The Approximate Capacity of the Gaussian N-Relay Diamond Network

Urs Niesen
Bell Laboratories, Alcatel-Lucent
Murray Hill, NJ 07974
urs.niesen@alcatel-lucent.com

Suhas Diggavi
University of California at Los Angeles
Los Angeles, CA 90095
suhas@ee.ucla.edu

***Abstract*—We consider the Gaussian "diamond" or parallel relay network, in which a source node transmits a message to a destination node with the help of N relays. Even for the symmetric setting, in which the channel gains to the relays are identical and the channel gains from the relays are identical, the capacity of this channel is unknown in general. The best known capacity approximation is up to an additive gap of order N bits and up to a multiplicative gap of order N^2, with both gaps independent of the channel gains.**

In this paper, we approximate the capacity of the symmetric Gaussian N-relay diamond network up to an additive gap of 1.8 bits and up to a multiplicative gap of a factor 14. Both gaps are independent of the channel gains, and, unlike the best previously known result, are also independent of the number of relays N in the network. Achievability is based on bursty amplify-and-forward, showing that this simple scheme is uniformly approximately optimal, both in the low-rate as well as high-rate regimes. The upper bound on capacity is based on a careful evaluation of the cut-set bound.

I. INTRODUCTION

Cooperation is a key feature of wireless communication. A simple canonical channel model capturing this feature is the "diamond" or parallel relay network, introduced by Schein and Gallager [1], [2]. This network consists of a source node connected through a broadcast channel to N relays; the relays, in turn, are connected to the destination node through a multiple-access channel (see Fig. 1). The objective is to maximize the rate achievable between the source and the destination with the help of the N relays. Throughout this paper, we will be interested in the Gaussian version of this problem, in which both the broadcast and the multiple-access part are subject to additive Gaussian noise. Moreover, for simplicity we will restrict attention to the symmetric case, in which the channel gains within the multiple-access part and the broadcast part of the network are identical (but are allowed to differ between the multiple-access and broadcast parts).

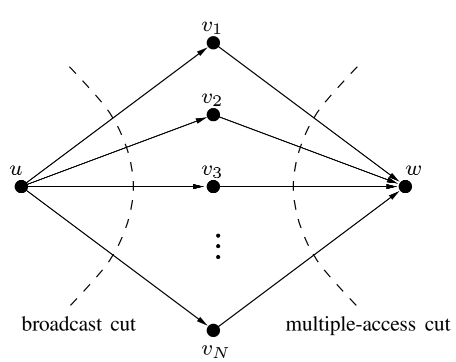

Fig. 1. The N-relay diamond network. The source node u transmits a message to the destination node w via the N relays $\{v_n\}_{n=1}^N$. The two cuts indicated in the figure are the broadcast cut (separating the source u from the relays $\{v_n\}$) and the multiple-access cut (separating the relays $\{v_n\}$ from the destination w).

For the Gaussian 2-relay diamond network, the rates achievable with decode-and-forward and with amplify-and-forward at the relays were analyzed in [2]. It is shown that these schemes achieve capacity in some regimes of signal-to-noise ratios (SNRs) of the broadcast and multiple-access parts of the diamond network. The asymptotic behavior of the N-relay Gaussian diamond network was investigated in [3]. In certain regimes of SNRs of the broadcast and multiple-access parts of the network, it is shown that amplify-and-forward is capacity achieving in the limit as $N \to \infty$. New achievable schemes for the Gaussian diamond network with bandwidth mismatch (i.e., the source and the relays have different bandwidth) were introduced in [4] and [5]. Perhaps surprisingly, these schemes lead to higher achievable rates than the ones obtained with amplify-and-forward and decode-and-forward even when the bandwidths at the source and the relays are identical. Half-duplex versions of the Gaussian diamond network, in which the relays cannot receive and transmit signals simultaneously, were considered in [6] and [7]. The capacity of a special class of 2-relay diamond networks is derived in [8]. For networks in this class, one relay receives the signal sent at the source without noise, and the destination node is connected to the relays by two orthogonal bit pipes of fixed rate. To the best of our knowledge, this is the only non-trivial example for which the capacity of the diamond network is known for all values of SNR. For the general Gaussian N-relay diamond network, the capacity is unknown.

