Exploiting Spatial Correlation in Pixel-Domain Distributed Image Compression

David Varodayan, Anne Aaron, and Bernd Girod*

Information Systems Laboratory, Department of Electrical Engineering
Stanford University, Stanford, CA 94305
{varodayan, amaaron, bgirod}@stanford.edu

Abstract. Contrary to convention, we construct distributed image compression codecs that operate in the pixel-domain, yet exploit spatial correlation at the decoder only. For lossless compression of binary images of text, we propose two novel decoders: one assumes the image to be a one-dimensional stationary Markov process, and the other assumes it to be a two-dimensional stationary Markov random field. We demonstrate that these decoders enable compression approaching their respective Slepian-Wolf limits and they perform better than the baseline pixel-domain decoder by factors of at least 2 and 4, respectively.

Index Terms Slepian-Wolf coding, distributed source coding, Low-Density Parity-Check codes, Baum-Welch algorithm

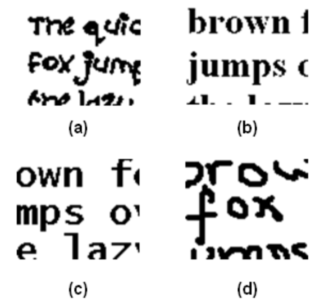

Fig. 1. Binary images of text (72 by 88 pixels)

1 INTRODUCTION

Conventional practice of distributed image compression (just like traditional image compression) suggests that performance is better if the encoder applies a decorrelating transform to the data than if it does not; consider implementations of distributed image and video codecs such as [1]. In this paper, we demonstrate that pixel-domain distributed coders can compress images without loss at rates very close to the theoretical limits by shifting the decorrelation task to the decoder. As a toy example, we consider binary text images as sources, such as those shown in Fig. 1. Our novel contributions are decoding algorithms that use spatial models of these images to aid their lossless recovery.

We begin with a version of the original formulation of distributed source coding and modify it to describe our problem. A memoryless finite-alphabet source X is to be transmitted without loss using the least average number of bits. Statistically dependent side information Y (not necessarily discrete) is available at the decoder only. The encoder must therefore compress X in the absence of Y, whereas the decoder uses Y to aid the recovery of X. Slepian and Wolf proved in 1973 that lossless compression is achievable at rates $R \geq H(X|Y)$, the conditional entropy of X given Y, for X and Y finite [2]. The difference in our setup is that the source has memory and is better represented as the random vector $\mathbf{X}$. We view the statistically dependent side information as the random vector $\mathbf{Y}$ since it is observed at the decoder in its entirety. The achievable rate region is given by $R \geq H(\mathbf{X}|\mathbf{Y})$, the conditional entropy rate of $\mathbf{X}$ given $\mathbf{Y}$ [3].

2 SYSTEM DESCRIPTION

The compression performance of the systems we discuss depends on the spatial model for $\mathbf{X}$ assumed at the decoder. We consider three models of increasing sophistication: independent identically distributed (i.i.d.) bits, one-dimensional (1D) stationary Markov process and two-dimensional (2D) stationary Markov random field.

2.1 Baseline i.i.d. decoding

The baseline system assumes $\mathbf{X}$ consists of i.i.d. bits. Our implementation, depicted in Fig. 2, follows the work by Liveris *et al.* [4]. The encoder sends the decoder the syndrome S of $\mathbf{X}$ with respect to a Low-Density Parity-Check (LDPC) code [5]. The decoder first calculates probability distributions $P(X_k|\mathbf{Y})$ on each of the bits of $\mathbf{X}$ given the observation of $\mathbf{Y}$. Then, since it assumes an i.i.d.

* This work is supported in part by the Max Planck Center for Visual Computing and Communication

model, the decoder recovers $\mathbf{X}$ by combining these distributions with the syndrome S by means of belief propagation iterations of LDPC decoding.

