Barrage Relay Networks

Thomas R. Halford*
*TrellisWare Technologies, Inc.
16516 Via Esprillo, Suite 300
San Diego, CA 92127–1708
Email: {thalford,kchugg}@trellisware.com

Keith M. Chugg†*
†Communication Sciences Institute
Ming Hsieh Dept. of Electrical Engineering
University of Southern California
Los Angeles, CA 90089–2565
Email: chugg@usc.edu

***Abstract*—A receiver-oriented perspective on capacity scaling in mobile ad hoc networks (MANETs) suggests that broadcast and multicast may be more natural traffic models for these systems than the random unicast pairs typically considered. Furthermore, traffic loads for the most promising near-term application for MANET technology – namely, networking at the tactical edge – are largely broadcast. The development of novel MANET approaches targeting broadcast first and foremost, however, has not been reported. Instead, existing system designs largely rely on fundamentally link-based, layered architectures, which are best suited to unicast traffic.**

In response to the demands of tactical edge communications, TrellisWare Technologies, Inc. developed a MANET system based on Barrage Relay Networks (BRNs). BRNs utilize an autonomous cooperative communication scheme that eliminates the need for link-level collision avoidance. The fundamental physical layer resource in BRNs is not a link, but a portion in space and time of a cooperative, multihop transport fabric. While initial hardware prototypes of BRNs were being refined into products by TrellisWare, a number of concepts similar to those that underlie BRNs were reported independently in the literature. That TrellisWare's tactical edge MANET system design and academic research reconsidering the standard networking approach for MANETs arrived at similar design concepts lends credence to the value of these emerging wireless network approaches.

I. Introduction

The past decade has seen a significant research focus on both the theoretical capabilities of and protocol designs for mobile ad hoc networks (MANETs). Efficient unicast data transport has been the primary focus of this effort. The most common formulation for capacity scaling analysis, for example, assumes sufficient node density so as to ensure link-level connectivity and considers the largest rate that can be guaranteed between randomly selected source-destination pairs. Gupta and Kumar's seminal result [1] shows that this rate falls quickly as the number of network nodes increases. One consequence of this result is that protocols seeking to maximize efficiency for randomly paired unicast traffic must be carefully designed owing to the effective reduction in available bandwidth with network growth. Indeed, MANET protocol designers have sought to do just this by carefully designing link-level medium access control (MAC) and multihop routing algorithms that seek to minimize the number of link transmissions while avoiding collisions. Another potential consequence of Gupta and Kumar's result is that networks without infrastructure are inherently ill-suited to the randomly-paired unicast traffic model. Although a simplification, this interpretation should not come as a surprise since interference associated with concurrent wireless transmissions must be avoided and routing of data packets utilizes intermediate nodes that are themselves sources and destinations.

Much of the impetus for the research activity described above has been provided by the potential applications for MANETs, one of the most compelling of which is communications at the *tactical edge*. A squadron of soldiers seeking to maintain connectivity in a challenging RF propagation environment (e.g., urban canyons, ships, subterranean structures, etc.) is an example of edge networking; first-responder communications, such as remote search and rescue, is another. This *tactical* MANET application assumes no supporting infrastructure and often no means for communicating outside of the squadron. Since nodes are both mobile and typically in rich scattering environments, link-level connectivity is unreliable and the network topology is highly dynamic. The traffic patterns and key performance metrics for tactical MANETs also differ substantially from those typically considered for ad hoc or sensor networks. Specifically, low-latency, network-wide broadcast and robust connectivity are the primary requirements of tactical MANET systems, while typical traffic includes interactive push-to-talk (PTT) voice and real-time video streaming from a small set of source nodes.

In light of the design goals for tactical MANET, it is useful to consider the throughput and latency scaling properties of broadcast traffic. It is described in Section II of this paper how reinterpreting existing scaling results suggests that ad hoc networks may be better suited for broadcast and multicast than for the more commonly considered unificast traffic models. In particular, the aggregate useful data rate received in the network scales linearly with the number of nodes for broadcast, which is not unexpected since all relaying nodes are also intended recipients. Despite this apparent match to broadcast traffic and the important tactical MANET application, there has been relatively little effort committed to designing MANET systems that target broadcast first and foremost.

Section III of this paper describes the basic concepts of Barrage Relay Networks (BRNs), which are a type of MANET designed from the ground up for the demands of tactical edge communications. BRNs utilize autonomous cooperative communications to enable packets to ripple out from source nodes rapidly and reliably through the network. Each node, upon receiving a packet, repeats the data as part of one or more teams of cooperative relays. In this manner, simultaneous wireless transmissions serve to improve reliability and are not avoided by any access control mechanism. BRNs do not utilize a point-to-point link abstraction; rather, a segment of the cooperative transport fabric in time and space is the fundamental physical layer resource to be controlled. Section III of this paper also describes access control for this cooperative transport resource. Section IV revisits broadcast scaling to show that BRNs scale optimally for this traffic model. Section V describes how the efficient barrage flooding mechanism can be contained for unicast or localized multicast traffic via *controlled barrage regions* (CBRs).

BRNs were developed by TrellisWare Technologies, Inc. to address tactical MANET applications with the first hardware demonstrations in 2004, followed by product development. During this same period a number of related concepts have been described in the open literature. Section VI of this paper summarizes this related work. In describing BRNs, the present paper details one approach for combining these suggested techniques in a functional tactical MANET design.

