Computation over Gaussian Multiple-Access Channels

Bobak Nazer and Michael Gastpar
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720, USA
Email: {bobak, gastpar}@eecs.berkeley.edu

***Abstract*— We consider the problem of computing the sum of independent Gaussian sources over a Gaussian multiple-access channel (MAC) with respect to a mean-squared error criterion. When the source and channel bandwidths are equal, the best separation-based solution to this problem performs exponentially worse in a distortion sense compared to the optimal solution: uncoded transmission. In this paper, we develop lattice codes for exploiting the structure of the Gaussian MAC when there are more channel uses than source symbols. We also demonstrate the usefulness of these codes for multicasting over a simple AWGN network. This version corrects an error that appears in the published version.**

I. Introduction

In the standard multiple-access problem, a central access point needs to reconstruct the signals observed by each user separately. Suppose now that the access point only reconstructs the sum of these observed signals. In previous work, we have shown that there is no separation principle even if the sources are independent [1]. We have developed a coding strategy, computation coding, for reliably computing an arbitrary function of discrete sources over a MAC [2]. We have also shown that computation coding can be useful in determining the multicast capacity of networks that include MACs [3]. In this paper, we examine the problem of recovering the sum of independent Gaussian sources over a Gaussian MAC with respect to a mean-squared error criterion.

Körner and Marton showed that if we are only interested in compressing the parity of two correlated binary sources, then the optimal performance can be achieved by having each terminal employ the same linear code [4]. The required sum rate is lower than that required by a Slepian-Wolf scheme so long as the sources are dependent. In [1], we showed that if each source is seen at one terminal of an appropriately matched MAC, then using the same linear source code and channel code at each terminal attains the optimal performance. We modify the above strategy for the Gaussian problem by using lattices as the underlying linear codes. Although we do not have an outer bound that matches our scheme, it performs significantly better than a separation-based strategy and well enough to characterize the multicast capacity of an AWGN butterfly network that includes a MAC.

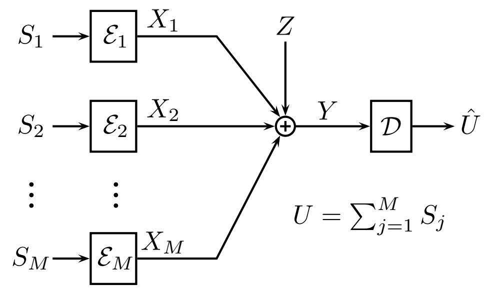

Fig. 1. Reliable Addition over a Gaussian MAC

II. Problem Statement

Each encoder, $\mathcal{E}_j$, sees an independent identically distributed (i.i.d.) Gaussian sequence $\{S_j[i]\}_{i=1}^k$ with mean 0 and variance σ_S^2. For every k source symbols, we are allotted $n = \ell k + r$ channel uses where $\ell, r \in \mathbb{Z}_+$ and $r < k$.

$$\mathcal{E}_j : \mathbb{R}^k \to \mathbb{R}^n \tag{1}$$

The encoders must satisfy average power constraints:

$$\frac{1}{n}\sum_{i=1}^{n} x_j[i]^2 \leq P \quad \forall j \in \{1, 2, \ldots, M\} \tag{2}$$

The channel output is just the sum of the channel inputs plus independent Gaussian noise:

$$Y[i] = \sum_{j=1}^{M} X_j[i] + Z[i] \tag{3}$$

where $\{Z[i]\}_{i=1}^n$ is an i.i.d. Gaussian sequence with mean 0 and variance σ_Z^2.

Our goal is to reconstruct the sum of the sources, $U = S_1 + S_2 + \cdots + S_M$, at the decoder with the lowest possible distortion. Distortion is measured by the usual mean-squared error criterion:

$$D = \frac{1}{k}\sum_{i=1}^{k} E[(U_i - \hat{U}_i)^2] \tag{4}$$

III. Lower Bound

We now give a lower bound on the best achievable distortion.

