Stealing Bits from a Quantized Source

Aaron S. Cohen Stark C. Draper Emin Martinian Gregory W. Wornell
Massachusetts Institute of Technology, Cambridge, MA 02139
{acohen,scd,emin,gww}@allegro.mit.edu

Abstract — **We analyze the efficiency of source requantization to reduce rate when starting with an arbitrary rate-distortion achieving codebook. In the quadratic-Gaussian case, it is possible to get within 0.5 bits/sample of the rate-distortion bound, i.e., all good codebooks automatically have close to successive-refinement structure. This same performance is achievable when rate is stolen by embedding bits in the source reconstruction.**

I. Introduction

Consider the scenario depicted in Fig. 1 in which a source is originally encoded at rate R_{orig} with distortion $d_0 = E\left[D(x^n, \hat{x}^n)/n\right]$. If a transcoder subsequently drops the encoding to a residual rate $R_{\text{res}} < R_{\text{orig}}$, the distortion in the ultimate reconstruction $\bar{x}^n$ is increased to $d > d_0$. Transcoding is efficient if R_{res} and d still lie on the rate distortion curve. If the source was originally encoded in a successively refinable manner (see, e.g., [3] and the references therein), this is possible by discarding least significant descriptions. We show, however, that near-efficiency is possible even without such a constraint on codebook design.

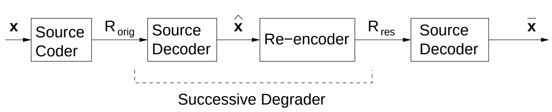

Figure 1: Successive Degradation.

II. Successive Degradation

For finite-alphabet i.i.d. sources with arbitrary distortion measures, we have that $R_{\text{res}}(d) = \inf I(\bar{x}; \hat{x})$ where the infimum is over all $\bar{x}$ such that $x \leftrightarrow \hat{x} \leftrightarrow \bar{x}$ and $E\left[D(x, \bar{x})\right] \leq d$.

In the case of an i.i.d. Gaussian source whose elements have variance σ_x^2 and a mean-square distortion (MSD) criterion we have the following stronger result: for an arbitrary original source code meeting the rate-distortion bound, the smallest distortion between $\bar{x}^n$ and x^n achievable by a transcoder is, in terms of the original distortion $d_0 = \sigma_x^2 2^{-2R_{\text{orig}}}$,

$$d(R_{\text{res}}) = d_0 + (\sigma_x^2 - d_0)2^{-2R_{\text{res}}}, \quad \text{for } R_{\text{res}} < R_{\text{orig}}, \qquad (1)$$

which is within 0.5 bits/sample of the rate-distortion function. Note the discontinuity at $R_{\text{res}} = R_{\text{orig}}$: to achieve a smaller distortion increment requires time-sharing; see Figure 2.

To show that (1) is achievable for any rate-distortion achieving source code, we first show that $E\left[\|\hat{x}^n\|^2/n\right] \approx \sigma_x^2 - d_0$ for large n, then analyze the distortion incurred by a Gaussian random re-encoder using [4, Thm. 3]. To establish the converse, we construct an original rate-distortion achieving source code that produces a reconstruction $\hat{x}^n$ that is indistinguishable from a noisy observation y^n of the source x^n

This work has been supported in part by NSF Grant No. CCR-0073520 and a grant from Microsoft Research.

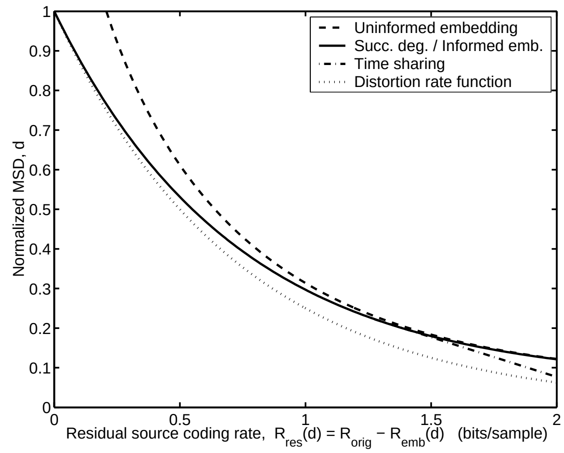

Figure 2: Quadratic-Gaussian bit stealing.

according to some $p_{y|x}$. Since the rate-distortion trade-off for this coding scheme is the same as for the problem of quantizing a source given a noisy observation, this codebook does not allow smaller distortions than (1) to be achieved.

III. Embedding in a Quantized Source

If the rate of a source is being reduced to accommodate an additional bit stream, one could alternatively reconstruct the source, embed the bit stream of rate R_{emb} into the reconstruction in the spirit of [2]—increasing the ultimate distortion to $d > d_0$—then requantize to the original codebook. The "effective" residual rate at which the source is encoded is then $R_{\text{res}}(d) = R_{\text{orig}} - R_{\text{emb}}(d)$. In contrast to successive degradation, the embedding approach does not require that the source decoder at the destination know $R_{\text{res}} \leq R_{\text{orig}}$.

For uninformed embedding, $R_{\text{res}}(d) = \inf I(\bar{x}; \hat{x})$ as in Section II, but with the addition constraint $p_{\bar{x}} = p_{\hat{x}}$, which results from the fact that the decoder is uninformed.

In the quadratic-Gaussian case, we have that $d(R_{\text{res}}) = 2\sigma_x^2 - d_0 - 2(\sigma_x^2 - d_0)\sqrt{1 - 2^{-2R_{\text{res}}}}$, which is close to (1) for small R_{emb} as depicted in Fig. 2. Furthermore, (1) can be achieved without requiring a new codebook to be generated at the point of embedding; it suffices to use the same codebook and have the source decoder apply simple post-decoding scaling [1].

References

[1] A. S. Cohen, S. C. Draper, E. Martinian, and G. W. Wornell, "Source requantization: Successive degradation and bit stealing," *Proc. DCC*, pp. 102–111, (Snowbird, UT), Mar. 2002.

[2] M. H. Costa, "Writing on dirty paper," *IEEE Trans. Inform. Theory*, v. 29, pp. 439–441, May 1983.

[3] W. H. Equitz and T. M. Cover, "Successive refinement of information," *IEEE Trans. Inform. Theory*, v. 37, pp. 269–275, Mar. 1991.

[4] A. Lapidoth, "On the role of mismatch in rate-distortion theory," *IEEE Trans. Inform. Theory*, v. 43, pp. 38–47, Jan. 1997.