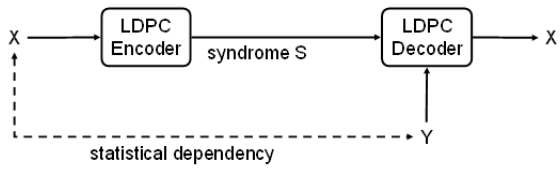

Fig. 2. The baseline system that models $\mathbf{X}$ as i.i.d. bits

2.2 1D Markov process decoding

Our first novel system, shown in Fig. 3, models the raster scan of $\mathbf{X}$ as a stationary first-order Markov process. Its encoder has the same lightweight structure as the baseline encoder. As before, the decoder calculates the distributions $P(X_k|\mathbf{Y})$ for each bit of $\mathbf{X}$. Since $\mathbf{X}$ is no longer assumed to be i.i.d., the decoder does not proceed with LDPC decoding iterations only. Instead, alternating iterations of the Baum-Welch algorithm [6] (without reestimation) and LDPC decoding exchange information until convergence.

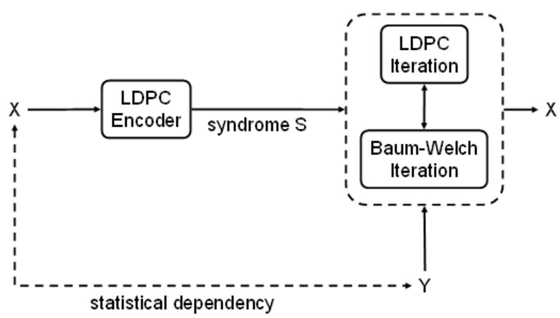

Fig. 3. The system that models the raster scan of $\mathbf{X}$ as a one-dimensional stationary Markov process

A Baum-Welch iteration, depicted in Fig. 4, applies the Markov model of $\mathbf{X}$ to create optimal extrinsic information $e_k(s)$ about each bit X_k. The Baum-Welch iteration receives as input the intrinsic information $i_k(s)$ about each bit X_k from the preceding LDPC decoding iteration as well as the side information $\mathbf{Y}$. The 2-state Markov model is parameterized by the transition probabilities p_{ij}. Since first-order Markovity guarantees that a bit of $\mathbf{X}$ is independent of all others given its immediate raster scan neighbors, the Baum-Welch iteration runs a forward pass and a backward pass through a 2-state trellis representing $\mathbf{X}$. The forward propagation is given by the recursion

$$\alpha_k(s_j) \propto \alpha_{k-1}(s_0)p_{0j} + \alpha_{k-1}(s_1)p_{1j}. \tag{1}$$

The backward pass on $\beta_k(s)$ is similar due to the reversibility of the stationary Markov process. Finally, the extrinsic information is computed as

$$e_k(s) \propto \alpha_k(s)\beta_k(s), \tag{2}$$

normalized so that $e_k(s_0) + e_k(s_1) = 1$.

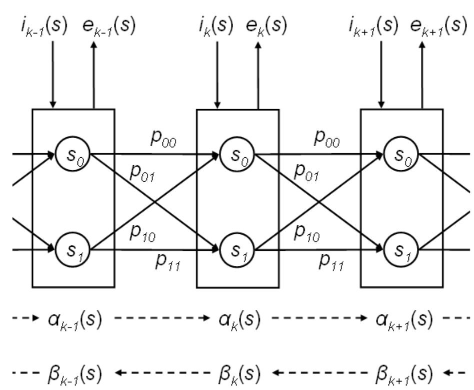

Fig. 4. A Baum-Welch iteration for the one-dimensional decoder

Based on this output of the Baum-Welch iteration, the side information $\mathbf{Y}$ and the syndrome S, the subsequent LDPC iteration generates its own extrinsic information about $\mathbf{X}$. In this way, the component decoders exchange information from iteration to iteration, until $\mathbf{X}$ is recovered. Similar one-dimensional decoders have been used to address related problems in source coding [7], channel coding [8] and encryption [9].

2.3 2D Markov random field decoding

Our second system extends the modeling to two dimensions, treating $\mathbf{X}$ as a stationary Markov random field. Under this assumption, a pixel of $\mathbf{X}$ is independent of all others given its four immediate neighbors in the image. Therefore, directly extending the 1D decoder of Section 2.2 would require replacing the first-order Markov model with a wth-order Markov model, where w is the width of the image. Consequently, each Baum-Welch iteration would operate on a trellis of 2^w states, which is unacceptable in terms of computation.