Given the difficulty of determining the capacity of communication networks in general and of the diamond network in particular, it is natural to ask if it can at least be approximated. For high rates, such an approximation should be additive in nature, i.e., we would like to determine capacity up to an additive gap. For low rates, such an approximation should be multiplicative, i.e., we would like to determine capacity up

978-1-4577-0595-3/11/$26.00 ©2011 IEEE

to a multiplicative gap. If a communication strategy can be shown to have both small additive as well as multiplicative gap, then that strategy is provably close to optimal both in the high rate as well as low rate regimes.

Additive approximations for channel capacity of communication networks were first derived in [9], where the capacity region of the two-user Gaussian interference channel is determined up to an additive gap of one bit. This approach of approximate capacity characterization was applied to general relay networks with single-source multicast in [10]. For such networks, the capacity is derived up to an additive gap of $13n$ bits, where n is the number of nodes in the network. This additive gap can be further sharpened to $3n$ [11] for the complex Gaussian case (or $1.5n$ for the real case). Since the N-user diamond network is a special case of a general relay network with a single source and destination and with $n = N+2$ nodes, these results yield an additive approximation up to a gap of $1.5N + 3$ bits for this network (assuming real channel gains).

Multiplicative approximations, pioneered in [12], were mostly analyzed for large wireless networks. For a network with n nodes, the emphasis was on finding capacity approximations up to a small multiplicative factor in n. Under a restricted model of communication, the capacity of a random wireless network was determined up to a constant multiplicative factor independent of n. Approximations for the equal rate point under a general Gaussian model were derived in [13] up to a multiplicative factor of $O(n^{\varepsilon})$ for any $\varepsilon > 0$. These approximation results were subsequently sharpened in [14], [15] to a factor $n^{O(1/\sqrt{\log(n)})}$. Approximations for the entire capacity region (as opposed to just the equal-rate point) were derived in [16] up to the same multiplicative factor of $n^{O(1/\sqrt{\log(n)})}$. Under some conditions on the node placement, this factor can further be sharpened to $O(\log(n))$ [17]. Multiplicative approximations for arbitrary relay networks with single-source multicast (as opposed to wireless networks with multiple unicast, i.e., multiple separate source-destination pairs) were derived in [10]. For a network with maximum degree d, the capacity is approximated to within a factor of $2d(d + 1)$. As pointed out earlier, the Gaussian N-relay diamond network is such a network with maximum degree $d = N$, and hence this result yields a multiplicative approximation up to a factor of $2N(N + 1)$.

To summarize, the capacity region of the general Gaussian N-relay diamond network is not known. The best known additive approximation is up to a gap of $1.5N + 3$ bits, and the best known multiplicative approximation is up to a factor of $2N(N + 1)$. In either case, the bounds degrade rather quickly as N increases. It is hence of interest to find approximation guarantees that behave better as a function of the number of relays N in the network. Ideally, we would like the approximation guarantees to be uniform in N.

As a main result of this paper, we show that such a uniform approximation is indeed possible. More precisely, we find an additive approximation of the capacity of the symmetric Gaussian N-relay diamond network of gap at most 1.8 bits for any SNR and number of relays N. Moreover, we find a multiplicative approximation to the capacity up to at most a factor 14, again for any SNR and number of relays N. This is a significant improvement over the previously best known additive approximation of $1.5N + 3$ bits and multiplicative approximation of a factor $2N(N + 1)$, especially for large values of N. In particular, as far as we know, this is the first such approximation result (in both multiplicative and additive approximation) that is independent of the number of network nodes for a nontrivial class of wireless networks. We further show that bursty amplify-and-forward with carefully chosen duty cycle is close to capacity achieving for the diamond network simultaneously in the sense of multiplicative and additive approximation up to the aforementioned gaps. Hence, bursty amplify-and-forward (with appropriately chosen duty cycle) is a good communication scheme for the Gaussian N-relay diamond network both at low and at high SNRs, and independently of the number of relays N.

