On Wyner-Ziv Networks

Michael Gastpar
Department of EECS
University of California, Berkeley
Berkeley, CA 94720-1770
gastpar@eecs.berkeley.edu

Abstract— **Wyner and Ziv [1] determined the rate-distortion function of source coding with side information. In this paper, we consider a network extension of their scenario: Many sources have to be compressed in a rate-distortion sense for a decoder that has access to uncompressed side information about the sources. This may model a sensor or multi-camera network situation where the data collection point itself also has a sensor/camera attached to it. Basic results for the case of two sources and side information have been presented recently [2], [3]. The present paper extends these results to the case of more than two sources. In particular, we derive an achievable rate-distortion region, and an outer bound to the best rate-distortion region. These two do not coincide in general, but they do in the special case where the sources are conditionally independent given the side information. We illustrate this result with applications. For example, we discuss two different kinds of rate losses of distributed compression as compared to joint (i.e., centralized) compression. Thereafter, we show that in some cases, our result permits to derive bounds to the rate-distortion region for the distributed compression problem without side information, and we also show how a certain source-channel separation theorem can be established.**

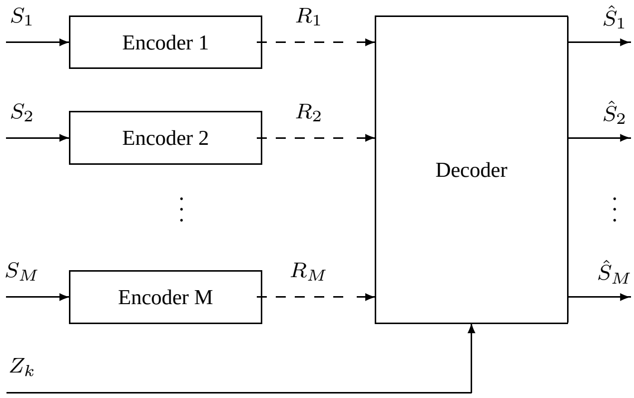

Fig. 1. The Wyner-Ziv network topology.

I. Introduction

In this paper we consider the problem of source encoding with fidelity criteria in a situation where the decoder has access to side information about the sources. To put the problem in perspective, consider the the source network topology in Figure 1.

The sequence $\{(S_{1,k}, S_{2,k}, \ldots, S_{M,k}, Z_k)\}_{k=1}^{\infty}$ represents independent copies of an $(M+1)$-tuple of dependent random variables $(S_1, S_2, \ldots, S_M, Z)$ which take values in finite sets $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_M, \mathcal{Z}$, respectively, distributed according to a fixed and known distribution

$$p(s_1, s_2, \ldots, s_M, z). \tag{1}$$

The encoder outputs are binary sequences which appear at rates $R_1, R_2, \ldots, R_M$ bits per input symbol, respectively. The decoder output is a sequence of M-tuples $\{(\hat{S}_{1,k}, \hat{S}_{2,k}, \ldots, \hat{S}_{M,k})\}_{k=1}^{\infty}$ whose components take values in finite reproduction alphabets $\hat{\mathcal{S}}_1, \hat{\mathcal{S}}_2, \ldots, \hat{\mathcal{S}}_M$. The encoding is done in blocks of length n, and the fidelity criteria, for $m = 1, 2, \ldots, M$, take the shape

$$E\left[\frac{1}{n}\sum_{k=1}^{n} d_m(x_{m,k}, \hat{x}_{m,k})\right], \tag{2}$$

where $d_m(s_m, \hat{s}_m) \geq 0, s_m \in \mathcal{S}_1, \hat{s}_m \in \hat{\mathcal{S}}_1$, is a given distortion function. We define $\mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M)$ as the set of rates for which the system of Figure 1 can operate when n is large and the average distortions given in (2) are arbitrarily close to D_m, for $m = 1, 2, \ldots, M$.

As for applications, the network shown in Figure 1 may be of interest in certain sensor network applications. In that case, the random variables $S_1, S_2, \ldots, S_M$ represent the sensor readings. Encoder 1 represents the first sensor, whose reading is S_1, and whose task is to encode this reading, using a fixed rate R_1. There are M sensors, and their goal is to inform a central data collector, who in turn attempts to estimate the sensor readings $S_1, S_2, \ldots, S_M$ at the highest possible fidelity. To simplify this task, the central data collector has access to "side information", namely, to an additional sensor reading Z.