Instead, we propose the system shown in Fig. 5 that employs a modified suboptimal Baum-Welch iteration at the decoder of roughly the same complexity as the Baum-Welch iteration of Section 2.2. Once again, the simple encoder structure is unchanged and the decoder begins by calculating the distributions $P(X_k|\mathbf{Y})$ for each bit of $\mathbf{X}$. Then decoding proceeds with alternating iterations of the modified Baum-Welch algorithm and LDPC decoding. The modified Baum-Welch forward recursion is given by

$$\alpha_k(s_j) \propto \sum_{h,i \in \{0,1\}} \alpha_{k-w}(s_h)\alpha_{k-1}(s_i)p_{hij}, \qquad (3)$$

where p_{hij} are the corresponding transition probabilities, and the backward recursion is similar. As before, the component decoders pass extrinsic information to each other from iteration to iteration, until $\mathbf{X}$ is recovered.

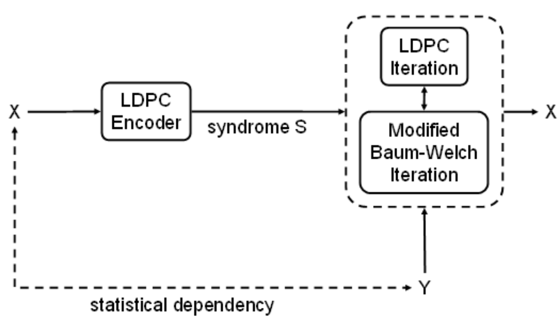

Fig. 5. The system that models $\mathbf{X}$ as a two-dimensional stationary Markov random field

3 SIMULATION RESULTS

We now compare the distributed compression performance and performance bounds of these three systems for different source images. We implement each system using sequences of rate-adaptive LDPC Accumulate codes [10] and apply rate control via feedback from the decoder. Every code is 6336 bits long and is regular of degree 3. Figs. 6, 7, 8 and 9 show these plots for the 6336 bit source images $\mathbf{X}$, shown in Figs. 1(a), (b), (c) and (d), respectively. The solid curves depict the average rate required over 50 trials for each system using random realizations of side information $\mathbf{Y}$, which are generated by the modulo 2 addition of i.i.d. binary noise of entropy $H(p)$ to $\mathbf{X}$. We denote the three average rate curves $R_{i.i.d.}$, R_{1D} and R_{2D} according to the decoder model for $\mathbf{X}$. The dashed curves represent the corresponding Slepian-Wolf performance bounds for each system, which we denote $H_{i.i.d.}(\mathbf{X}|\mathbf{Y})$, $H_{1D}(\mathbf{X}|\mathbf{Y})$ and $H_{2D}(\mathbf{X}|\mathbf{Y})$. Note that evaluating $H_{1D}(\mathbf{X}|\mathbf{Y})$ and $H_{2D}(\mathbf{X}|\mathbf{Y})$ is an open problem, so instead we calculate a good lower bound for $H_{1D}(\mathbf{X}|\mathbf{Y})$ and empirically estimate a good lower bound for $H_{2D}(\mathbf{X}|\mathbf{Y})$.

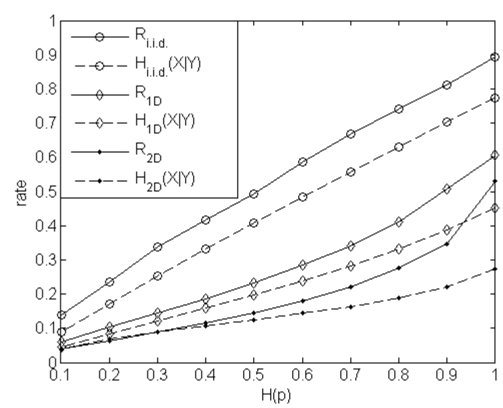

Fig. 6. Distributed compression performance and bounds when the source $\mathbf{X}$ is the image in Fig. 1(a)

These experimental results show that pixel-domain distributed source compression of text images can improve by factors of at least 2 and 4 by modeling the images as one-dimensional Markov processes and two-dimensional Markov random fields, respectively. Moreover, the systems endowed with spatial models are proportionally more efficient with respect to their Slepian-Wolf bounds than the baseline i.i.d. system, when the side information $\mathbf{Y}$ is very dependent on $\mathbf{X}$ (i.e. for $0.1 \leq H(p) \leq 0.4$).