II. Broadcast as a Fundamental Service

A. Transmitter-Oriented Scaling

Gupta and Kumar's [1] seminal research on capacity scaling in wireless networks has spurred much interest in the fundamental limits of such systems (cf., [2] for a recent summary). Specifically, assuming a node density sufficiently high to provide connectivity with an abstract physical (PHY) layer model and a traffic model forming n random unicast pairs in an n node network, it was shown in [1] that the throughput obtainable by each node scales as

$$\Theta\left(\frac{1}{\sqrt{n\log n}}\right). \tag{1}$$

Wang, *et al.* [3] presented a framework for throughput capacity computation that unifies Gupta and Kumar's results with those for broadcast (e.g., [4]) and multicast (e.g., [5]) traffic models. Summed over all source nodes, the sum throughput capacity of m-destination multicast is

$$t(n,m) = \begin{cases} \Theta\left(\sqrt{\frac{n}{\log n}}\right) & m = 1 \\ \Theta\left(\sqrt{\frac{n}{m\log n}}\right) & 1 < m \leq \Theta(n/\log n), \\ \Theta(1) & m > \Theta(n/\log n) \end{cases} \tag{2}$$

where the $m = 1$ case corresponds to unicast and is simply the sum throughput version of (1).

These transmitter-oriented scaling results characterize the rate at which data can be injected into to the network by source nodes rather than the rate at which intended nodes receive the data. Expressions such as (1) and (2) can be viewed as fundamentally *negative* results since they state that the transmitted data rate available for a given node decreases rapidly with growing network size. Intuitively, it is no surprise that unicast data rates do not scale well since nodes must *locally* coordinate transmissions to avoid interference on the wireless medium.[1] Furthermore, when a node relays a packet for which it is not an intended recipient, it cuts into the rate available for data relevant to the relaying node. Given this context, (1) and (2) can also be viewed as defining a challenge: how to design access control and routing protocols that make most efficient use of the diminishing data rate capabilities.

B. Receiver-Oriented Scaling

For broadcast traffic the intuitive reasons for poor scaling given above do not hold – i.e., all relay nodes are also intended destinations. To capture this effect it is useful to reinterpret the transmitter-oriented scaling results in terms of the total amount of data delivered to the intended destination nodes. For the randomly paired unicast traffic model, this is the same as the aggregate transmitted rate given in (2). For the m-destination multicast case, however, (2) implies

$$r(n,m) = \begin{cases} \Theta\left(\sqrt{\frac{n}{\log n}}\right) & m = 1 \\ \Theta\left(\sqrt{\frac{nm}{\log n}}\right) & 1 < m \leq \Theta(n/\log n) \\ \Theta(m) & m > \Theta(n/\log n) \end{cases}, \tag{3}$$

where $r(n,m)$ is the aggregate received data rate when n sources each transmit to m randomly chosen destinations.

For the case of broadcast, $r(n,n) = \Theta(n)$ so that the aggregate useful received data rate increases linearly with the size of the network. Thus, from the perspective of sum information reception, broadcast is clearly more appealing than unicast. The simple conversion of (2) to (3) is not intended to uncover anything fundamental about wireless network scaling *per se*. Rather, it is meant to highlight the fact that, since broadcast and multicast deliver useful information to many destination nodes simultaneously, ad hoc networks may be inherently better matched to these traffic models than to randomly paired unicast. Furthermore, identifying applications that are dominated by broadcast traffic and designing protocols that target it from the ground up may prove fruitful.

The challenge posed to the networking community by this receiver-oriented scaling result is the design of low-overhead broadcast protocols that scale optimally. It is likely that broadcast applications will require less than n contemporaneous data sources; scalability therefore requires that the large m case of (3) be met when any subset of $s \leq n$ source nodes transmit. Although asymptotically optimal broadcast protocols have been proposed in the literature (e.g., [7]), the tendancy to rely on a virtual backbone (i.e., a connected dominating set of nodes) to limit redundant transmissions presents a larger control overhead management challenge than does unicast [8].

C. Usage Considerations and Tactical MANETs

Asymptotic scaling results may have limited relevance to practical MANET deployment for several reasons. First, traffic models and node distributions are likely to deviate from those assumed in tractable analyses. For example, scalability could

[1] Note that the focus here is on networks with topological dynamics induced by mobility and link outage that likely preclude the use of those complex coordination schemes shown to achieve greater unicast capacity (e.g., [6]).

be expected to improve in networks with localized unicast traffic. Second, in many applications the size of the ad hoc network is limited, with connectivity on a larger scale provided by hierarchical designs. Third, control overhead is neglected in scaling analysis but can be a significant, even limiting, factor in practical networks. Specifically, it is unclear whether the potential improvement in wireless network scaling for unicast traffic afforded by extremely sophisticated medium access and routing protocol designs justifies the control overhead penalties that will invariably be incurred in order to support their implementation. While the mathematical tools to address this question formally are as yet nascsent (cf., [9]), anecdotal evidence from practical implementations suggest that the answer may be no [9, 10]. Finally, in some applications metrics other than throughput for a highly loaded network (e.g., latency, fidelity, robustenss, etc.) are paramount.

Notwithstanding these limitations, the scaling results afforded by asymptotic analysis are valuable in guiding the top-level design and deployment of wireless networks. In particular, as described above, ad hoc networks may be a good match to broadcast intensive applications. Despite this, the utility of broadcasting in wireless networks is commonly motivated by its role in network maintenance such as route table distribution (cf., [11]) – i.e., broadcast is often viewed as important inasmuch as it supports unicast transmission.