In Section II, we briefly quote the result for the single-source case, i.e., for the situation of Figure 1 with $M = 1$. In Section III, we extend the arguments leading to the single-source result to the case of multiple sources. In particular, we present inner and outer bounds to the rate-distortion region, and we note that the two do not coincide in general. In Section IV, we study the special case where the sources are conditionally independent given the side information. For that case, we find that the inner and outer bounds of Section III coincide, which enables us also to analyze the question of the *rate loss* of distributed encoding with respect to joint encoding. In Section V, we discuss a few applications and consequences of the result of Section IV. Finally, we conclude and outline extensions in Section VI.

II. The Single-Source Case

Wyner and Ziv [1] solved the problem shown in Figure 1 for the case of a single source (i.e., $M = 1$) and side information Z. Denote the smallest rate R required to achieve a prescribed distortion of D by $R^{WZ}_{S_1|Z}(D)$. For the case where the source has a discrete alphabet, the solution is established in [1]. It can be expressed as

$$R^{WZ}_{S_1|Z}(D) \quad = \quad \min_{p(w|s_1)} I(S_1; W) - I(Z; W), \tag{3}$$

where the minimum is taken over all $p(w|s_1)$ and functions $g(\cdot)$ satisfying

$$Ed(S_1, g(Z, W)) \quad \leq \quad D. \tag{4}$$

The result is extended to a class of sources with continuous alphabets (and under appropriate distortion measures) in [4].

III. The Multi-Source Case

The extension considered in this paper is illustrated in Figure 1. As pointed out in Section I, we assume that the involved sources $S_m, m = 1, 2, \ldots, M$ as well as the side information Z take values in finite alphabets. Extensions to certain classes of sources with continuous alphabets, under appropriate distortion measures, may be established along the lines of [4]. The goal of this paper is to find the rate-distortion region

$$\mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M). \tag{5}$$

In words, this is the set of rate tuples $(R_1, R_2, \ldots, R_M)$ that permit to reconstruct the source S_m at a distortion of at most D_m, for $m = 1, 2, \ldots, M$. This rate-distortion region is not known to date, and its determination must be suspected to be a hard problem since the problem illustrated in Fig. 1 can be seen as an extension of the standard distributed compression problem, simply by omitting the side information Z. However, the latter is an old open problem; the best general achievable rates and outer bounds appear in [5], and do not coincide. We briefly discuss this problem, and its relationship to the Wyner-Ziv problem studied in this paper, in Section V-A.

In this paper, we present general inner and outer bounds to the region (5). These bounds to not generally match. Thereafter, we consider the special case where the sources are conditionally independent given the side information. For that case, the presented inner and outer bounds coincide, establishing a rate-distortion result. Moreover, it can be shown that in this special case, the *rate loss* of distributed versus joint compression vanishes.

A. Achievable Rate Tuples

In [2], [3], we extended the coding idea of [1] (based on [6]) to derive a set of achievable rate pairs for the networks of Figure 1 in the special case where $M = 2$. These arguments can be extended easily to provide an achievable rate-distortion region for the case of M sources, as follows.

Theorem 1 (achievable rate tuples) An inner bound to the rate-distortion region is,

$$\begin{aligned}&\mathcal{R}_a(D_1, D_2, \ldots, D_M)\\ &\quad\subseteq \quad \mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M),\end{aligned} \tag{6}$$

where $\mathcal{R}_a(D_1, D_2, \ldots, D_M)$ is the set of all rate tuples $(R_1, R_2, \ldots, R_M)$ for which there exist discrete random variables $(W_1, W_2, \ldots, W_M)$ whose joint distribution satisfies

$$\begin{aligned}&p(w_1, w_2, \ldots, w_M, s_1, s_2, \ldots, s_M, z)\\ &\quad= \quad p(s_1, s_2, \ldots, s_M, z) \prod_{m=1}^{M} p(w_m|s_m)\end{aligned} \tag{7}$$

for which, for each $\mathcal{T} \subseteq \{1, 2, \ldots, M\}$, the following condition is satisfied:

$$\sum_{m \in \mathcal{T}} R_m \quad \geq \quad I(S_{\mathcal{T}}; W_{\mathcal{T}} | W_{\mathcal{T}^c}, Z) \tag{8}$$

and for which there exist functions $g_m(\cdot)$, for $m = 1, 2, \ldots, M$, such that

$$Ed_m(S_m, g_m(W_1, W_2, \ldots, W_M, Z)) \quad \leq \quad D_m, \tag{9}$$

for $m = 1, 2, \ldots, M$.