The main technical contribution of this paper is the upper bound on capacity. The standard way to obtain upper bounds on the capacity of the diamond network is to evaluate two particular cuts in the wireless network, namely the one separating the source from the relays (called the *broadcast cut* in the following) and the one separating the relays from the destination (called the *multiple-access cut* in the following) as depicted in Fig. 1. This approach is taken, for example, in [3]–[5]. In fact, for the type of symmetric Gaussian N-relay diamond networks considered in this paper, whenever the capacity is known, it coincides with the minimum of these two cuts. We show in this paper that, in order to obtain uniform additive or multiplicative approximations for the capacity of this network, considering just these two cuts is not sufficient. Instead we need to *simultaneously* optimize over *all* possible 2^N cuts separating the source from the destination have to be taken into account. Without this careful outer bound evaluation, we believe that the uniform (in network size) approximation would not have been possible.

The remainder of this paper is organized as follows. Section II introduces the precise problem statement. Section III presents the main results. Section IV contains concluding remarks. Due to space constraints, all results are presented without proofs. These can be found in the full version [18] of the paper.

II. Problem Statement

Consider the symmetric Gaussian N-relay diamond network as depicted in Fig. 1. The source node u transmits a message to the destination node w with the help of N parallel relays $\{v_1, \ldots, v_N\}$. The inputs at time $t \in \mathbb{N}$ at nodes u and v_n will be denoted by $X[t]$ and $X_n[t]$, respectively. The channel outputs at time $t \in \mathbb{N}$ at nodes w and v_n will be denoted by $Y[t]$ and $Y_n[t]$. The channel inputs and outputs are related as

$$Y_n[t] \triangleq \sqrt{g}X[t] + Z_n[t],$$

$$Y[t] \triangleq \sqrt{h}\sum_{n=1}^{N} X_n[t] + Z[t],$$

where $\{Z[t]\}_t, \{Z_n[t]\}_{n,t}$ are independent and identically distributed Gaussian random variables with mean zero and variance one, independent of the channel inputs. The channel gains g and h are assumed to be real positive numbers, constant as a function of time, and known throughout the network.

A *T-length block code* for the diamond network is a collection of functions

$$\begin{aligned} f &: \{1,\dots,M\} \to \mathbb{R}^T, \\ f_n &: \mathbb{R}^T \to \mathbb{R}^T, \forall n \in \{1,\dots,N\}, \\ \phi &: \mathbb{R}^T \to \{1,\dots,M\}. \end{aligned}$$

The encoding function f maps the message W, assumed to be uniformly distributed over the set $\{1,\dots,M\}$, to a channel input

$$\{X[t]\}_{t=1}^T \triangleq f(W)$$

at the source node u. The function f_n maps the channel outputs $\{Y_n[t]\}_{t=1}^T$ to the channel inputs

$$\{X_n[t]\}_{t=1}^T \triangleq f_n\big(\{Y_n[t]\}_{t=1}^T\big)$$

at relay v_n.[1] The decoding function ϕ maps the channel outputs $\{Y[t]\}_{t=1}^T$ at the destination node w into a reconstruction

$$\hat{W} \triangleq \phi\big(\{Y[t]\}_{t=1}^T\big).$$

We say the code satisfies a *unit average power constraint* if

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}(X^2[t]) \le 1, \quad \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}(X_n^2[t]) \le 1, \forall n.$$

The *rate* of the code is $\log(M)/T$, and its *average probability of error* $\mathbb{P}(\hat{W} \neq W)$. A rate R is *achievable* if there exists a sequence of T-length block codes with unit average power constraint and rate at least R such that the average probability of error approaches zero as $T \to \infty$. The *capacity* $C(N,g,h)$ is the supremum of all achievable rates.

Throughout this paper, we use bold font to denote vectors and matrices. $\log$ and $\ln$ denote the logarithms to base 2 and e, respectively. All capacities and rates are given in bits per channel use.

III. Main Results

A natural scheme for the diamond network is *amplify-and-forward*, in which each relay transmits a scaled version of the received signal. Formally,

$$X_n[t] = \alpha Y_n[t] = \alpha\sqrt{g}X[t] + \alpha Z_n[t],$$

where the constant α is chosen to satisfy the power constraint at the relay. Denote by $R_1(N,g,h)$ the rate achieved by amplify-and-forward.