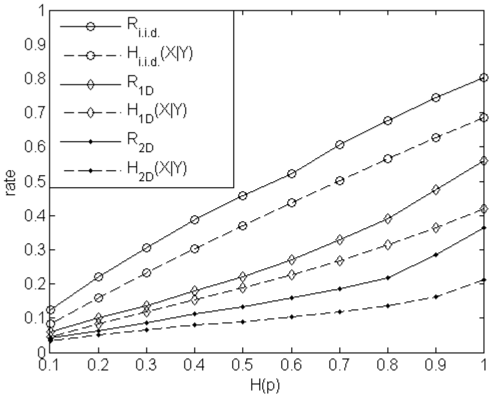

Fig. 7. Distributed compression performance and bounds when the source **X** is the image in Fig. 1(b)

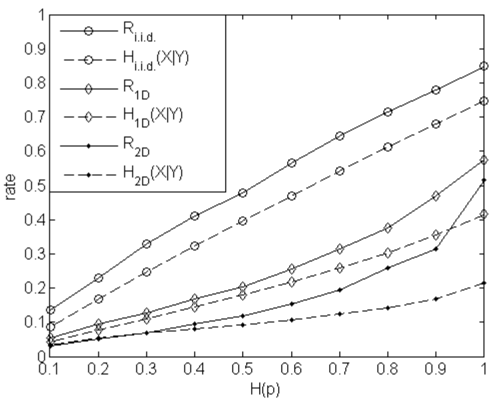

Fig. 8. Distributed compression performance and bounds when the source **X** is the image in Fig. 1(c)

4 CONCLUSION

We develop novel pixel-domain distributed source compression algorithms for binary text images. The proposed encoders are lightweight and the decoders model the images as either one-dimensional stationary Markov processes or two-dimensional stationary Markov random fields. These two systems compress the images without loss with at most half and quarter the rate, respectively, of the baseline pixel-domain system that does not model the images spatially. In fact, when the side information is good, the proposed systems are proportionally more efficient with respect to their bounds than their baseline counterpart.

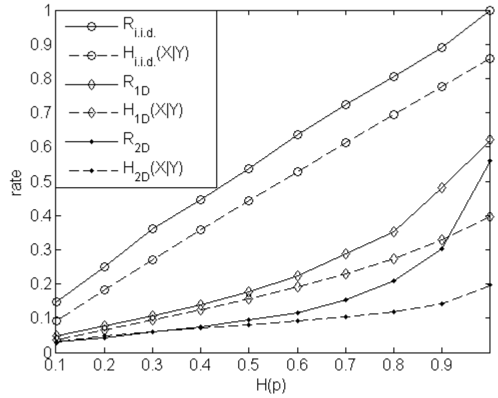

Fig. 9. Distributed compression performance and bounds when the source **X** is the image in Fig. 1(d)

References

1. Aaron, A., Rane, S., Setton, E., Girod, B.: Transform-domain Wyner-Ziv codec for video. In: SPIE Visual Communications and Image Processing Conf., San Jose, CA (2004)
2. Slepian, D., Wolf, J.K.: Noiseless coding of correlated information sources. IEEE Trans. Inform. Theory **19** (1973) 471–480
3. Cover, T.: A proof of the data compression theorem of Slepian and Wolf for ergodic sources. IEEE Trans. Inform. Theory **21** (1975) 226–228
4. Liveris, A., Xiong, Z., Georghiades, C.: Compression of binary sources with side information at the decoder using LDPC codes. IEEE Commun. Lett. **6** (2002) 440–442
5. Gallager, R.G.: Low-density parity-check codes. Cambridge MA: MIT Press (1963)
6. Baum, L.E., Petrie, T., Soules, G., Weiss, N.: A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. Ann. Math. Stat. **41** (1970) 164–171
7. García-Frías, J., Zhong, W.: LDPC codes for compression of multi-terminal sources with hidden Markov correlation. IEEE Commun. Lett. **7** (2003) 115–117
8. García-Frías, J.: Decoding of low-density parity check codes over finite-state binary Markov channels. IEEE Trans. Commun. **52** (2004) 1840–1843
9. Schonberg, D., Draper, S.C., Ramchandran, K.: On blind compression of encrypted correlated data approaching the source entropy rate. In: Proc. Allerton Conf. Commun., Contr. and Comput., Allerton, IL (2005)
10. Varodayan, D., Aaron, A., Girod, B.: Rate-adaptive distributed source coding using low-density parity-check codes. In: Proc. Asilomar Conf. on Signals, Syst., Comput., Pacific Grove, CA (2005)