Remark. It can also be shown that it is sufficient to consider auxiliary random variables $(W_1, W_2, \ldots, W_M)$ over alphabets $\mathcal{W}_m$ of cardinalities

$$|\mathcal{W}_m| \quad \leq \quad |\mathcal{S}_m| + 1, \tag{10}$$

for $m = 1, 2, \ldots, M$. Larger alphabets will not reveal any new rate-distortion region points.

The proof of this theorem is a simple extension of the proof given in [2]. For an outline of the proof, the coding scheme that achieves this performance involves a "binning" operation at each of the M encoders. Roughly speaking, each encoder first applies a suitable vector quantizer to its source, obtaining a respective quantization index. At the time of code design, the set of all quantization indices that can occur at encoder m is (e.g., randomly) partitioned (sometimes referred to as "binned"), separately for each m, $m = 1, 2, \ldots M$. Instead of providing the decoder with the full quantization index, each encoder m only furnishes the index of the partition to which the actually observed quantization index belongs. If the partitions ("bins") are small enough, then the decoder can exploit their dependence and the side information to recover, inside each of the M partitions, which quantization index was actually observed by encoder m, $m = 1, 2, \ldots M$. The rate conditions (8) ensure that the partitions are small enough. Finally, the side information Z can be used a second time in order to undo a part of the quantization error.

From Theorem 1, it is easy to determine achievable rate-distortion points: one simply picks *any* auxiliary random variables that satisfies (7). The corresponding rate-distortion point is found by evaluating (8) for all $\mathcal{T} \subseteq \{1, 2, \ldots, M\}$, which yields the rates, and then finding the functions $g_m(\cdot)$, for $m = 1, 2, \ldots, M$, which yields the corresponding distortions.

B. Unachievable Rate Tuples

A converse (outer) bound can be given in extension of the outer bound in [5, Thm.6.2], as follows:

Theorem 2 (unachievable rate tuples) A converse bound to the rate-distortion region is

$$\begin{aligned}&\mathcal{R}_o(D_1, D_2, \ldots, D_M)\\ &\quad\supseteq \quad \mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M)\end{aligned} \tag{11}$$

where $\mathcal{R}_o(D_1, D_2, \ldots, D_M)$ is the set of all rate tuples $(R_1, R_2, \ldots, R_M)$ such that there exist discrete random variables $(W_1, W_2, \ldots, W_M)$ over alphabets $\mathcal{W}_m$ of cardinalities

$$\prod_{m=1}^{M} |\mathcal{W}_m| \quad \leq \quad 1 + \prod_{m=1}^{M} |\mathcal{S}_m|, \tag{12}$$

with

$$p(w_m|s_1, s_2, \ldots, s_M, z) = p(w_m|s_m), \quad (13)$$

for all $m = 1, 2, \ldots, M$, for which, for each $\mathcal{T} \subseteq \{1, 2, \ldots, M\}$, the following condition is satisfied:

$$\sum_{m \in \mathcal{T}} R_m \geq I(S_1, S_2, \ldots, S_M; W_{\mathcal{T}}|W_{\mathcal{T}^c}, Z) \quad (14)$$

and for which there exist functions $g_m(\cdot)$, for $m = 1, 2, \ldots, M$, such that

$$Ed_m(S_m, g_m(W_1, W_2, \ldots, W_M, Z)) \leq D_m, \quad (15)$$

for $m = 1, 2, \ldots, M$.

The proof of Theorem 2 is given in [2] for the case of two sources, but it is readily extended to the case of M sources.

The evaluation of Theorem 2 is an involved task in general: to get a valid converse bound, the minimization has to be carried out over *all* possible conditional distributions with alphabet sizes according to (12).

C. Rate-distortion Region

As pointed out earlier, the rate-distortion region for the general lossy source coding network illustrated in Figure 1 is unknown. In the results presented here, this is apparent in the difference between (7) and (13): the latter allows many more degrees of freedom. For the scenario without the side information Z, this point is elegantly explained in [5, p.223]. Rather than quantifying the difference between Theorems 1 and 2, the goal of this paper is to discuss a special case for which the two regions can indeed be shown to coincide.