[1]Note that the functions $\{f_n\}$ at the relays are not causal. This is to simplify notation; due to the layered nature of the network all results remain the same if causality is imposed.

If the SNR at the relays is low (i.e., $g \ll 1$), it turns out that simple amplify-and-forward is arbitrarily suboptimal. This is because the received signal power g at the relay is much smaller than the noise power 1, and therefore the relay amplifies mostly noise. This effect can be mitigated by using *bursty amplify-and-forward*. For a constant $\delta \in (0,1]$, called the *duty cycle* in the following, we communicate for a fraction δ of time at average power $1/\delta$ using the amplify-and-forward scheme and stay silent for the remaining time. This satisfies the overall average unit power constraint. The resulting achievable rate is denoted by $R_\delta(N,g,h)$. This notation is consistent, i.e., for $\delta = 1$ the simple and bursty amplify-and-forward scheme coincide and achieve both rate $R_1(N,g,h)$.

Our first result lower bounds the rate achievable by using either simple or bursty amplify-and-forward with optimized duty cycle δ.

Theorem 1. *For any $N \ge 2$, $g,h > 0$, there exists $\delta^\star \in (0,1]$ such that Equation* (4) *(on the top of the next page) holds.*

Note that the optimal duty cycle $\delta^\star$ is allowed to depend on N, g, and h. In the high-rate regime, i.e., the first and third case in (4), the duty cycle achieving the lower bound is $\delta^\star = 1$, and hence the bursty amplify-and-forward scheme reduces to simple amplify-and-forward. On the other hand, in the low-rate regime, i.e., the second, fourth, and fifth case in (4), $\delta^\star < 1$ and (genuine) bursty amplify-and-forward is used.

Having established an achievable rate, the next theorem provides an upper bound on the capacity of the diamond network.

Theorem 2. *For any $N \ge 2$, $g,h > 0$, Equation* (5) *(on the top of the next page) holds.*

As a corollary to Theorems 1 and 2, we obtain that bursty amplify-and-forward is close to optimal, in the sense that it achieves capacity both up to a constant additive as well as multiplicative gap, where the constants are independent of the number of relays N and the channel gains g and h. This shows that optimized bursty amplify-and-forward is a good communication scheme for the diamond network both at low rates (due to the small multiplicative gap to capacity) as well as high rates (due to the small additive gap to capacity).

Corollary 3. *For any $N \ge 2$, $g,h > 0$, there exists $\delta^\star \in (0,1]$ such that*

$$C(N,g,h) - R_{\delta^\star}(N,g,h) \le 1 + \frac{1}{2}\log(3) \le 1.8 \textit{ bits},$$

and

$$\frac{C(N,g,h)}{R_{\delta^\star}(N,g,h)} \le \frac{4}{\ln(4/3)} \le 14.$$

We point out that choosing the duty cycle $\delta^\star$ as a function of N, g, and h, is not necessary to obtain the additive approximation result in Corollary 3. In fact, using only simple amplify-and-forward achieves the same additive approximation guarantee, i.e.,

$$C(N,g,h) - R_1(N,g,h) \le 1.8 \text{ bits}$$

$$R_{\delta^\star}(N,g,h) \geq \begin{cases} \frac{1}{2}\log\big(1+\frac{1}{3}N\min\{g,Nh\}\big), & \text{if } \max\{g,Nh\}\geq 1\\ \frac{1}{2}\ln(4/3)\log(1+Ng), & \text{if } \max\{g,Nh\}<1,\ g\leq h\\ \frac{1}{2}\log\big(1+\frac{1}{3}N^2gh\big), & \text{if } \max\{g,Nh\}<1,\ g\in(h,N^2h),\ N\sqrt{gh}\geq 1\\ \frac{1}{2}\ln(4/3)\log(1+N\sqrt{gh}), & \text{if } \max\{g,Nh\}<1,\ g\in(h,N^2h),\ N\sqrt{gh}<1\\ \frac{1}{2}\ln(4/3)\log(1+N^2h), & \text{if } \max\{g,Nh\}<1,\ g\geq N^2h. \end{cases} \tag{4}$$