Lemma 1: The best achievable distortion for sending k sums of independent Gaussian sources over a Gaussian MAC with n channel uses is lower bounded by:

$$D_{\text{LOWER}} = M\sigma_S^2 \left(\frac{\sigma_Z^2}{\sigma_Z^2 + MP} \right)^{\frac{n}{k}} \tag{5}$$

Proof: Using steps from the converse to the multiple-access problem (see [5, pp. 399-407]) as well as the independence of the sources, we can get that $I(X_1, X_2, \ldots, X_M; Y) \leq \frac{1}{2}\log\left(1 + \frac{MP}{\sigma_Z^2}\right)$. It is also clear that the rate distortion function for jointly compressing the sum is given by $R_U(D) = \frac{1}{2}\log\left(\frac{M\sigma_S^2}{D}\right)$. By applying the data processing inequality, we get the desired lower bound. ■

Lemma 2: If $k = n$, then uncoded transmission is optimal for sending k sums of independent Gaussian sources over a Gaussian MAC with n channel uses.

Essentially, this is the distributed computation extension of the famous fact that uncoded transmission is optimal for sending a Gaussian source over an additive white noise Gaussian (AWGN) channel. In the point-to-point setting, we are able to use the separation theorem to optimally send a Gaussian source over an AWGN channel for any ratio $\frac{k}{n}$. However, for computing a sum over a Gaussian MAC, there is no separation theorem. As in the discrete case, we can improve our communication system by taking advantage of the structure of the MAC.

IV. Separation-Based Computation

We now show that if we want to reconstruct the sum of independent Gaussian sources by separation, the encoders can do no better than send their sources to the decoder.

Lemma 3: The sum rate distortion function for the distributed compression of the sum of i.i.d. Gaussian sources with mean 0 and variance σ_S^2 is:

$$R(D) = \frac{M}{2}\log\left(\frac{M\sigma_S^2}{D}\right) \tag{6}$$

Remark 1: Note that joint source coding of the sum requires a much lower rate, $R(D) = \frac{1}{2}\log\left(\frac{M\sigma_S^2}{D}\right)$. This penalty for distributedness is similar to that seen in the CEO problem.

Proof: (*Converse.*) Let $f_j : \mathbb{R}^k \rightarrow \{1, 2, \ldots, 2^{kR_j}\}$ be the j^{th} source encoding function and let $W_j = f_j\left(\{S_j[i]\}_{i=1}^k\right)$ be the message output by this encoder for a length k block of source symbols. Given $\mathbf{W} = (W_1, W_2, \ldots, W_M)$ at the decoder, the minimum-mean squared estimate (MMSE) of U is given by the conditional expectation.

$$D \geq \frac{1}{k}\sum_{i=1}^{k} E[(U[i] - E[U(i)|\mathbf{W}])^2] \tag{7}$$

$$\stackrel{(a)}{=} \frac{1}{k}\sum_{i=1}^{k} E\left[\left(\sum_{j=1}^{M} S_j[i] - E\left[\sum_{j=1}^{M} S_j[i] \middle| W_j\right]\right)^2\right] \tag{8}$$

$$\stackrel{(b)}{=} \frac{1}{k}\sum_{j=1}^{M}\sum_{i=1}^{k} E[(S_j[i] - E[S_j[i]|W_j])^2] \tag{9}$$

$$\stackrel{(c)}{\geq} \sum_{j=1}^{M} \sigma_S^2 2^{-2R_j} \tag{10}$$

(a), (b) by independence of S_i and S_j for all $i \neq j$
(c) by the single source rate distortion converse (see [5, pp. 350-351])

Minimizing the function $\sum_{j=1}^{M} \sigma_S^2 2^{-2R_j}$ is just a convex optimization problem subject to the convex constraint $\sum_{j=1}^{m} R_j = R$. It easily follows that the minimizing solution satisfies $R_1 = R_2 = \cdots = R_M$. We finally get that:

$$D \geq M\sigma_S^2 2^{-2\frac{R}{M}} \tag{11}$$

$$R(D) \geq \frac{M}{2}\log\left(\frac{M\sigma_S^2}{D}\right) \tag{12}$$

(*Achievability*.) Each encoder simply uses a standard Gaussian rate distortion code for its source with distortion target $D_j = \frac{D}{M}$. Such a code requires a rate of at least $\frac{1}{2}\log\left(\frac{M\sigma_S^2}{D}\right)$ per encoder. See [5, pp. 351-358] for the derivation of such a code. The decoder recovers each source and sums the individual estimates to get an estimate of the desired sum. ■