IV. CONDITIONALLY INDEPENDENT SOURCES

In this section, we study a special case of the scenario introduced in Section I and illustrated in Figure 1: Suppose that the M sources are conditionally independent given the side information Z, i.e., their joint distribution (1) satisfies

$$p(s_1, s_2, \ldots, s_M, z) = p(z) \prod_{m=1}^{M} p(s_m|z). \quad (16)$$

Note that this special case is non-trivial: it does *not* imply that the sources $S_1, S_2, \ldots, S_M$ are independent.

A. Rate-distortion Region

It is intuitive that the situation is considerably simplified if the sources are conditionally independent given the side information. More precisely, it can be shown that the regions described by Theorems 1 and 2 coincide in this case. We summarize this in the following theorem:

Theorem 3: If $(S_1, S_2, \ldots, S_M)$ are conditionally independent given Z, then

$$\mathcal{R}_a(D_1, D_2, \ldots, D_M) = \mathcal{R}_o(D_1, D_2, \ldots, D_M) = \mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M), \quad (17)$$

thus establishing the rate-distortion region entirely.

The proof of Theorem 3 is given in [2], but we take this opportunity to emphasize that even under the assumption of conditional independence, conditions (7) and (13) *do not* coincide. Counterexamples can be constructed along the lines of the example in [5, p.223] by introducing a mixing random variable. Rather, the reason why the regions of Theorems 1 and 2 coincide is because the additional degrees of freedom in the choice of the auxiliary random variables $(W_1, W_2, \ldots, W_M)$ in Theorem 2 (i.e., in Eqn. (13)) over Theorem 1 (i.e., Eqn. (7)) leave the involved mutual information and distortion functionals unaffected.

B. The Rate Loss for Conditionally Independent Sources (I): Separate vs. Joint Encoding

An interesting question in distributed source coding is that of the loss due to the fact that compression is distributed rather than joint. The particular interest in this question is due to the famous result of Slepian and Wolf [6], showing that in the case of *lossless* compression, separate encoding of two (correlated) sources requires the same (sum) rate as joint encoding.

In this section, we study the rate loss due to the fact that the sources have to be encoded separately (Figure 1), rather than jointly (Figure 2). Since the rate region

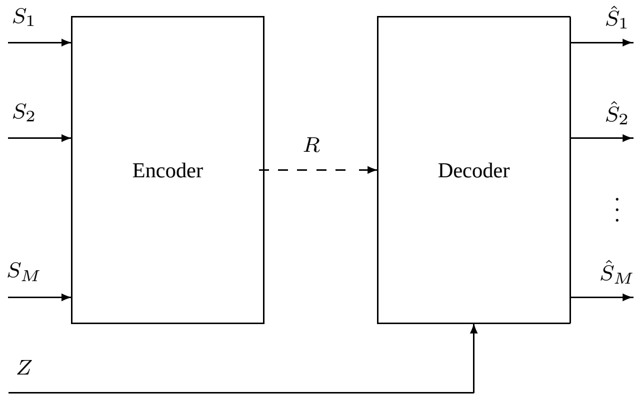

Fig. 2. Joint compression of all sources, with side information at the decoder.

$\mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1, D_2, \ldots, D_M)$ is not generally known, we cannot analyze the rate loss for arbitrary sources, but we can do it for the special case for which Theorem 3 provides full rate-distortion region. For that special case, we find that there is no rate loss, i.e., the minimum sum rate needed to achieve a distortion tuple $(D_1, D_2, \ldots, D_M)$ in the system of Figure 1 is the same as the minimum rate needed in the system of Figure 2. The following theorem can be proved.

Theorem 4: If $S_1, S_2, \ldots, S_M$ are conditionally independent given the side information Z, then

$$\min_{(R_1,R_2,\ldots,R_M) \in \mathcal{R}^{WZ}_{S_1,S_2,\ldots,S_M|Z}(D_1,D_2,\ldots,D_M)} \sum_{m=1}^{M} R_m = R^{WZ}_{(S_1,S_2,\ldots,S_M)|Z}(D_1, D_2, \ldots, D_M),$$

where $R^{WZ}_{(S_1,S_2,\ldots,S_M)|Z}(D_1, D_2, \ldots, D_M)$ denotes the Wyner-Ziv rate-distortion function for the vector source illustrated in Figure 2.