$$C(N,g,h) \leq \begin{cases} \frac{1}{2}\log\big(1+N\min\{g,Nh\}\big), & \text{if } \max\{g,Nh\}\geq 1\\ \frac{1}{2}\log(1+Ng), & \text{if } \max\{g,Nh\}<1,\ g\leq h\\ \frac{1}{2}\log\big(1+2N^2gh\big)+\frac{1}{2}, & \text{if } \max\{g,Nh\}<1,\ g\in(h,N^2h),\ N\sqrt{gh}\geq 1\\ \log(1+2N\sqrt{gh}), & \text{if } \max\{g,Nh\}<1,\ g\in(h,N^2h),\ N\sqrt{gh}<1\\ \frac{1}{2}\log(1+N^2h), & \text{if } \max\{g,Nh\}<1,\ g\geq N^2h. \end{cases} \tag{5}$$

for all $N\geq 2$, $g,h>0$. However, the same is not true if we are also interested in multiplicative approximation guarantees (at least in the low-rate regime). To achieve a constant additive as well as multiplicative approximation, the duty cycle $\delta^\star$ is required to vary as a function of N, g, and h.

From Theorems 1 and 2, the capacity of the diamond network has three distinct regimes, depending on whether $g\leq h$, $h<g<N^2h$, or $g\geq N^2h$. In the first regime ($g\leq h$), the channel gain to the relays is weak compared to the channel gain to the destination, and the achievable rate is constrained by the broadcast part of the diamond network. The capacity in this regime is given approximately by

$$C(N,g,h)\approx\frac{1}{2}\log(1+Ng),$$

where the approximation is in the sense of Corollary 3, namely up to a multiplicative gap of factor 14 in the low-rate regime ($g\ll N^{-1/2}$) and up to an additive gap of 1.8 bits in the high-rate regime ($g\gg N^{-1/2}$). This is precisely the capacity of a single-input multiple-output channel with one transmit antenna, N receive antennas, and channel gain $\sqrt{g}$ between each of them. Thus, the broadcast cut in Fig. 1 is approximately tight in this regime.

In the third regime ($g\geq N^2h$), the channel gain to the relays is strong compared to the channel gain to the destination, and the achievable rate is now constrained by the multiple-access part of the channel. The capacity in the third regime is given approximately by

$$C(N,g,h)\approx\frac{1}{2}\log(1+N^2h).$$

This is the capacity of the multiple-input single-output channel with N transmit antennas, one receive antenna, and channel gain $\sqrt{h}$ between each of them. Thus, the multiple-access cut in Fig. 1 is approximately tight in this regime. Observe that to achieve this rate, the signals sent by the relays must be highly correlated and add up coherently at the destination.

The most interesting regime is the second one ($h<g<N^2h$). If $\max\{g,Nh\}\geq 1$, then the capacity is given approximately by

$$C(N,g,h)\approx\frac{1}{2}\log\big(1+N\min\{g,Nh\}\big),$$

and again either the broadcast cut or the multiple-access cut are tight. If $\max\{g,Nh\}<1$ the situation is more complicated. If $N\sqrt{gh}\geq 1$, then the capacity of the diamond network is approximately

$$C(N,g,h)\approx\frac{1}{2}\log(1+N^2gh),$$

and, if $N\sqrt{gh}<1$,

$$C(N,g,h)\approx\frac{1}{2}\log(1+N\sqrt{gh}).$$

In both cases, the capacity depends on the product of g and h, and not merely on the minimum of g and Nh. Hence neither the broadcast cut nor the multiple-access cut are tight in this case. In fact, these bounds can be arbitrarily bad, both in terms of additive as well as multiplicative gap, as the next two examples illustrate.

For the additive gap, consider $g=N^{-5/8}$ and $h=N^{-9/8}$. Then $\max\{g,Nh\}=N^{-1/8}<1$, $g=N^{1/2}h\in(h,N^2h)$, and $N\sqrt{gh}=N^{1/8}\geq 1$, so that

$$C(N,g,h)\approx\frac{1}{2}\log(1+N^2gh)=\frac{1}{2}\log(1+N^{1/4}).$$

On the other hand, the minimum of the broadcast and multiple-access cuts yields

$$\frac{1}{2}\log\big(1+N\min\{g,Nh\}\big)=\frac{1}{2}\log(1+N^{3/8}),$$

resulting in an additive gap of order $\Theta(\log(N))$ bits, which is unbounded as the number of relays $N\to\infty$.