Lemma 4: The best achievable distortion for a separation-based scheme for sending k sums of independent Gaussian sources over a Gaussian MAC with n channel uses is given by:

$$D_{\text{SEP}} = M\sigma_S^2 \left(\frac{\sigma_Z^2}{\sigma_Z^2 + MP} \right)^{\frac{n}{Mk}} \tag{13}$$

Proof: We employ the source code detailed in the proof of Lemma 3 in conjuction with a channel code that achieves the maximum sum rate. The converse follows from combining the converse steps from Lemma 3 with the standard MAC converse (see [5, pp. 399-407]). ■

V. Lattices for Computation

Lattice codes are the AWGN equivalent of linear codes for discrete memoryless channels. A great deal work has gone into showing that lattice codes can achieve capacity on an AWGN channel and the rate distortion bound for a Gaussian source [6]–[11]. Like linear codes, the primary appeal of lattices has been their lower complexity compared to purely random coding strategies. However, their structural properties can also be exploited for distributed computation. First, we will need some definitions from [10].

A. Lattice Preliminaries

Definition 1: An n-dimensional *lattice*, Λ, is a set of points in $\mathbb{R}^n$ such that if $\mathbf{x}, \mathbf{y} \in \Lambda$, then $\mathbf{x} + \mathbf{y} \in \Lambda$, and if $\mathbf{x} \in \Lambda$, then $-\mathbf{x} \in \Lambda$. A lattice can always be written in terms of a generator matrix $\mathbf{G} \in \mathbb{R}^{n \times n}$:

$$\Lambda = \{\mathbf{x} = \mathbf{z}\mathbf{G} : \mathbf{z} \in \mathbb{Z}^n\} \tag{14}$$

where $\mathbb{Z}$ represents the integers.

Definition 2: A *lattice quantizer* is a map, $Q : \mathbb{R}^n \to \Lambda$, that sends a point, $\mathbf{x}$, to the nearest lattice point in Euclidean distance:

$$\mathbf{x_q} = Q(\mathbf{x}) = \arg\min_{\mathbf{l} \in \Lambda} ||\mathbf{x} - \mathbf{l}||_2 \tag{15}$$

Definition 3: Let $\mathbf{x} \bmod \Lambda = \mathbf{x} - Q(\mathbf{x})$. The mod Λ operation satisfies:

$$((\mathbf{x} \bmod \Lambda) + \mathbf{y}) \bmod \Lambda = (\mathbf{x} + \mathbf{y}) \bmod \Lambda \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n \tag{16}$$

Definition 4: The *fundamental Voronoi region*, $\mathcal{V}$, of a lattice, is the set of all points that are closest to the zero vector: $\mathcal{V}_0 = \{\mathbf{x} : Q(\mathbf{x}) = \mathbf{0}\}$.

Definition 5: The *normalized second moment* of a lattice is:

$$G(\Lambda) = \frac{1}{k} \frac{\int_{\mathcal{V}} ||\mathbf{x}||^2 d\mathbf{x}}{\left[\int_{\mathcal{V}} d\mathbf{x}\right]^{1+\frac{2}{k}}} \tag{17}$$

Erez, Litsyn and Zamir showed in [9] that there exist lattices that are simultaneously good source codes and good channel codes. These will be extremely useful in our distributed refinement scheme.

Lemma 5 (Erez-Litsyn-Zamir): There exist sequences of lattices, Λ_k, such that the normalized second moment, $G(\Lambda_k)$, is asymptotically optimal:

$$\lim_{k \to \infty} G(\Lambda_k) = \frac{1}{2\pi e} \tag{18}$$

and with high probability an i.i.d. Gaussian sequence, $\{S[i]\}_{i=1}^k$, with mean zero and variance $G(\Lambda_k)$ falls within the fundamental Voronoi region, $\mathcal{V}_{0,k}$:

$$\lim_{k \to \infty} \Pr\left(\{S[i]\}_{i=1}^k \in \mathcal{V}_{0,k}\right) = 1 \tag{19}$$

See [9] for a full proof.