A number of applications are candidates for broadcast-oriented wireless network designs. Coordinated and distributed computing in sensor networks provides one example [7]; applications that currently flood at the *application* layer (e.g., query response and target tracking) provide another. The tactical edge applications that are the focus of the present paper provide a particularly important venue for one-to-many communication. For example, real-time video streaming over multi-hop wireless networks for robotics platform control and surveillance is a critical emerging application for tactical MANETs operating in restricted environments

A typical illustrative tactical MANET scenario is that of a squadron of soldiers exploring an urban environment and seeking to maintain real-time communication capabilities. The environment may include multistory buildings and subterranean structures. One critical service is push-to-talk, multi-hop voice, for which there is a strict latency requirement. Another is sharing of real-time sensor data – e.g., one node streaming audio and video data throughout the network. Since the soldiers are expected to move through the environment either by foot or motor vehicle, the propagation channel is challenging. Specifically, time-varying, frequency-selective fading channels can be expected, as well as abrupt path loss changes due to corner effects. In additional to physical layer challenges, this induces rapid changes to the network topology.

In summary, tactical MANETs typically[2] have sizes of 5 to 100 nodes and operate without any fixed infrastructure (often in isolation). The mission criticality of the data drives latency and robustness as the key performance metrics. Furthermore, the traffic in tactical MANETs is primarily broadcast and multicast, although a limited amount unicast traffic may need to be supported concurrently. This application therefore motivates the development of protocols that are broadcast-oriented and highly robust to dynamics in the network topology.

[2]The term "tactical MANET" has also been used to describe networks for tactical applications comprising a heterogeneous mix of local radio networks, satellite networks, and gateways.

III. Barrage Relay Networks

In this section, barrage relay networks are proposed as an efficient scheme for broadcasting in mobile ad hoc networks. BRNs are a simplified abstraction of a tactical MANET system developed by TrellisWare Technologies, Inc. with hardware system demonstrations beginning in 2004 and leading to a current product line. This section summarizes the basic concepts underlying BRNs and relies heavily on [12–14].

The barrage relay concept is best understood through the simple example which follows. A number of network capabilities are required for BRNs. While these are described in more detail following the example, the two most critical assumptions for the example are those of a time division multiple access (TDMA) network and a method of *autonomous cooperative communications*. The assumption that all nodes utilize a common *TDMA framing* format[3] requires coarse slot-level synchronization, which can be accomplished using low overhead pilot signaling [13] and aided by the Global Positioning System (GPS) when available. Autonomous cooperation ensures that concurrent transmission of identical packets results not in collision at all nodes within range, but in a form of cooperative diversity; this is done with no coordination beyond the TDMA-level synchronization.

A. Barrage Relay via Example

With these assumed capabilities, the BRN broadcast mechanism is illustrated in Figure 1. Generally, M-slot TDMA framing is assumed for some $M \geq 3$; Figure 1 assumes that $M = 3$ with slots labeled A, B, and C. Suppose the black node transmits a packet on slot A of the first TDMA frame. All nodes that successfully receive this packet are, by definition, one hop away from the source node. These nodes then transmit the same packet on slot B, thus relaying to the nodes that are two hops away from the source, which in turn transmit the same information on slot C. Nodes that are 3 hops away from the source relay on slot A of the second TDMA frame. Packets thus propagate outward from the source via a decode-and-forward approach. To prevent the relay transmissions from propagating back towards the source, each node relays a given packet only once. For example, one-hop nodes will receive the first broadcast packet on slot A and again on slot C but only relay on slot B. This can be enforced by an explicit mechanism (e.g., by maintaining a history of received packets) or implicitly via protocol header design.

[3]Conceptually, the barrage relay approach can be applied to channelization methods other than time division. Time division provides the most intuitive illustration of BRNs and also has some practical advantages in terms of interference isolation in the absence of any precise power control mechanism.

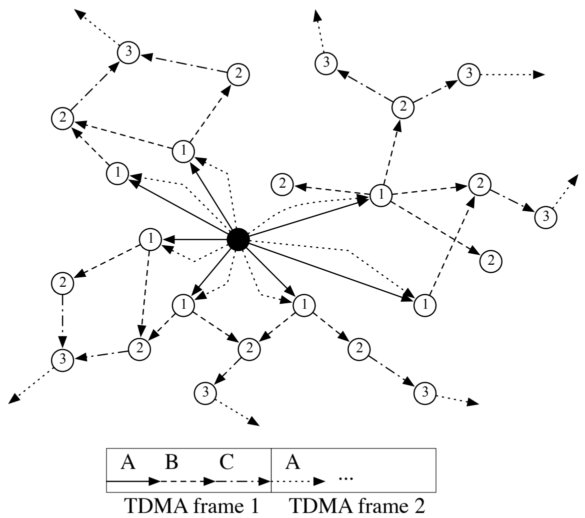

Fig. 1. Barrage relay network broadcast protocol for a three slot ($M = 3$) TDMA frame format. The source node is black while relay nodes are numbered by their distance in hops from the source.

Observe that a number of two-hop nodes in Figure 1 receive the same packet on the same slot, but from different one-hop nodes. Due to the assumed autonomous cooperative communication scheme, these packets neither collide nor result in destructive interference. An arbitrary number of relays can thus be cooperatively transmitting to one or more nodes without any coordination other than TDMA slot synchronization. Furthermore, a given node may participate in several such cooperating transmission *teams*, without requiring knowledge about, or even of, its possible teammates.

