ($M = 1$), in which case broadcasting is equivalent to receiver-directed transmission. The energy efficiency of the receiver-directed strategy improves slightly with increasing M , but only because as the cloud becomes denser, the forwarding

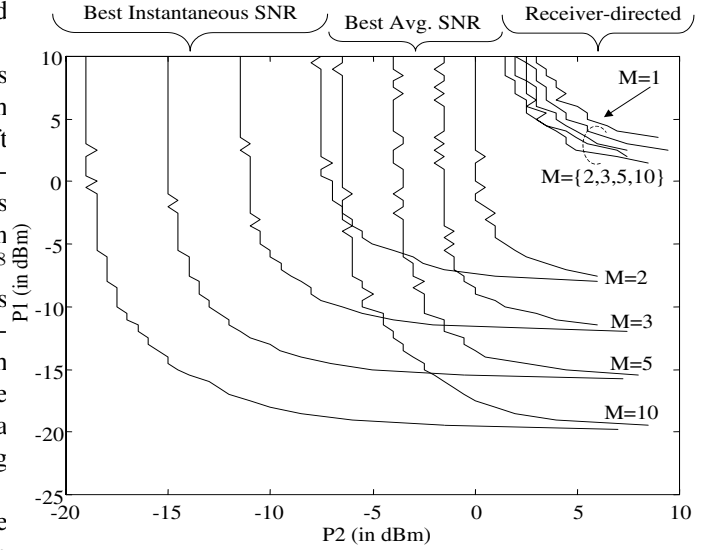


Fig. 2. Contours showing the minimum combination of source and relay powers (P_1, P_2) required to achieve an end-to-end packet error probability of 10^{-2} over the relay channel.

M	receiver-directed	best avg. SNR	best inst. SNR
1	6.692 mW	6.692 mW	6.692 mW
2	5.968 mW	1.611 mW	0.632 mW
3	5.623 mW	0.990 mW	0.284 mW
5	5.047 mW	0.642 mW	0.134 mW
10	4.798 mW	0.382 mW	0.067 mW

Table I: Minimum total transmit power $P_1 + P_2$ required to guarantee an average end-to-end packet error probability of 10^{-2} in a relay channel for several values of M (number of intermediate nodes in the cloud) for the receiver-directed system and the broadcast-oriented systems (using either the relay with best average SNR or with the best instantaneous SNR).

relay selected by the source is more likely to be near the cloud's center. This improvement is not due to any diversity advantage.

On the other hand, the energy efficiency of the two proposed broadcast-oriented strategies improve dramatically with increasing M . While some of this improvement is due to the clustering of relays in the center of the cloud, most of it comes from the diversity effect of receiving the source transmission at all M relays. Because both broadcast-oriented strategies implement receiver diversity in the source-relay transmission, the asymptotic source transmit power requirements P_1 are roughly the same as $P_2 \rightarrow \infty$. However, the system which uses the potential relay with best instantaneous SNR has a better diversity effect for the relay-destination transmission and thus requires smaller relay transmit powers P_2 than the system which just uses the best average SNR. However, using best average SNR is more practical because it only requires knowledge of the average SNR (which could be related to the node's distance to the destination) rather than the instantaneous SNR (which would be difficult to measure/predict in a nonreciprocal channel).

By selecting an appropriate (P_1, P_2) operating point, the system can divide the power requirements between the source and relay in a manner that takes into account the energy available at each node. The results of Fig. 2 indicate that the proposed strategy is robust to the power allocation strategy. For instance,

¹The quasi-static Rayleigh fading assumption implies that the channel SNR remains constant for the duration of each packet, but varies independently from packet to packet according to an exponential distribution. Short packets were chosen because they are common in wireless sensor networks and because direct evaluation of (2) becomes numerically unstable for $N > \approx 100$. The bit error probability of NFSK in AWGN is $p_b = \frac{1}{2}e^{-\mathcal{E}_b/2N_o}$ [11].

when $M = 10$ and the potential relay with best instantaneous SNR is used to forward the message, the source could decrease its power to -18 dBm if the relay increases its power to 0 dBm. This is because as the source reduces its transmit power, the elected relay tends to be closer to it and therefore requires a higher transmit power to reach the destination. Conversely, if the source increases its power, the forwarding relay will tend to be closer to the destination and could afford to transmit at lower power.