The proof of this theorem appears in [2].

An application of this theorem is the situation where $S_1, S_2, \ldots, S_M$ represents a Gaussian *vector* source that has to be compressed *jointly*. It is easy to see that in this case, there exists a linear and orthogonal transform of the vector $(S_1, S_2, \ldots, S_M)$ such that the transformed components are conditionally independent given the side information Z. This, and extensions thereof, is studied in further detail in the context of the distributed Karhunen-Loève transform, introduced in [7], [8], [9]. For this problem, Theorem 4 establishes that an optimal coding scheme is to first transform the vector $(S_1, S_2, \ldots, S_M)$ such as to obtain M conditionally independent components (given the side information), and then to apply a separate Wyner-Ziv code to each of the transformed components.

C. The Rate Loss for Conditionally Independent Sources (II): No Side Information at the Encoder

In this section, we study the rate loss due to the fact that the side information is not available at the encoders, i.e., we compare the rate-distortion regions for the source-coding problems of Figure 1 and 3, respectively. For the case of a single source, this problem has been studied in [1], [10]. In this paper, we show that their results can be used directly to assess the corresponding rate loss in the problem with multiple sources when the sources are conditionally independent given the side information. Theorem 1 can be used to make this precise. As asserted by Theorem 3, for the case of conditionally independent sources, both Theorem 1 and Theorem 2 describe the rate-distortion region. But the rate-distortion region described by Theorem 1 can be achieved by encoding each source with respect to the side information, but completely ignoring the remaining sources. Hence, the rate loss between the systems of Figure 1 and Figure 3 in the case of conditionally independent sources is characterized precisely by the rate losses for each of the sources separately.

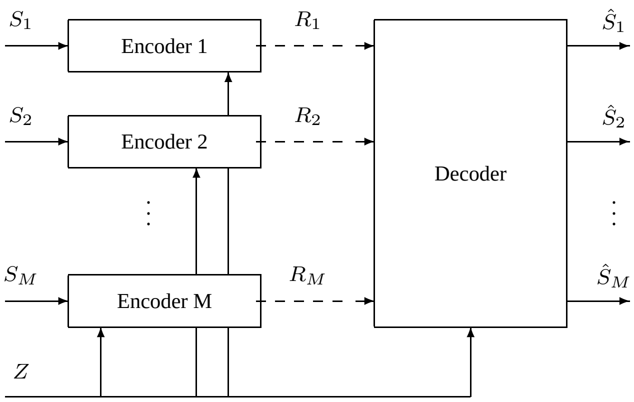

Fig. 3. Side information both at the encoder and at the decoder.

This argument can be used to prove for example the following statement.

Theorem 5: When $S_1, S_2, \ldots, S_M$ are conditionally independent given Z, the rate loss is zero if and only if the rate losses for the corresponding single-source Wyner-Ziv problems are zero.

For the cases where the rate loss for one or more of the sources is non-zero, it can at least be bounded using the results of Zamir [10].

V. APPLICATIONS

A. Bounds on Distributed Compression

Consider the source network topology illustrated in Figure 4: The problem is to find the set of rate tuples

$$(R_1, R_2, \ldots, R_M) \tag{18}$$

that permit to achieve a prespecified set of fixed target distortions

$$D_m = Ed_m(S_m, \hat{S}_m), \tag{19}$$

for $m = 1, 2, \ldots, M$. By analogy to Equation (5), we denote the corresponding rate-distortion region by

$$\mathcal{R}_{S_1, S_2, \ldots, S_M}(D_1, D_2, \ldots, D_M). \tag{20}$$

This problem is unsolved to date. The best known outer and inner bounds can be found e.g. in [5], and they do not coincide in general. Nevertheless, the answer to the problem of Figure 4 is known for a number of special cases: The case of perfect reconstruction has been solved by Slepian and Wolf in [6]. For $M = 2$, other special cases include the situation when $R_1 \to \infty$ (or $R_2 \to \infty$), which is the Wyner-Ziv problem [1], extended to $M \geq 2$ in this paper; the case when $D_1 = 0$ (or $D_2 = 0$), see [11]; and the case when one random variable has to be reconstructed perfectly [12]. The Gaussian case where only one source is reconstructed (but both have to be encoded, using rates R_1 and R_2, respectively) was treated in [13].