The spatial reuse of time slots enables packets to be pipelined into the source for transmission every M slots. For example, in Figure 1, the one-hop nodes will not receive the packet transmitted by the three-hop nodes during slot A of the second TDMA frame. Thus, the source can safely transmit a second packet during that slot. It is readily verified that M must be at least 3 in order to allow such *spatial pipelining*. Larger values of M can be chosen so as to trade throughput for enhanced robustness to topological variation.

B. Autonomous Cooperation

The *autonomous cooperative communication* capability assumed can be provided by the method described in [12] which is based on phase dithering (cf., [15]) and modern, turbo-like error correction (cf., [16]). Specifically, if each transmitting node pseudo-randomly dithers its carrier phase, then the superposition of these signals will induce a time-varying channel characteristic at a receiving node. A modern error correction code can then be used to extract the time diversity provided by this induced time-varying fading channel. Other parameters of the signal could be dithered to achieve similar benefits; however, if no dithering were used, destructive combining would be observed at the receiver in some instances.

Dithering on a symbol-by-symbol basis complicates the channel estimation process at the receiver. As described in [12], each code block can be split into a number of bursts with the phase dithered only on a burst-by-burst basis. If all transmitting nodes use the same training pattern, then the receiver can estimate the composite channel on each burst based on this known sequence. The receiver thus does not know how many nodes are teaming to send and processes the received signal in the same manner as it would in the non-cooperitive case. Furthermore, each transmitter does not have any channel state information for its channel, nor does it even require knowledge of the existence of other teaming transmitters. Finally, a node may cooperate in more than one team to relay to multiple receiving nodes and the teams may change on a packet by packet basis according to connectivity.

Autonomous cooperation is a key component of BRNs as it enables highly reliable relay of packets. Indeed, it is well-understood that the use of cooperative communication schemes in ad hoc networks can lead to significant increases in network capacity (cf., [17] and the references therein). Two major approaches to cooperative communication have thus far been studied: distributed space-time coding (D-STC) (cf., [18]) and distributed beamforming (D-BF) (cf., [19]). Both approaches, however, present significant implementation challenges due to the required inter-node coordination. For example, both D-STC and D-BF require that the number of cooperating transmitters be known *a priori*. D-STC requires that cooperating users coordinate their encoding and transmission sequence according to some space-time code. D-BF requires that the composite transmission channel be known by all cooperating transmitters as well as the receiver.

The coordination required for D-STC and D-BF is particularly problematic in the context of *tactical* MANETs because it reduces robustness while increasing latency and overhead. In networks with highly dynamic topologies, frequent updates of the inter-node coordination information would be required, thereby impacting overhead and increasing the latency required for packet relay. Also, it is unclear how to efficiently design coordinated cooperation schemes that allow for multiple cooperating teams with some shared members. The team transmission could be serialized, but this again induces more delay. Finally, it is likely that inaccurate coordination information would be used on some occasions, resulting in degraded reliability. The autonomous cooperative scheme described above provides the majority of the performance benefits of these more complicated approaches while eliminating the hurdles associated with practical tactical MANET implementation.

There are several other assumptions regarding autonomous cooperation in BRNs that are more subtle. First, there are *hop-based protocol header constraints* so as to ensure that each node transmits exactly the same data packet. Specifically, networking stacks typically modify protocol header information at each relaying node so that even if two relay nodes transmit identical payload data, the resulting on-air packets may be different. In order to support autonomous cooperation, node-specific packet transformation must be suppressed: protocol headers can be modified only in a manner that is hop-dependent. Second, in addition to relaying identical packets, cooperating nodes must also relay on identical time slots. Relaying decisions are traditionally made at the network layer, introducing processing delays that are unpredictable and node-dependent (cf., [20]). To minimize delay and ensure that relays occur on the same slot, relay decisions can be made at the physical layer, which we term *PHY-layer switching*. Finally, cooperating nodes transmitting on the same TDMA slot may incur different propagation delay due to the heterogeneous propagation environment. This effect induces a *cooperative delay spread* and the receiver must be capable of reliable reception in this environment. In most tactical MANET scenarios, the point-to-point link is expected to be frequency selective so the receiver has the capability to handle delay spread. The cooperative delay spread is additive to the physical delay spread of the channels. Thus, the receiver should accommodate the total effective delay spread induced by the channel and the expected cooperative delay spread. Autonomous cooperation is therefore compatible with, for example, narrowband single carrier waveforms [13], orthogonal frequency division multiplexing (OFDM) waveforms, and direct sequence waveforms (cf., [21]), with sufficient equalizer capability, cyclic prefix length, and Rake receiver finger span, respectively.

C. Barrage Access Control

The barrage broadcast mechanism described in Section III-A provides a low-latency, robust mechanism for flooding from a *single* source node. In order to provide networking with multiple sources, some form of access control is required. The simplest scenario is PTT voice in which an operator accesses the channel by keying down the transmitter and collisions are resolved by standard manual operational protocols. For data transmission, this operator control is not sufficient. Nonetheless, it illustrates the fundamental PHY-layer resource to be controlled: the entire barrage broadcast fabric for some period of time. This is in stark contrast to the standard layered architecture for network design which controls access to a node-to-node link via a collision avoidance MAC protocol.

While MAC may still be an applicable term to describe algorithms controlling access to barrage resource, the term *Barrage Access Control (BAC)* was introduced in [14] to emphasize that BRNs are not based on a link abstraction. As an example of the basic structure for a BAC protocol, assume that the TDMA structure employed by the BRN partitions time slots into three independent logical channels:

(i) the *request* logical channel (RLC);
(ii) the *confirmation* logical channel (CLC); and,
(iii) the *data* logical channel (DLC).