We will now use these lattices to develop two schemes: one that completely rejects the channel noise and another that makes use of an uncoded transmission phase.

B. Achievable Scheme

In [10], Kochman and Zamir develop an elegant joint source-channel lattice scheme for sending a Wyner-Ziv Gaussian source over a dirty paper channel. Our distributed refinement scheme consists of two main steps. First, we use uncoded transmission to send a noisy sum to the decoder. Then, we have each encoder run a version of the Kochman-Zamir scheme targeted at the desired sum, U. Unfortunately, there is a penalty for this form of distributedness. The lattice at each encoder results in channel outputs that violate the power constraint by a factor of M. Therefore, we must scale down our inputs to meet the power constraint and accept the resulting increase in distortion at the decoder.

Theorem 1: The following distortion is achievable for sending k sums of independent Gaussian sources over a Gaussian MAC with $n = \ell k, \ \ell \in \mathbb{Z}_+$ channel uses if $P > \frac{M-1}{M}\sigma_Z^2$:

$$D_{\text{wz}} = M\sigma_S^2 \left(\frac{\sigma_Z^2}{\sigma_Z^2 + MP}\right)\left(\frac{M\sigma_Z^2}{\sigma_Z^2 + MP}\right)^{\ell - 1} \tag{20}$$

Proof: We will first show the achievable scheme for $\ell = 2$. We thus have $2k$ channel uses to convey k sums. We will use the first k channel uses for an uncoded transmission phase as in Lemma 2. The decoder will then form an MMSE estimate $\hat{U}$ of the sum $U = S_1 + \cdots + S_M$ and use this as side information for the next phase. Thus, $U = Q + \hat{U}$ where Q is an i.i.d. Gaussian sequence with mean 0 and variance $\sigma_Q^2 = M\sigma_S^2 \frac{\sigma_Z^2}{\sigma_Z^2 + MP}$. Choose a sequence of good lattices, Λ_k, using Lemma 5 and scale them such that the normalized second moment of the lattice is MP. Let $\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_M$ be independent dither vectors drawn uniformly over the fundamental Voronoi region, $\mathbf{d}_j \sim \text{Unif}(\mathcal{V}_{0,k})$, and made available to the encoders and decoder.

Each encoder transmits $\frac{1}{\sqrt{M}}\mathbf{x}_j$ where:

$$\mathbf{x}_j = [\gamma \mathbf{s}_j + \mathbf{d}_j] \bmod \Lambda_k \tag{21}$$

The channel output is given by:

$$\mathbf{y} = \frac{1}{\sqrt{M}} \sum_{j=1}^{M} \mathbf{x}_j + \mathbf{z} \tag{22}$$

The decoder then computes:

$$\mathbf{t} = \alpha \mathbf{y} - \left(\sum_{j=1}^{M} \mathbf{d}_j + \gamma \hat{\mathbf{u}}\right) \tag{23}$$

$$\begin{aligned} \mathbf{r} &= \mathbf{t} \bmod \Lambda_k & (24) \\ &= \left[\frac{\alpha}{\sqrt{M}} \sum_{j=1}^{M} \mathbf{x}_j + \alpha \mathbf{z} - \sum_{j=1}^{M} (\mathbf{d}_j + \gamma \mathbf{s}_j) + \gamma \mathbf{q}\right] \bmod \Lambda_k \\ &= \left[\left(\frac{\alpha}{\sqrt{M}} - 1\right) \sum_{j=1}^{M} \mathbf{x}_j + \alpha \mathbf{z} + \gamma \mathbf{q}\right] \bmod \Lambda_k & (25) \end{aligned}$$

If the second moment of the term inside the modulo operation does not exceed MP, the second moment of the lattice, then we can guarantee that:

$$\lim_{k \to \infty} \Pr\left(\mathbf{r} = \left(\frac{\alpha}{\sqrt{M}} - 1\right) \sum_{j=1}^{M} \mathbf{x}_j + \alpha \mathbf{z} + \gamma \mathbf{q}\right) = 1 \tag{26}$$