For the multiplicative gap, consider $g=N^{-2}$ and $h=N^{-3}$. Then $\max\{g,Nh\}=N^{-2}<1$, $g=Nh\in(h,N^2h)$, and $N\sqrt{gh}=N^{-3/2}<1$, so that

$$C(N,g,h)\approx\frac{1}{2}\log(1+N\sqrt{gh})\approx\frac{1}{2}\log(e)N^{-3/2}.$$

On the other hand, the minimum of the broadcast and multiple-access cuts yields

$$\frac{1}{2}\log\big(1+N\min\{g,Nh\}\big)\approx\frac{1}{2}\log(e)N^{-1}$$

resulting in a multiplicative gap of order $\Theta(\sqrt{N})$, which is again unbounded as the number of relays $N \to \infty$.

In the second regime, we thus need to take cuts other than the broadcast and multiple-access ones into account. The need for this can be understood as follows. Consider a general cut separating the source node u from the destination node w in the diamond network as shown in Fig. 2. Formally, let $S \subset \{1, \ldots, N\}$, and consider the cut from $u \cup \{v_n\}_{n \in S}$ to $w \cup \{v_n\}_{n \in S^c}$. Assume the signals $\{X_n\}_{n=1}^N$ sent from the relays to the destination are highly correlated. This results in the signal summing up coherently at the receiver, increasing the rate across the cut. At the same time, if the signals sent from the relays are highly correlated, then the signal $\{X_n\}_{n \in S^c}$ available at the relays on the other side of the cut can be used to estimate the signal received at the destination node. This decreases the rate across the cut. Thus, for general cuts, there is a tradeoff between the gain from coherent reception and the loss from prediction that come with increased signal correlation. This tradeoff is absent if we only consider the broadcast and multiple-access cuts. It is precisely this tradeoff that determines the behavior of the capacity of the diamond network in the second regime.

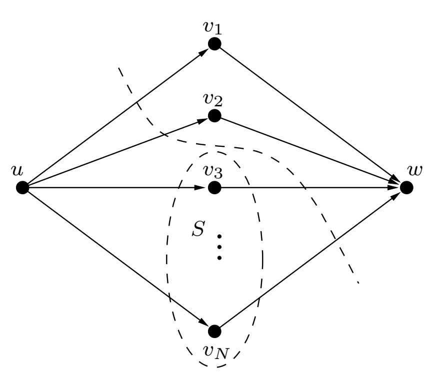

Fig. 2. A general cut in the diamond network. Here $S \subset \{1, \ldots, N\}$, and the cut separates $u \cup \{v_n\}_{n \in S}$ from $w \cup \{v_n\}_{n \in S^c}$.

Finally, we point out that a (partial) decode-and-forward strategy is not sufficient to provide a uniform capacity approximation as in Corollary 3. Indeed, due to symmetry, *all* relays would be able to decode the source in any such strategy, which implies that decode-and-forward and partial decode-and-forward coincide in this case. The rate achievable with decode-and-forward is given by

$$\frac{1}{2} \log\big(1 + \min\{g, N^2 h\}\big).$$

Comparing this with Corollary 3, we see that (partial) decode-and-forward has an additive gap of at least $\Omega(\log(N))$ bits and a multiplicative gap of at least a factor $\Omega(N)$ to capacity. Similarly, as was pointed out earlier, the traditional amplify-and-forward strategy does not yield a constant factor approximation of capacity. In fact, in can be shown that simple amplify-and-forward results in unbounded multiplicative gap even for $N = 2$. Therefore the bursty amplify-and-forward scheme advocated in this work has the nice property of being uniformly approximately optimal in both the additive and multiplicative sense, as well as being a simple modification of the traditional amplify-and-forward scheme.