The BAC protocol uses control traffic transmitted on the RLC and CLC in order to coordinate access to the DLC. The bandwidth reserved for the RLC and CLC therefore constitute the overhead associated with this protocol. In practice the overhead required by a specific BAC protocol depends heavily on the design constraints it must meet (e.g., how rapidly new sources must be able to access the DLC). Future access to the data logical channel is tabulated in a *DLC schedule* that lists which node (if any) is the designated source on each DLC slot. When a node is the (sole) designated source, it injects packets every M time slots; otherwise, nodes relay what they receive as per the protocol described in Section III-A. In this manner, multiple data flows are time-multiplexed on the DLC so as to provide shared access to the efficient barrage broadcast mechanism (the overhead associated with switching between sources on the DLC is discussed in Section IV below.)

The RLC and CLC are used to maintain the DLC schedule in response to time-varying traffic demands. Nodes announce their desire to source data (and how much) via a message broadcast on the RLC. A common *schedule control* node responds to this message by broadcasting a scheduling update message on the CLC that incorporates the new data source. Collisions are precluded on the CLC since there is a dedicated source for all DLC schedule updates. Collisions on the RLC, however, can result when multiple candidate sources transmit simultaneously. In this case, the receipt (or not) of a concomitant scheduling update can be used as an implicit positive (or negative) acknowledgement of the successful transmission of a request message to the schedule control node, thereby enabling the use of Ethernet-like backoff and retry schemes.

IV. Broadcast Scaling in BRNs

The scalability of the abstract protocol for broadcast coordination described above is apparent: sources time-multiplex access to a broadcast mechanism whose single-source throughput is independent of the network size. The scaling of broadcast latency in particular can be readily ascertained. Following Gupta and Kumar's standard formulation, consider a network comprising n nodes randomly distributed in an operational region with unit area. The point-to-point transmission range of the nodes scales with n as $\Theta(\sqrt{\log n/n})$ so as to guarantee connectivity with high probability [1]. The hop diameter $h(n)$ of the network scales as the ratio of the network Euclidean diameter to the radio range: $\Theta(\sqrt{n/\log n})$. Neglecting the range extension[4] that is afforded by cooperative communications, the total number of time slots required to broadcast a singe packet through the network is simply $h(n)$. That is to say, the latency of barrage broadcast scales at most as:

$$\Theta\left(\sqrt{\frac{n}{\log n}}\right) \tag{4}$$

In order to assess the scaling of broadcast throughput, it is first necessary to specify what is in fact being measured as n grows. In scaling analyses of traditional MANET architectures, the growth in overhead associated with maintaining optimal communication routes as the network size increases is typically neglected [9]. In keeping with this convention, the bandwidth allocated to the request and confirmation logical channels is neglected in the following. However, the overhead associated with switching between sources can be accounted for analytically and sheds light on an important practical

[4]When cooperative range extension is considered, $h(n)$ is replaced by the minimum number of slots required to fully broadcast from a given source, maximized over all sources. A recent cooperative flooding study demonstrated that this factor grows logarithmically with $h(n)$ under certain conditions [22].

design consideration for BAC protocols. Specifically, there is a certain amount of time required to allow the last broadcast packet of one source to ripple out of the network before another node can start sourcing packets. Returning to Figure 1, the black node can clearly transmit a new packet on slot A of the second frame. However, if control of the barrage broadcast mechanism switched immediately from the black node to the one of the one-hop nodes on frame 2, then a packet injected by that node on slot A of frame 2 would collide with the packet sourced by the black node in frame 1.

Returning to the same framework used to study broadcast latency above, suppose that source nodes are allocated sufficient slots L to transmit P packets. Again neglecting cooperative range extension effects, L must be at least

$$L(n, P) = M(P - 1) + h(n) \tag{5}$$

in order to prevent collisions on the DLC with M-slot TDMA framing. The sum broadcast throughput is therefore:

$$\frac{P}{L(n, P)} = \frac{P}{M(P - 1) + h(n)}. \tag{6}$$

Provided the number of packets P per contiguous source slot block is scaled appropriately with increasing network size n, $\Theta(n)$ receiver-oriented sum broadcast throughput scaling is therefore achieved for an arbitrary number of source nodes. In summary, BRNs scale optimally for two of the most critical metrics for tactical MANET: latency and broadcast throughput.

V. UNICAST FLOWS IN BRNS

As discussed in Section II, while tactical MANET traffic is largely broadcast or multicast, support for unicast is a reasonably expected requirement. Unicast, or multicast for that matter, can of course be achieved via network-wide flooding (i.e., all nodes relay but only intended recipients pass the received packet up the stack from the PHY-layer); however, the ability to support multiple, contemporaneous localized unicast streams can only serve to enhance the network capacity offered by barrage relay networks. This section describes controlled barrage regions (CBRs), which are a building block for just such spatially separated unicast transmissions.

A. Cooperative Links, Paths, and CBRs

In a traditional network, a link is defined by a transmit/receive radio pair that share a suitably reliable point-to-point communications channel. A path is a series of these links or hops connecting a source and destination. A unicast route is typically a single path, although it is possible to establish several paths to form a route with robustness through redundancy – e.g., the backup routing variant of the ad hoc on-demand distance vector (AODV) routing protocol [23].