See [7] for a detailed discussion of the effect of the dither in this step. The second moment can be controlled by requiring that:

$$\left(\frac{\alpha}{\sqrt{M}} - 1\right)^2 (M^2 P) + \alpha^2 \sigma_Z^2 + \gamma^2 \sigma_Q^2 \le MP \tag{27}$$

This equation will be satisfied by our final choice of the constants α,β, and γ. The decoder's estimate of the sum is given by:

$$\hat{\hat{\mathbf{u}}} = \beta\mathbf{r} + \hat{\mathbf{u}} \quad (28)$$

$$= \beta\left(\left(\frac{\alpha}{\sqrt{M}} - 1\right)\sum_{j=1}^{M}\mathbf{x}_j + \alpha\mathbf{z} + \gamma\mathbf{q}\right) + \hat{\mathbf{u}} \quad (29)$$

$$= \beta\left(\left(\frac{\alpha}{\sqrt{M}} - 1\right)\sum_{j=1}^{M}\mathbf{x}_j + \alpha\mathbf{z}\right) - (1-\beta\gamma)\mathbf{q} + \hat{\mathbf{u}} \quad (30)$$

This estimate gives the following mean-squared error:

$$\beta^2\left(\left(\frac{\alpha}{\sqrt{M}} - 1\right)^2 M^2P + \alpha^2\sigma_Z^2\right) + (1-\beta\gamma)^2\sigma_Q^2 \quad (31)$$

We define the following constants:

$$\alpha = \frac{MP\sqrt{M}}{MP + \sigma_Z^2} \quad (32)$$

$$\gamma_0 = \sqrt{\frac{MP}{\sigma_Q^2}\left(1 - \frac{M\sigma_Z^2}{MP + \sigma_Z^2}\right)} \quad (33)$$

and let $\gamma \to \gamma_0$ from below as $k \to \infty$. This ensures that Equation (27) is always satisfied. We also set:

$$\beta = \frac{\sigma_Q^2\gamma}{MP} \quad (34)$$

As $k \to \infty$, we get that the achieved distortion is:

$$D = M\sigma_S^2\frac{\sigma_Z^2}{\sigma_Z^2 + MP}\frac{M\sigma_Z^2}{\sigma_Z^2 + MP} \quad (35)$$

This proves the theorem for $\ell = 2$. For all higher values of ℓ, the scheme can be repeated with the final estimate from the last refinement taken as side information for the next stage. ■

Remark 2: A simple achievable scheme for situations where we do not have an integer number of channel uses per source symbol is to time share between two integers whose average gives the proper ratio.

Remark 3: The published version of this paper contains two different schemes for computing over the Gaussian MAC with lattices. The "reliable refinement" scheme originally presented as Theorem 1 had a mistake in the proof. The original claim was that the distortion could be made to fall like $\left(\frac{\sigma_Z^2}{\sigma_Z^2+P}\right)^\ell$. Unfortunately, after fixing the proof, the distortion rises to $\left(\frac{M\sigma_Z^2}{\sigma_Z^2+MP}\right)^\ell$ which is inferior to the scheme that was originally presented as Theorem 2 (now Theorem 1).

For comparison, note that with repetition coding we can achieve the following distortion:

$$D_{\text{REP}} = M\sigma_S^2\frac{\sigma_Z^2}{\sigma_Z^2 + \ell MP} \quad (36)$$

where $n = \ell k$ for some $\ell \in \mathbb{Z}_+$.

Example 1: Let $M = 5$, $P = 10$, $\sigma_Z^2 = 1$, and $\sigma_S^2 = 1$. The end-to-end distortions for repetition coding, separation-based coding, our computation coding scheme, and our lower bound are given in Figure 2 (from equations (36), (13), (20), and (5) respectively). Note that as the number of channel uses per source symbol increases, our scheme performs exponentially better than separation-based coding.