The minimum total power requirement $P_1 + P_2$ for each system is given in Table I. This table assumes that no premium on energy consumption is placed on one node over the other, and that instead the goal is to minimize the *sum* of the two powers (although this minimization could be easily modified to account for unequal energy reserves). By just increasing the number of intermediate nodes from $M = 1$ to $M = 2$, the total transmit power requirement is reduced from 6.692 mW to 1.611 mW or 0.632 mW, when using the best average or instantaneous SNRs, respectively. When $M = 10$ the requirement is reduced to 0.382 mW or 0.067 mW, which indicates that a reduction in total transmit power of one or two orders-of-magnitude is possible by properly exploiting macrodiversity in a dense relay channel.

V. PRACTICAL CONSIDERATIONS

The goal of the previous section was to illustrate the potential energy savings that could be achieved by exploiting macrodiversity, but is not meant to suggest a practical means of achieving these savings. Several issues must be taken into account in order to approach the promised energy efficiencies. Several of the more important issues, which are the topic of ongoing research, are listed below:

Channel models: The previous analysis assumed that the M channels experienced flat Rayleigh fading and were both spatially and temporally (on a per-packet basis) uncorrelated. While Rayleigh fading is often used to gauge worst case performance, it assumes the presence of only specular energy [12]. Because many channels also have an additional line-of-sight (LOS) component, Rician or Nakagami fading distributions could provide a more accurate model [13]. For the low channel signaling rates expected of embedded networks, frequency-selectivity is not anticipated for unspread systems. However, as discussed below, there are advantages to using direct-sequence spread-spectrum (DS-SS) in embedded networks, and the corresponding analysis will require frequency-selectivity to be taken into account [14]. In some situations, such as when the LOS component is very strong or channel coding is able to overcome temporal variations, performance may be dominated by (log-normal) shadowing rather than small-scale fading [12]. In this case, the channel might not be independent from packet to packet, and additional temporal decorrelation may be necessary in the form of slow frequency hopping (as in Bluetooth) or interleaving. In a sparse network, it is reasonable to assume that the M channels are uncorrelated, but as the density increases there may be peer-to-peer spatial correlation [15]. Thus, the advancement of this research will require that the impact of these channel effects be incorporated into the analysis with the primary goal of determining what types of channels gain the most from

macrodiversity-combining (and which don't) and a secondary goal of suggesting practical methods of using macrodiversity (possibly in conjunction with other simple signal processing techniques) to improve performance over actual channels.

Combined MAC/network-layer protocol: The previous analysis assumed that the potential relays were able to automatically elect an appropriate forwarding relay. Ideally, the relay with the best instantaneous relay-destination SNR should be used, but in practice the one with the best average SNR could be used. In the absence of shadowing, this is the node closest to the destination. While such a negotiation does not require the actual packet to be shared among the relays, it does require that each potential relay be aware of its distance to the destination and if any nodes closer to the destination also received the source's transmission correctly. If there is a closer potential relay, the node defers transmission, otherwise it forwards its message. If multiple nodes are simultaneously at the closest distance to the destination, they would have to agree on which one forwards. All the necessary information to perform this negotiation could be explicitly conveyed through a control channel, and for large source packets the resulting overhead could be acceptable. However, when the source transmits shorter packets (as in a sensor network) the overhead would quickly become excessive and an alternative mechanism is desirable. Furthermore, it is also desirable for the negotiation policy to take into account the amount of energy that each relay can afford to expend when forwarding.

The solution that we advocate is to incorporate the negotiation process into the medium access control (MAC) protocol. For purposes of exposition, assume the following: (1) Acknowledgments are broadcast from the destination to prospective relays over an error- and delay-free return channel, (2) Each relay perfectly knows its distance to the source and destination, (3) Each relay assigns itself an *energy penalty* that quantifies the importance of conserving energy during the next data transmissions (which could be inversely proportional to its energy reserve), and (4) The destination can detect the collision of signals intended for it. To resolve the channel contention that arises when multiple nodes correctly receive the source message, each potential relay should wait a specified amount of time before attempting to forward the message. As carrier sense multiple access (CSMA) also requires nodes to wait before (re)transmitting, these two mechanisms could be combined (along with the ARQ and power control strategy) within the MAC layer. However, unlike traditional CSMA algorithms, the wait time should be equal to some (possibly nonlinear) weighted function of the node's distance to the destination and its energy penalty. This policy will assure that nodes closer to the destination will generally try to relay the message first, as will nodes with more plentiful energy sources. However, there could still be contention when multiple nodes have the same energy-penalty/destination-distance function. With narrowband signaling, this initial collision is inevitable, but the likelihood of subsequent collisions can be reduced by having the prospective nodes use a different function to reset their timers. This function could include side information, for instance the distance of the node to the *source*.