In a BRN, however, links are *cooperative* and comprise one or more transmitters and one receiver as illustrated in Figure 2(a). A cooperative path, as illustrated in Figure 2(b) is a formed by series of braided cooperative links between a source and destination node. A controlled broadcast region is simply one or more cooperative paths (of distinct lengths) from a source node to a destination node. Figure 2(c) illustrates a CBR comprising two cooperative paths where the longer cooperative path goes around a large obstruction.

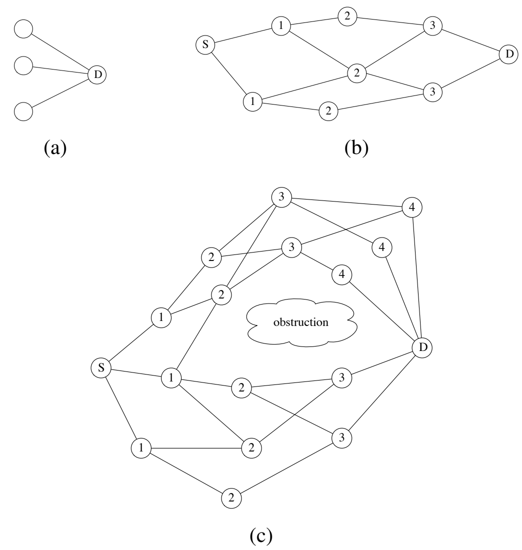

Fig. 2. Examples of (a) a (3 to 1) cooperative link; (b) a (length 4) cooperative path; and (c) a controlled barrage region comprising two cooperative paths, one of length 5 hops and another of length 4.

These definitions directly illustrate the robustness of unicast transport via CBRs. First, a cooperative link may remain viable even if one or more transmitting nodes becomes nonfunctional. Second, a cooperative path may remain reliable even if some of its cooperative links are lost. Finally, a CBR may remain viable even if some of the cooperative paths are lost. This makes the likelihood of loosing source-destination connectivity low even in the presence of significant variation of link reliability and/or network topology. It may be possible, though unlikely, that a CBR could be lost due to single node failure. For example, if the middle node labeled “1” in Figure 2(c) was lost, it is conceivable that unicast transmission along the CBR would fail. While this description of robustness is qualitative, these benefits are shown quantitatively in [14] and [24].

B. Establishing CBRs for Unicast & Multicast

Controlled barrage regions can be established by specifying a set of buffer (i.e., sentry) nodes around the cooperating nodes. The relay function of buffer nodes is suppressed so that external packets do not propagate into the control region, nor do internal packets propagate to the rest of the network. In this way, multiple unicast transmissions may be established in different portions of the network as illustrated in Figure 3. It is precisely the action of these buffer nodes that inspired the name *controlled* barrage region. This concept is readily extended to multicast by establishing sentries around a portion of the network containing the source and all destination nodes.

The buffer nodes can be specified through the broadcast of request-to-send (RTS) / clear-to-send (CTS) messages on the BAC control channels. For example, if the source broadcasts

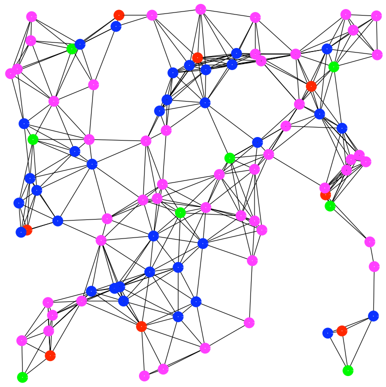

Fig. 3. Eight unicast flows in a randomly-generated barrage relay network topology. Source, destination, relay, and buffer nodes are colored green, red, blue, and pink, respectively.

an RTS message and each relaying node increments a hop-count field, the destination node will obtain knowledge of the shortest cooperative path from the source. If the destination broadcasts a CTS message containing the shortest path length and each relaying node increments a hop-count field from the destination, then the source will also know the distance from the destination. Furthermore, each relay node will know the shortest cooperative path length (S) and the length of the cooperative path on which it lies. A CBR containing only the shortest cooperative path can be established by making all nodes on cooperative paths one larger than this value ($S+1$) buffer nodes. A CBR comprising cooperative paths of length S and $S+1$, can be established by making nodes on paths of length $S+2$ buffer nodes. This method can be used to include cooperative paths of length $S+M-1$, where S is the shortest path and M is the number of TDMA slots per frame (e.g., 3 in Figure 1). While only described at a high level here, the correctness of a formal specification of this CBR establishment protocol can be proven for unicast and multicast [25].

C. Coordinating CBRs

The abstract BAC protocol described in Section III-C must clearly be extended in order to support unicast and multicast transmission via CBRs. In particular, the *global* DLC schedule described for broadcast traffic can instead be replaced by *local*, node-specific schedules. While these schedules will not be identical at all nodes – e.g., two nodes may be unicast sources on the same time slot – they must be made consistent so as to avoid collisions between distinct flows.

VI. Related Work

Independent of TrellisWare's BRN development, researchers have suggested many approaches analogous to those described herein. Much of this work has focused on component technology (e.g., cooperation) without consideration of how various methods could be integrated into a full practical MANET design. There are a few notable exceptions described below.

In [20], Ramanathan articulated a design concept for next-generation MANETs that shares a number of common features with BRNs including cooperative transmission and a multi-packet, multi-hop medium reservation scheme. In particular, the *hop-centric* nature of data transport in current-generation MANET systems – i.e., each packet experiences a large amount of processing, queueing, and access contention delay at each hop – was identified in [20] as introducing sufficient end-to-end latency so as to make wireless networks latency-limited under certain conditions. A *cut-through physical layer* that is similar to the PHY-layer switching assumed in BRNs was consequently described in [20, 26]. While this paper succeeded in providing a set of concepts key to next generation MANET design, the compatibility of the components was not demonstrated (e.g., while the importance of cooperation is described, the multi-hop, path-oriented MAC is based on traditional links rather than an analog of CBRs). Nevertheless, Ramanathan demonstrated extraordinary vision in his radical redefinition of wireless network design; that these concepts were independently formulated, implemented, and validated in a *tactical* MANET system by TrellisWare lends credence to their potential relevance to a wider application space.