IV. Conclusion

We presented an approximation of the capacity of the symmetric Gaussian N-relay diamond network. The capacity was characterized up to a 1.8 bit additive gap and a factor 14 multiplicative gap uniformly for all channel gains and number of relays. The inner bound in this approximate characterization relies on bursty amplify-and-forward, showing that this scheme is good simultaneously at low and high rates, uniformly in the number of relays N. The upper bound resulted from a careful evaluation of the cut-set bound. We argued that all 2^N possible cuts in the diamond network need to be evaluated simultaneously, and that the standard approach of only considering the minimum of the broadcast and multiple-access cuts is insufficient to derive uniform capacity approximations.

References

[1] B. Schein and R. Gallager, "The Gaussian parallel relay network," in *Proc. IEEE ISIT*, p. 22, June 2000.

[2] B. Schein, *Distributed Coordination in Network Information Theory*. PhD thesis, Massachusetts Institute of Technology, 2001.

[3] M. Gastpar and M. Vetterli, "On the capacity of large Gaussian relay networks," *IEEE Trans. Inf. Theory*, vol. 51, pp. 765–779, Mar. 2005.

[4] Y. Kochman, A. Khina, U. Erez, and R. Zamir, "Rematch and forward for parallel relay networks," in *Proc. IEEE ISIT*, pp. 767–771, July 2008.

[5] S. S. C. Rezaei, S. O. Gharan, and A. K. Khandani, "A new achievable rate for the Gaussian parallel relay channel," in *Proc. IEEE ISIT*, pp. 194–198, June 2009.

[6] F. Xue and S. Sandhu, "Cooperation in a half-duplex Gaussian diamond relay channel," *IEEE Trans. Inf. Theory*, vol. 53, pp. 3806–3814, Oct. 2007.

[7] H. Bagheri, A. S. Motahari, and A. K. Khandani, "On the capacity of the half-duplex diamond channel," *arXiv:0911.1426 [cs.IT]*, Nov. 2009. submitted to IEEE Transactions on Information Theory.

[8] W. Kang and S. Ulukus, "Capacity of a class of diamond channels," *arXiv:0808.0948 [cs.IT]*, Aug. 2008. submitted to IEEE Transactions on Information Theory.

[9] R. Etkin and D. N. C. Tse, "Gaussian interference channel capacity to within one bit," *IEEE Trans. Inf. Theory*, vol. 54, pp. 5534–5562, Dec. 2008.

[10] A. S. Avestimehr, S. N. Diggavi, and D. N. C. Tse, "Wireless network information flow: A deterministic approach," *IEEE Trans. Inf. Theory*, Apr. 2011.

[11] A. Özgür and S. Diggavi, "Approximately achieving Gaussian relay network capacity with lattice codes," in *Proc. IEEE ISIT*, pp. 669–673, June 2010. see also arXiv:1005.1284 [cs.IT].

[12] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, pp. 388–404, Mar. 2000.

[13] A. Özgür, O. Lévêque, and D. N. C. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. Inf. Theory*, vol. 53, pp. 3549–3572, Oct. 2007.

[14] J. Ghaderi, L.-L. Xie, and X. Shen, "Hierarchical cooperation in ad hoc networks: Optimal clustering and achievable throughput," *IEEE Trans. Inf. Theory*, vol. 55, pp. 3425–3436, Aug. 2009.

[15] U. Niesen, P. Gupta, and D. Shah, "On capacity scaling in arbitrary wireless networks," *IEEE Trans. Inf. Theory*, vol. 56, pp. 3959–3982, Sept. 2009.

[16] U. Niesen, P. Gupta, and D. Shah, "The balanced unicast and multicast capacity regions of large wireless networks," *IEEE Trans. Inf. Theory*, vol. 56, pp. 2249–2271, May 2010.

[17] U. Niesen, "Interference alignment in dense wireless networks," *IEEE Trans. Inf. Theory*, vol. 57, pp. 2889–2901, May 2011.

[18] U. Niesen and S. Diggavi, "The approximate capacity of the Gaussian N-relay diamond network," *arXiv:1008.3813 [cs.IT]*, Aug. 2010. Submitted to IEEE Transactions on Information Theory.