An alternative solution to the problem of concurrent trans-

missions would be to use direct sequence spread spectrum (DS-SS) signaling. Due to the capture effect of DS-SS [10], only one of the concurrently transmitted packets will typically be detected when using a conventional receiver. Using a coherent or noncoherent Rake receiver at the destination is particularly appealing for this type of a network, because it will allow *transmit diversity* in the final hop. With a Rake receiver it could be advantageous to actually *encourage* many prospective relays to transmit simultaneously in order to achieve richer transmit diversity. A simple policy would be to have each prospective node immediately forward the packet if it is within a certain radius of the destination, regardless of what the other nodes are doing.

In our simulations, the set of intermediate nodes were confined to a 7 m wide circular cloud. This region was chosen because we observed no advantage to using relays outside of it. Relays that are too far from the destination require too much power, while if the relay is too close to the destination then it would generally be better for the source to simply transmit directly to the destination rather than over a cascade of two binary symmetric channels. However, this region was only chosen heuristically with no attempt to optimize it. A more detailed study should involve a more formal definition of the region of diversity advantage. In an actual implementation, the source will broadcast to all nodes within a certain radius. Each node that receives the message should only participate in the negotiation process if it is inside the region of diversity advantage. This decision need not be hard, as nodes could use their position relative to the center of the region as additional side information during the implied negotiation. Note that this policy entails a *distributed routing* algorithm, and when extrapolated to ad hoc networks, the policy could be modified to incorporate various position-aware routing strategies [16].

The performance of the system will be highly dependent on the reliability of the return channel, which until now has been assumed to be error- and delay-free. When the return channel is not perfect, it is important to ensure that none of the more distant potential relays transmit after the packet has already arrived correctly at the destination. One option would be to simply transmit a return message at high power and/or with strong coding, although this is wasteful of network resources. An alternative approach would involve the use of *passive acknowledgments* [10], i.e. when a node forwards the packet, its neighboring nodes will overhear the transmission and know not to transmit. An actual implementation of a macrodiversity-combining protocol is likely to involve some combination of passive and active acknowledgments.

Position location: Throughout the discussion, it has been assumed that every transceiver knows its relative location. While location estimates could be obtained by integrating a Global Positioning System (GPS) receiver into each node, it is too expensive to do this for every node and not all nodes will be in a position to receive GPS signals. An alternative is to use the received signal strength indicator (RSSI) at each node to estimate position. In [17], a maximal likelihood (ML) solution is presented for the problem of obtaining position estimates of all nodes in the presence of log-normal shadowing when the absolute position of just a small fraction of the nodes are known. The estimation error of this approach is also given in [17], and

should be incorporated into the analysis of the macrodiversity-combining protocol (although the analysis will need to be extended to accommodate other channel models). One drawback of the ML approach is that it requires the solution to be solved at a central location, which is not desirable for large embedded networks. Thus, the algorithm should be modified to allow a more scalable distributed solution of the ML position estimation problem. One possible implementation would involve the use of the turbo-principle [18] to decompose the problem into several locally optimal solutions and then use iteration to obtain the global solution.

Other energy costs: So far, the focus has been on minimizing *transmit* energy in the network. However the energy required to *receive* data in an embedded network can often be nearly as much as it is to transmit [19]. Furthermore, the transmitter circuitry will consume more power than just that radiated by its antenna. Thus, these extra energy expenditures should also be factored into the model. [20] discusses energy issues for receiver-directed low power wireless networks, which could be extended for broadcast-oriented macrodiversity-combining networks. This modification is expected to influence the design of the protocol and provide some guidance regarding when the proposed strategy should be used and when it shouldn't. For transmit distances smaller than a certain threshold, it might make sense to only use a receiver-directed transmission, while for longer range transmissions the proposed broadcast strategy may prove advantageous (e.g. to create a *virtual bridge* between two widely separated clusters of nodes).