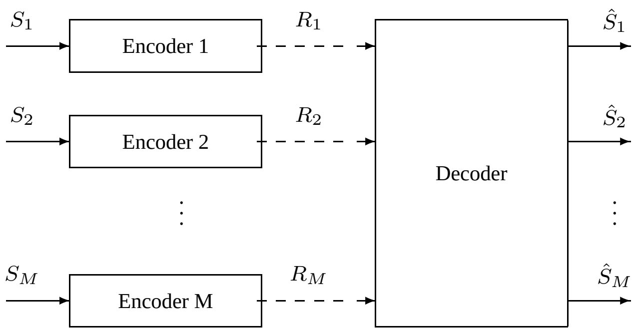

Fig. 4. The distributed compression problem.

The arguments developed here can be used to determine lower bounds to the achievable performance for the scenario of Figure 4: Clearly, providing the receiver with the exact values of a subset of the sources $(S_1, S_2, \ldots, S_M)$ can only improve the situation (i.e., at fixed distortions, it can only reduce the rates required for the remaining sources). This insight thus provides the following theorem:

Theorem 6: If the rate tuple $(R_1, R_2, \ldots, R_M)$ lies inside the achievable rate-distortion region for the problem illustrated in

Figure 4, i.e., if

$$\begin{aligned}&(R_1, R_2, \ldots, R_M)\\ &\in \quad \mathcal{R}_{S_1,S_2,\ldots,S_M}(D_1, D_2, \ldots, D_M)\end{aligned} \tag{21}$$

then, for each subset $\mathcal{T} \subseteq \{1, 2, \ldots, M\}$, the rate tuple with indices in $\mathcal{T}$, denoted by $(R_i)_{i\in\mathcal{T}}$, must lie in the Wyner-Ziv rate-distortion region of the sources with indices in $\mathcal{T}$, denoted by $S_\mathcal{T}$, given the remaining sources, denoted by $S_{\mathcal{T}^c}$, i.e.,

$$(R_i)_{i\in\mathcal{T}} \quad \in \quad \mathcal{R}^{WZ}_{S_\mathcal{T}|S_{\mathcal{T}^c}}(D_1, D_2, \ldots, D_M). \tag{22}$$

This theorem thus permits to find lower bounds to the desired rate-distortion region. However, to actually compute the lower bound in the general case is difficult: the multi-source Wyner-Ziv rate-distortion region is not known in general, i.e., as remarked above, Theorems 1 and 2 do not generally coincide. To still find a valid lower bound to the rate-distortion region (20), we can substitute the *lower bound* to the Wyner-Ziv rate-distortion region in Theorem 6. Unfortunately, as pointed out above, the actual evaluation of our lower bound to the Wyner-Ziv rate-distortion (Theorem 2) is generally a difficult problem, since all possible distributions for the auxiliary random variables have to be taken into account. Clearly, interesting situations arise whenever a subset of the involved sources is conditionally independent given the remaining sources: Here, the lower bound established in Theorem 6 can actually be evaluated.

B. Source-Channel Separation For Wyner-Ziv Networks

In Figure 1, suppose that instead of having abstract channels ("bit pipes") of fixed width $R_1, R_2, \ldots, R_M$ bits each, we have point-to-point noisy channels. Is the best possible distortion performance achieved by using capacity-achieving channel coding on each of the point-to-point noisy channels, effectively turning them into error-free bit pipes of rates $C_1, C_2, \ldots, C_M$, respectively, and using these bit pipes to convey the optimum source codewords for the Wyner-Ziv problem (assuming it is known)? In other words, is there a source-channel separation theorem for Wyner-Ziv source-channel networks?

In general network situations, it is not true that the optimum performance can be described by combining the rate-distortion region for the source network with the capacity-cost region for the channel network. A famous counterexample appears in [14, p.449], see [15] for further details.

However, it has been noted that for certain simple network topologies, such a source-channel separation theorem does indeed apply, see e.g. [16]. In extension of [16, Thm.1.10], it can be shown that if the sources $S_1, S_2, \ldots, S_M$ are conditionally independent given the side information Z, an optimal scheme is to encode the sources using the source code that leads to Theorem 1, and transmit the resulting source codewords across the M noisy point-to-point channels, using capacity-approaching codes.