Opportunistic large arrays (OLAs) were introduced in [27] to describe an efficient physical layer flooding algorithm for asynchronous sensor networks in which nodes accumulate power from many received copies of the same symbol. Subsequently, a number of protocol design issues for OLAs have been investigated (cf., [28] and the references therein). The key differences between BRNs and OLAs lie in the packet size, data rates, and assumed RF environments of the target applications. Specifically, the time synchronization and receiver signal processing assumed in BRNs allows packets to be many bits long (rather than single symbols) and data rates to be larger than the inverse of the maximum relative delay spread (i.e., multipath effects can be mitigated).

Recently, Yackoski, *et al.* studied the use of cooperative communications in practical wireless networks with a particular focus on the control cost associated with cooperation [29]. Following a multi-hop RTS/CTS procedure that uses non-cooperative transmission over a reactively maintained route to establish which nodes will cooperatively relay, a mechanism remarkably similar to flooding in a CBR – including spatial pipelining – is used for data transport. Although the decentralized D-STC scheme used for node cooperation in [29] is not autonomous in the sense described above, it indeed addresses many of the overhead issues associated with cooperative communications in mobile ad hoc networks.

Finally, a number of authors have studied cooperative communications in the context of minimizing the sum energy (cf., [30, 31]) or latency (cf., [22, 31]) required for network-wide broadcast. Mergen, *et al.* analyzed critical parameters that govern the performance of cooperative flooding under a continuum approximation of dense networks [32]. With the notable exception of [33], however, the majority of cooperative flooding studies have focused on the single source case – i.e., analogs of barrage access control for more general cooperative transmission schemes have not been widely examined.

VII. Conclusion

Unlike traditional approaches to wireless networking – wherein networking protocols are based on a point-to-point link abstraction – barrage relay networks utilize an autonomous cooperative communications scheme that affords a radically different building block for protocol design: robust, low-latency broadcast. In this paper, it was demonstrated that by time-multiplexing this flooding mechanism between competing data sources, $\Theta(1)$ broadcast scaling of the transmit data rate can be achieved, implying linear scaling in the useful received data rate with network size. Extensions of the barrage flooding mechanism to support traffic that can be spatially contained (e.g., unicast) were also discussed.

The BRN technology that underlies TrellisWare's tactical MANET system can be seen as a pragmatic fusion of a number of design elements: autonomous cooperation, PHY-layer switching, barrage access control, and controlled barrage regions. While developed internally at TrellisWare, many similar concepts have appeared in the literature in isolation (e.g., [20, 29]). Some of the design trades made in this synthesis reflect a focus on comprehensive implementation and deployment rather than strict theoretical optimality.

A formal information theoretic treatment of the BRN concept may illuminate methods to characterize and enhance the scalability, performance, and extensibility of networking stacks which employ it. In particular, it remains to be seen to what extent the barrage flooding mechanism can be employed in network designs with different optimization criteria than latency and robustness (e.g., delay tolerant unicast traffic, energy conservation.) Furthermore, the application of advanced techniques for traditional MANET performance enhancement (e.g., power control, rate adaptation, hierarchical routing) to the BRN framework is an interesting area for future research.

References

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.

[2] F. Xue and P. R. Kumar, "Scaling laws for ad hoc wireless networks: An information theoretic approach," *Foundations and Trends in Networking*, vol. 1, no. 2, pp. 145–270, 2006.

[3] Z. Wang, H. R. Sadjadpour, and J. J. Garica-Lua-Aceves, "A unifying perspective on the capacity of wireless ad hoc networks," in *Proc. IEEE Conf. Computer Comms.*, Phoenix, AZ, April 2008, pp. 753–761.

[4] B. Tavli, "Broadcast capacity of wireless networks," *IEEE Communications Letters*, vol. 10, no. 2, pp. 68–69, 2006.

[5] X. Y. Li, S. J. Tang, and O. Frieder, "Multicast capacity for large scale wireless ad hoc networks," in *Proc. ACM/IEEE Int'l Conf. Mobile Comp. and Networking*, Montral, Canada, September 2007.

[6] A. Ozgur, O. Leveque, and D. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. Information Theory*, vol. 53, no. 10, pp. 2549–2572, 2007.

[7] A. Keshavarz-Haddad, V. Ribeiro, and R. Riedi, "Broadcast capacity in multihop wireless networks," in *Proc. ACM/IEEE Int'l Conf. Mobile Comp. and Networking*, Los Angeles, CA, 2006.

[8] H. Liu, P. Wan, X. Jia, X. Liu, and F. Yao, "Efficient flooding scheme based on 1-hop information in mobile ad hoc networks," in *Proc. IEEE Conf. Computer Comms.*, Barcelon, Spain, 2006.

[9] J. Andrews, N. Jindal, M. Haenggi, R. Berry, S. Jafar, D. Guo, S. Shakkottai, R. Heath, M. Neely, S. Weber, and A. Yener, "Rethinking information theory for mobile ad hoc networks," *IEEE Comms. Mag.*, vol. 46, no. 12, pp. 94–101, December 2008.