Node activity: In a wireless sensor network, nodes are often placed into a sleep state and only awoken in response to external stimulus or according to a rendezvous schedule [21]. Thus, although the network may have a high density of devices, the fraction of those that are awake and can serve as relays may be low. The irregular node activity could actually be viewed as a type of *time diversity* that should be exploited. The analysis could be extended to model the effects of irregular node activity by modelling the cloud of relays as a set of nodes with random activity. In fact, the analysis presented earlier could be viewed as already incorporating a notion of node activity; in the simulation, the M intermediate nodes were in a different location during each source broadcast and this repositioning could be interpreted as changes in node activity rather than actual node movement. However, the simulation assumed that the same number of nodes were present during each source transmission while in practice this number will itself be a random variable.

Extension to multihop networks: While the discussion has focused on the relay channel, the proposed macrodiversity-combining concepts can be applied to multihop networks. Related work has been pursued by Wieselthier et al in [22] which proposes multicasting protocols that exploit the broadcast-oriented nature of wireless. One key issue is that of range and power control, i.e. how power should be allocated among the many nodes that form the chain of relays. One analytical tool that could be used to answer this question is a generalization of the two-dimensional contour plot shown in Fig. 2 to a multi-dimensional region of operation for the entire network (one dimension per hop). This analysis is related to earlier work on

optimal range control for receiver-directed multihop networks [23], [24], and should take into consideration the possibility of hop-to-hop correlation in the channel [15]. Our earlier analysis showed that the system was robust to the allocation of power between source and relay, and it is anticipated that a multihop system would have a similar robust behavior. In fact, strict power/range control might not even be required by the system. An alternative strategy would be to simply let each node broadcast with as much energy as it can afford and let whichever node(s) that receive it forward it.

Furthermore, a multihop network will be sensitive to unreliable return channels. When the return channel is unreliable there is the potential for stray messages to begin propagating through the network along multiple routes (i.e. the so-called broadcast storm problem [25]) because distant nodes didn't know that a closer node successfully forwarded its message. While it seems that this is a behavior that should be discouraged by the protocol, it may in fact be beneficial in a dense network because of the *routing* diversity it provides [26].

VI. CONCLUSION

The results presented in this paper indicate that an order-of-magnitude reduction in transmit power is possible by properly exploiting the macrodiversity inherent in dense ad-hoc networks, and that these gains can be achieved with simple decentralized processing and cross-layer protocols. However, these highly encouraging findings must be taken with a grain of salt. While *transmitter* power is indeed conserved, this comes at the potential cost of increased *receiver* power. In addition, the extremely low activity cycle of embedded networks is an issue, and the actual channel may not be characterized by quasi-static Rayleigh fading. However, there will certainly be situations (intentional or unintentional) when a node's transmission is overheard by more than just its intended receiver, and the overhearing nodes may be in a reasonable position to act as relays. At this point, each overhearing node has already consumed energy by attempting to receive the message, and it would be wasteful for it to do nothing if the message is correct. Therefore, it is our assertion that *whenever the opportunity to exploit macrodiversity presents itself, it should be taken*. A conservative approach would be to simply make opportunistic use of redundant nodes whenever they happen to be available, although a more aggressive approach would be to proactively manage the network (e.g. by putting the nodes that comprise a virtual antenna array into the same sleep schedule) to intentionally set up situations where macrodiversity could be exploited.