VI. Conclusions

In this paper, we have extended the results of Wyner and Ziv [1] to the case of M sources. The coding scheme for this extension involves, beyond the methods developed in [1], also the fact that the sources $S_1, S_2, \ldots, S_M$ are dependent, i.e., it involves additional techniques developed in [5], [6]. Our considerations lead to an achievable rate-distortion region, and an outer bound to the full rate-distortion region, but these two do not generally coincide. Hence, our arguments do not establish the full rate-distortion region for the Wyner-Ziv problem with multiple source.

Thereafter, we identify a special scenario for which our arguments *do* establish the full rate-distortion region. This is the case when the sources $S_1, S_2, \ldots, S_M$ are conditionally independent given the side information Z. This special case is non-trivial in the sense that the sources may still be dependent. For that case, we determine the rate losses with respect to joint compression, and we identify a source-channel separation theorem.

Our results may have applications to sensor networks, where the sources $S_1, S_2, \ldots, S_M$ represent sensor readings, to be communicated to a central data collector. The data collector also has a sensor, with sensor reading Z. Extensions of our results therefore address the question of the best and worst cases of side information, leading to criteria for the optimal placement of the data collector.

Acknowledgment

Fruitful discussions with Dr. Gerhard Kramer (Bell Labs), Dr. Pier Luigi Dragotti (Imperial College), and Prof. Martin Vetterli (EPFL) are gratefully acknowledged.

References

[1] A. D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Transactions on Information Theory*, vol. IT–22, pp. 1–11, January 1976.

[2] M. Gastpar, "The Wyner-Ziv problem with multiple sources," *submitted to IEEE Transactions on Information Theory*, June 2002.

[3] M. Gastpar, "On the Wyner-Ziv problem with two sources," in *Proc IEEE Int Symp Info Theory*, (Yokohama, Japan), July 2003.

[4] A. D. Wyner, "The rate-distortion function for source coding with side information at the decoder-II: General sources," *Information and Control*, vol. 38, pp. 60–80, 1978.

[5] T. Berger, "Multiterminal source coding," *Lectures presented at CISM Summer School on the Information Theory Approach to Communications*, July 1977.

[6] D. Slepian and J. K. Wolf, "Noiseless coding of correlated information sources," *IEEE Transactions on Information Theory*, vol. IT–19, pp. 471–480, 1973.

[7] M. Gastpar, P. L. Dragotti, and M. Vetterli, "The distributed Karhunen-Loève transform," in *Proc 2002 International Workshop on Multimedia Signal Processing*, (St. Thomas, US Virgin Islands), December 2002.

[8] M. Gastpar, P. L. Dragotti, and M. Vetterli, "The distributed, partial, and conditional Karhunen-Loève transforms," in *Proc 2003 IEEE Data Compression Conference (DCC)*, (Snowbird, UT), March 2003.

[9] M. Gastpar, P. L. Dragotti, and M. Vetterli, "The distributed Karhunen-Loève transform," *EPFL I&C Technical Report IC/2003/12*, March 2003. To be submitted to IEEE Trans Info Theory, Nov. 2003.

[10] R. Zamir, "The rate loss in the Wyner-Ziv problem," *IEEE Transactions on Information Theory*, vol. IT–42, no. 4, 1996.

[11] A. H. Kaspi and T. Berger, "Rate-distortion for correlated source with partially separated encoders," *IEEE Transactions on Information Theory*, vol. IT–28, pp. 828–840, November 1982.

[12] T. Berger and R. W. Yeung, "Multiterminal source encoding with one distortion criterion," *IEEE Transactions on Information Theory*, vol. IT–35, pp. 228–236, March 1989.

[13] Y. Oohama, "Gaussian multiterminal source coding," *IEEE Transactions on Information Theory*, vol. IT–43, pp. 1912–1923, November 1997.

[14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[15] M. Gastpar, "On source-channel communication in networks," in *DIMACS Workshop on Network Information Theory* (P. Gupta, G. Kramer, and A. van Wijngarden, eds.), AMS DIMACS, March 2003.

[16] M. Gastpar, *To Code Or Not To Code*. PhD thesis, Ecole Polytechnique Fédérale (EPFL), Lausanne, Switzerland, 2002.