[10] J. K. Burbank, P. F. Chimento, B. K. Haberman, and W. T. Kasch, "Key challenges of military tactical networking and the elusive promise of MANET technology," *IEEE Comms. Mag.*, pp. 39–45, November 2006.

[11] B. Williams and T. Camp, "Comparison of broadcasting techniques for mobile ad hoc networks," in *Proc. ACM Int'l Symp. Mobile Ad Hoc Networking and Computing*, Lausanne, Switzerland, 2002, pp. 194–205.

[12] D. K. Lee and K. M. Chugg, "Pragmatic cooperative diversity communications," in *Proc. IEEE Military Comm. Conf.*, Washington, DC, October 2006.

[13] A. Blair, T. Brown, K. M. Chugg, and M. J. Johnson, "Tactical mobile mesh network system design," in *Proc. IEEE Military Comm. Conf.*, Orlando, FL, October 2007.

[14] A. Blair, T. Brown, K. M. Chugg, T. R. Halford, and M. Johnson, "Barrage relay networks for cooperative routing in tactical MANETs," in *Proc. IEEE Military Comm. Conf.*, San Diego, CA, November 2008.

[15] I. Hammerström, M. Kuhn, and A. Wittneben, "Cooperative diversity by relay phase rotations in block fading environments," in *Proc. IEEE Workshop on Signal Proc. Advances in Wireless Comm.*, Lisbon, Portugal, July 2004, pp. 293–297.

[16] K. M. Chugg, P. Thiennviboon, G. D. Dimou, P. Gray, and J. Melzer, "A new class of turbo-like codes with universally good performance and high-speed decoding," in *Proc. IEEE Military Comm. Conf.*, Atlantic City, NJ, October 2005.

[17] A. Scaglione, D. L. Goeckel, and J. N. Laneman, "Cooperative communications in mobile ad hoc networks," *IEEE Signal Processing Mag.*, vol. 23, pp. 18–29, September 2006.

[18] J. N. Laneman and G. W. Wornell, "Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks," *IEEE Trans. Information Theory*, vol. 49, pp. 2415–2425, October 2003.

[19] E. Erkip, A. Sedonaris, A. Stefanov, and B. Aazhang, *Advances in Network Information Theory*. AMS DIMACS Series, 2004, ch. Cooperative communication in wireless systems, pp. 303–320, eds. P. Gupta, G. Kramer, and A. J. van Wijngaarden.

[20] R. Ramanathan, "Challenges: A radically new architecture for next generation mobile ad hoc networks," in *Proc. ACM/IEEE Int'l Conf. Mobile Comp. and Networking*, Cologne, Germany, August 2005, pp. 132–139.

[21] N. H. Lu and C. Nelson, "Cooperative team communications," in *Proc. IEEE Military Comm. Conf.*, Boston, MA, October 2009.

[22] M. Baghie and B. Krishnamachari, "Fast flooding using cooperative transmissions in wireless networks," in *Proc. International Conf. Comms.*, Dresden, Germany, June 2009.

[23] S. J. Lee and M. Gerla, "AODV-BR: Backup routing in ad hoc networks," in *Wireless Comm. and Networking Conf.*, Chicago, IL, September 2000, pp. 1311–1316.

[24] T. R. Halford and K. M. Chugg, "The reliability and stability of unicast transmission in barrage relay networks," submitted to MobiCom 2010.

[25] T. R. Halford, K. M. Chugg, and A. Polydoros, "Barrage relay networks: System and protocol design," to appear at PIMRC 2010 (invited).

[26] R. Ramanathan and F. Tchakountio, "Channel access over path segments for ultra low latency MANETs," in *Proc. IEEE Military Comm. Conf.*, Orlando, FL, October 2007.

[27] A. Scaglione and Y.-W. Hong, "Opportunistic large arrays: Cooperative transmission in wireless multihop ad hoc networks to reach far distances," *IEEE Trans. Signal Proc.*, vol. 51, no. 8, pp. 2082–2092, 2003.

[28] L. Thanayankizil, A. Kailas, and M. A. Ingram, "Routing for wireless sensor networks with an opportunistic large array physical layer," *Ad Hoc & Sensor Wireless Networks*, vol. 8, no. 1-2, pp. 79–117, 2009.

[29] J. Yackoski, L. Zhang, C.-C. Shen, L. Cimini, and B. Gui, "Networking with cooperative communications: Holistic design and realistic evaluation," *IEEE Comms. Mag.*, pp. 113–119, August 2009.

[30] I. Maric and R. D. Yates, "Cooperative multihop broadcast for wireless networks," *IEEE J. Select. Areas Commun.*, vol. 22, no. 6, pp. 1080–1088, August 2004.

[31] G. Jakllari, S. V. Krishnamurthy, M. Faloutsos, and P. V. Krishnamurthy, "On broadcasting with cooperative diversity in multi-hop wireless networks," *IEEE J. Select. Areas Commun.*, vol. 25, no. 2, pp. 484–496, February 2007.

[32] B. Sirkeci-Mergen, A. Scaglione, and G. Mergen, "Asymptotic analysis of multistage cooperative broadcast in wireless networks," *IEEE Trans. Information Theory*, vol. 52, no. 6, pp. 2531–2550, June 2006.

[33] S. Kirti and A. Scaglione, "Cooperative broadcast in dense networks with multiple sources," in *Proc. IEEE Intl. Workshop on Signal Proc. Advances for Wireless Comms.*, Perugia, Italy, June 2009, pp. 514–518.