REFERENCES

- [1] National Research Council (Computer Science and Telecommunications Board), *Embedded, Everywhere*. Washington, DC: National Academy Press, 2001.
- [2] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Commun.*, vol. 6, pp. 311–335, Mar. 1998.
- [3] S. Sendonaris, E. Erkip, and B. Aazhang, "Increasing uplink capacity via user cooperation diversity," in *Proc. IEEE Int. Symp. on Inform. Theory (ISIT)*, (Cambridge, MA), p. 156, Aug. 1998.
- [4] J. N. Laneman, "Cooperative diversity in wireless networks: Algorithms and architectures," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA, Aug. 2002.
- [5] Z. J. Haas and C.-P. Li, "The multiply-detected macrodiversity scheme for wireless cellular systems," *IEEE Trans. Veh. Tech.*, vol. 47, pp. 506–530, May 1998.
- [6] J. Jubin and J. Tornow, "The DARPA packet radio network protocol," *Proc. IEEE*, vol. 75, pp. 21–32, Jan. 1987.
- [7] J. N. Laneman and G. W. Wornell, "Energy-efficient antenna sharing and relaying for wireless networks," in *IEEE Wireless Commun. and Networking Conf.*, (Chicago), pp. 7–12, Sept. 2000.
- [8] T. M. Cover and A. A. El Gamal, "Capacity theorems for the relay channel," *IEEE Trans. Inform. Theory*, vol. 25, pp. 572–584, Sept. 1979.
- [9] J. Boyer, D. Falconer, and H. Yanikomeroglu, "A theoretical characterization of the multihop wireless communications channel with diversity," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, (San Antonio, TX), pp. 841–845, Nov. 2001.
- [10] G. Lauer, "Packet-radio routing," in *Routing in Communications Networks* (M. Steenstrup, ed.), ch. 11, pp. 351–396, Englewood Cliffs, NJ: Prentice Hall, 1995.
- [11] M. K. Simon and M. S. Alouini, *Digital Communications over Fading Channels: A Unified Approach to Performance Analysis*. New York: John Wiley & Sons, Inc., 2000.
- [12] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall PTR, second ed., 2002.
- [13] M. D. Yacoub, "Fading distributions and co-channel interference in wireless systems," *IEEE Antennas and Propagation Magazine*, vol. 42, pp. 150–159, Feb. 2000.
- [14] A. A. M. Saleh and R. A. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Select. Areas Commun.*, vol. 5, pp. 128–137, Feb. 1987.
- [15] N. Patwari, Y. Wang, and R. J. O'Dea, "The importance of the multipoint-to-multipoint indoor radio channel in ad hoc networks," in *IEEE Wireless Commun. and Networking Conf.*, (Orlando, FL), pp. 608–612, Mar. 2002.
- [16] I. Stojmenovic, "Position-based routing in ad hoc networks," *IEEE Commun. Magazine*, vol. 40, pp. 128–134, July 2002.
- [17] N. Patwari, R. J. O'Dea, and Y. Wang, "Relative location in wireless networks," in *Proc. IEEE Veh. Tech. Conf. (VTC)*, (Rhodes, Greece), pp. 1149–1153, May 2001.
- [18] J. Hagenauer, "The turbo principle: Tutorial introduction and state of the art," in *Proc., Int. Symp. on Turbo Codes and Related Topics*, (Brest, France), pp. 1–11, Sept. 1997.
- [19] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *IEEE INFOCOM*, (New York, NY), June 2002.
- [20] J. Rabaey, J. Ammer, J. L. da Silva, and D. Patel, "PicoRadio: Ad-hoc wireless networking of ubiquitous low-energy sensor/monitor nodes," in *Proc. IEEE Comp. Society Workshop on VLSI*, (Orlando, FL), pp. 9–12, Apr. 2000.
- [21] T. D. Todd, F. Bennett, and A. Jones, "Low power rendezvous in embedded wireless networks," in *Proc. Int. Workshop on Mobile Ad Hoc Networking & Computing*, Aug. 2000.
- [22] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "Resource management in energy-limited, bandwidth-limited, transceiver-limited wireless networks for session-based multicasting," *Computer Networks*, vol. 39, pp. 113–131, 2002.
- [23] T.-C. Hou and V. Li, "Transmission range control in multihop packet radio networks," *IEEE Trans. Commun.*, vol. 34, pp. 38–44, Jan. 1986.
- [24] H. Takagi and L. Kleinrock, "Optimal transmission ranges for randomly distributed packet radio terminals," *IEEE Trans. Commun.*, vol. 32, pp. 246–257, 1984.
- [25] S. Y. Ni, Y. C. Tseng, Y. S. Chen, and J. P. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proc. ACM MobiCOM*, pp. 151–162, Aug. 1999.
- [26] N. Roux, J. S. Pegon, and M. W. Subbarao, "Cost adaptive mechanisms to provide network diversity for MANET reactive routing protocols," in *Proc. IEEE Military Commun. Conf. (MILCOM)*, (Los Angeles), pp. 387–291, Oct. 2000.