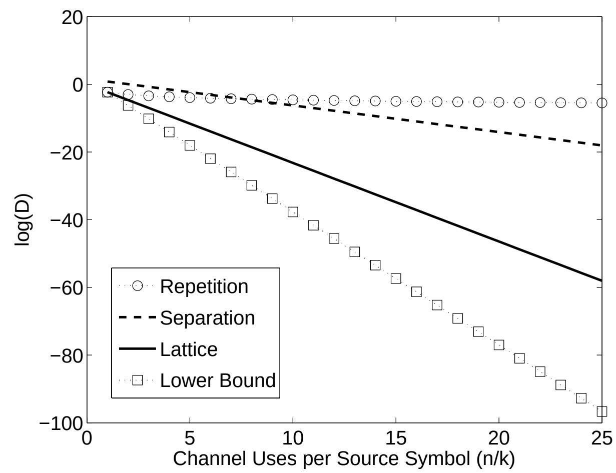

Fig. 2. Refining the Sum of Gaussian Sources over a Gaussian MAC, $M = 5$, $P = 10$, $\sigma_Z^2 = 1$, $\sigma_S^2 = 1$

VI. Multicasting over an AWGN Butterfly Network

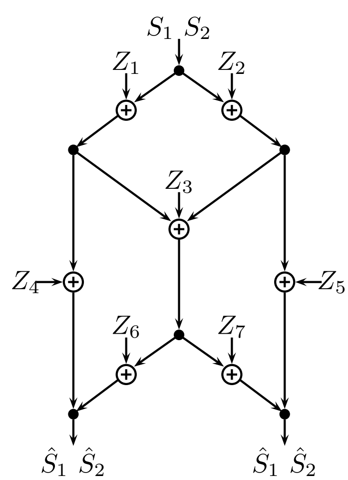

Fig. 3. Network Coding Example

Consider the AWGN channel network in Figure 3. This is essentially the butterfly network introduced in [12] with a Gaussian MAC at the center. Each vertex on the graph represents a decoder/encoder pair. All encoders must satisfy an average power constraint, $\frac{1}{n}\sum_{i=1}^{n} x_j[i]^2 \leq P$. The $Z_m,\ m = 1, 2, \ldots, 7$ are drawn i.i.d. according to a Gaussian distribution with mean 0 and variance σ_Z^2. We would like to determine the multicast capacity of this channel network. Note that the source node is at the top of the graph and the two receivers are at the bottom.

Theorem 2: The following multicast rate is achievable for the channel network in Figure 3:

$$C = \frac{1}{2}\log\left(1+\frac{P}{\sigma_Z^2}\right) + \frac{1}{2}\log\left(\frac{1}{2}+\frac{P}{\sigma_Z^2}\right) \qquad (37)$$

Proof: For ease of analysis, we will first consider sending two i.i.d. Gaussian sequences, $\{S_1[i]\}_{i=1}^k$ and $\{S_2[i]\}_{i=1}^k$, with mean 0 and variance σ_S^2 using $n = \ell k$ channel uses for some $\ell \in \mathbb{Z}_+$ to be specified later. We choose a k dimensional lattice according to the proof of Theorem 1. Using this lattice, we successively quantize the sources ℓ times just as in the proof of Theorem 1 such that we can reconstruct the sum to distortion $D \triangleq 2\sigma_S^2 \left(\frac{2\sigma_Z^2}{\sigma_Z^2+2P}\right)^\ell$. We transmit the quantized S_1 sequence down the left path and the quantized S_2 sequence down the right path. Note that there is still some leftover rate on the left and right paths. At the multiple-access point, each encoder runs the encoding scheme according to Theorem 1 such that we receive the sum at the decoder at distortion D. This sum is requantized to distortion at most D so that it can be conveyed to the receivers.

At the left receiver we have estimates $\hat{\mathbf{u}}$ and $\hat{\mathbf{s}}_\mathbf{1}$ and we form an estimate $\hat{\hat{\mathbf{s}}}_\mathbf{2} = \hat{\mathbf{u}} - \hat{\mathbf{s}}_\mathbf{1}$. The left receiver knows the source vectors at distortions $\sigma_S^2\left(\frac{\sigma_Z^2}{\sigma_Z^2+P}\right)^\ell$ and $9\sigma_S^2\left(\frac{\sigma_Z^2}{\sigma_Z^2+P}\right)^\ell$, respectively. (The distortions for the right receiver follow by analogy.) Now, if we instead consider using these Gaussian sequences to communicate codewords from a random Gaussian codebook, we can multicast at a rate $R = \frac{1}{2}\log\left(1+\frac{P}{\sigma_Z^2}\right) + \frac{1}{2}\log\left(\left(\frac{1}{9}\right)^{\frac{1}{\ell}}\left(1+\frac{P}{\sigma_Z^2}\right)\right)$. This approaches $\log\left(\frac{1}{2}+\frac{P}{\sigma_Z^2}\right)$ as ℓ increases. Finally, we note that we have $\frac{1}{2}\log\left(1+\frac{P}{\sigma_Z^2}\right) - \frac{1}{2}\log\left(\frac{1}{2}+\frac{P}{\sigma_Z^2}\right)$ rate remaining on both the left and right paths and we use this to transmit a common message. ∎

Remark 4: In the original version of this paper, the above theorem claimed to achieve the multicast capacity of the network in Figure 3. This was because we made use of the scheme that refined the sum like $\left(\frac{\sigma_Z^2}{\sigma_Z^2+P}\right)^\ell$. Without that scheme, we have to accept a slightly lower rate for multicasting over the butterfly network.

VII. Discussion

For complete recovery of independent sources, a separation-based scheme is clearly optimal. However, when we are only interested in the sum then the summation performed by the Gaussian MAC can be quite useful. This MAC property has already been exploited by uncoded transmission for the case when the source and channel have the same bandwidth. By repeatedly applying a lattice code to refine our sources, we have shown that we can continue to reap some of the benefits of uncoded transmission even when there are many more channel uses than source symbols. Instead of designing channel codes that avoid the interference created by other users, our lattice computation codes make use of collisions to reliably add signals over the MAC.

By using our computation coding scheme as part of an overall network code, we can decode and retransmit the sum of messages. Using uncoded transmission to achieve the same effect results in noise building up as we traverse the network. We are currently investigating a generalized version of this strategy in order to determine the multicast capacity of general AWGN networks.

Acknowledgment

The authors would like to thank K. Eswaran, V. Prabhakaran, and A. D. Sarwate for helpful discussions. In particular, K. Eswaran provided a proof of Lemma 3. B. Nazer was supported by an NSF Graduate Fellowship. M. Gastpar was supported by NSF Grant CCR 0347298.

References

[1] B. Nazer and M. Gastpar, "Reliable computation over multiple-access channels," in *43rd Annual Allerton Conference, Monticello, IL, September, 2005.*

[2] B. Nazer and M. Gastpar, "Computation over multiple-access channels," *IEEE Transactions on Information Theory*. Submitted Aug. 2006.

[3] B. Nazer and M. Gastpar, "Computing over multiple-access channels with connections to wireless network coding," in *IEEE ISIT 2006, Seattle, WA*.

[4] J. Körner and K. Marton, "How to encode the modulo-two sum of binary sources," *IEEE Trans. Inform. Theory*, vol. 25, pp. 219–221, March 1979.

[5] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley-Interscience, 1991.

[6] R. Urbanke and B. Rimoldi, "Lattice codes can achieve capacity on the AWGN channel," *IEEE Trans. Inform. Theory*, vol. 44, pp. 273–278, January 1998.

[7] U. Erez and R. Zamir, "Achieving $\frac{1}{2}\log(1+\mathrm{SNR})$ on the AWGN channel with lattice encoding and decoding," *IEEE Trans. Inform. Theory*, vol. 50, pp. 2293–2314, October 2004.

[8] R. Zamir, S. Shamai, and U. Erez, "Nested linear/lattice codes for structured multiterminal binning," *IEEE Trans. Inform. Theory*, vol. 48, pp. 1250–1276, June 2002.

[9] U. Erez, S. Litsyn, and R. Zamir, "Lattices which are good for (almost) everything," *IEEE Trans. Inform. Theory*, vol. 51, pp. 3401–3416, October 2005.

[10] Y. Kochman and R. Zamir, "Analog matching of colored sources to colored channels," in *IEEE ISIT 2006, Seattle, WA*.

[11] S. D. Servetto, "Lattice quantization with side information," in *IEEE DCC 2000, Snowbird, UT*.

[12] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, pp. 1204–1216, July 2000.