Optimal Multiple Description and Multiresolution Scalar Quantizer Design

Michelle Effros
California Institute of Technology

***Abstract*— I present new algorithms for fixed-rate multiple description and multiresolution scalar quantizer design. The algorithms both run in time polynomial in the size of the source alphabet and guarantee globally optimal solutions. To the author's knowledge, these are the first globally optimal design algorithms for multiple description and multiresolution quantizers.**

I. Introduction

Multiple description and multiresolution source codes are data compression algorithms wherein a single source is described in multiple descriptions by a single encoder. In a multiple description code, each description can be decoded on its own or together with any subset of other descriptions; the goal of code design is to minimize the expected distortion of the reconstruction with respect to a known distribution on the description receipts or losses. In a multiresolution code, the descriptions are ordered, and the decoder can decode the first i descriptions for any value of i; the result is a family of nested descriptions of increasing reconstruction quality or "resolution," and the goal of code design is to minimize the expected distortion of the reconstruction with respect to a known distribution on the resolutions. This work focuses on fixed-rate multiple description and multiresolution scalar quantizer design.[1]

Prior work on fixed-rate multiple description scalar quantizer design includes both iterative descent algorithms [1] and approaches that rely on shortest path algorithms [2], [3].[2] Algorithms in the first family guarantee locally optimal code design and generalize easily from scalar to vector quantizers [5], [6]. Unfortunately there may exist multiple local optima, and the difference between the performance of the best and worst of these is unbounded; as a result, it is difficult to make strong, theoretical statements about the quality of solutions designed through iterative descent algorithms. Further, the run times for iterative descent algorithms depend on the number of iterations required to achieve convergence, which makes bounding algorithmic complexity difficult.

Algorithms in the second family yield the best code among all multiple description scalar quantizers that meet a certain convexity constraint. Unfortunately, optimal codes do not satisfy the convexity constraint in general, and the difference between the best code that meets the constraint and the truly optimal code can be arbitrarily large [2]. Shortest path algorithms run in time polynomial in the size of the source alphabet. The computational feasibility of these algorithms relies heavily on the restriction to scalar rather than vector quantizer design.

Prior work on fixed-rate multiresolution scalar quantizer design likewise includes both iterative descent algorithms [7], [8], [9] and shortest path algorithms [2], [10].[3] As in the corresponding multiple description codes, iterative descent techniques generalize easily to vector quantizer design, but they are susceptible to local optimality problems and their complexity is difficult bound. The shortest path algorithms rely heavily on both the restriction to scalar codes and the convexity constraint to obtain computational feasibility; the convexity constraint is not satisfied in general by optimal fixed-rate multiresolution scalar quantizers [2]. As a result, the performance difference between the codes designed using either approach and the corresponding optimal codes can again be arbitrarily large.

The algorithms described in this paper run in time polynomial in the size of the source alphabet and guarantee a globally optimal solution. The code design algorithm relies heavily on the restriction to scalar quantizers to obtain computational feasibility. The key innovation results from an observation about convexity discussed in Section III. The following section gives the definitions required for that discussion.

II. Preliminaries

Let $\mathcal{X} = \{x_1, \ldots, x_N\}$ denote a finite, real-valued source alphabet with $x_1 \le x_2 \le \cdots \le x_N$. Associated with each symbol x_n, $n \in \{1, \ldots, N\}$, is a symbol probability $p(n)$ with $p(n) > 0$ for all n and $\sum_{n=1}^{N} p(n) = 1$. The distortion between a source x and its reproduction x' is measured using the squared-error distortion measure $d(x, x') = (x - x')^2$. The rates for the M descriptions of a multiple description or multiresolution scalar quantizer are integers $R_1, \ldots, R_M$ greater than or equal to 1. The constants

$$K_m = 2^{R_m}, \quad m \in \{1, \ldots, M\}$$
$$K = \prod_{m=1}^{M} K_m$$

are useful in the discussion that follows.

[1] For fixed-rate coding, the given expected distortion performance criterion is equivalent to earlier Lagrangian performance measures.

[2] The multiple description scalar quantizer design algorithm in [2] originally appeared in [4]; [3] focuses on fast design under restrictive assumptions on the code parameters.

[3] The multiresolution scalar quantizer design algorithm in [2] also first appeared in [4].

A. Multiple Description Scalar Quantizers

An M-description scalar quantizer comprises a family of encoders $\alpha = \{\alpha_m\}_{m=1}^M$ and a family of decoders $\beta = \{\beta_b\}_{b\in\{0,1\}^M}$. Let

$$\mathcal{M} = \{1, \ldots, M\}$$

and for each $b = (b_1, \ldots, b_M) \in \{0,1\}^M$, let

$$\mathcal{M}(b) = \{m \in \mathcal{M} : b_m = 1\}.$$

Each encoder maps a source symbol to its M descriptions. Each decoder maps a distinct subset of received descriptions to its corresponding reproduction; decoder β_b uses the descriptions in $\mathcal{M}(b)$. Formally, the mappings are defined as follows

$$\begin{aligned} \alpha_m : &\quad \mathcal{X} \to \{0, \ldots, K_m - 1\} \quad \forall m \in \mathcal{M} \\ \beta_b : &\quad \prod_{m\in\mathcal{M}(b)} \{0, \ldots, K_m - 1\} \to \mathbb{R} \quad \forall b \in \{0,1\}^M. \end{aligned}$$

Finally, let

$$\alpha_b(x) = (\alpha_m(x) : m \in \mathcal{M}(b)).$$

Together, the M encoders define 2^M partitions of alphabet $\mathcal{X}$, here denoted by $\{\mathcal{P}_b\}_{b\in\{0,1\}^M}$. Partition $\mathcal{P}_b$ breaks the alphabet $\mathcal{X}$ into subsets such that symbols x and x' are in the same subset if and only if $\alpha_b(x) = \alpha_b(x')$. For each $c \in \mathcal{P}_b$, let $\alpha_b(c)$ denote that shared description; then $\alpha_b(c) \in \prod_{m\in\mathcal{M}(b)}\{0, \ldots, K_m - 1\}$ for all $c \in \mathcal{P}_b$. We call each $c \in \mathcal{P}_b$ a *codecell* from $\mathcal{P}_b$ and each reproduction $\beta_b(\alpha_b(c))$ a *codeword*.

Given a fixed rate vector $(R_1, \ldots, R_M)$, the goal of the multiple description scalar quantizer design algorithm is to minimize the expected distortion with respect to a known distribution $\{q(b)\}_{b\in\{0,1\}^M}$ on the 2^M possible patterns of received and lost descriptions. Thus an optimal code $(\alpha^\star, \beta^\star)$ satisfies $(\alpha^\star, \beta^\star) = \arg\min_{(\alpha,\beta)} J_D(p, q, \alpha, \beta)$, where

$$\begin{aligned} &J_D(p, q, \alpha, \beta) \\ &= \sum_{b\in\{0,1\}^M} q(b) \sum_{n=1}^{N} p(n) d(x_n, \beta_b(\alpha_b(x_n))) \\ &= \sum_{b\in\{0,1\}^M} q(b) \sum_{c\in\mathcal{P}_b} \sum_{x_n\in c} p(n) d(x_n, \beta_b(\alpha_b(c))). \end{aligned}$$

In light of these two representations, we can equivalently view multiple description scalar quantizer design as requiring the design of encoders $\{\alpha_m\}_{m=1}^M$ and decoders $\{\beta_b\}_{b\in\{0,1\}^M}$ or the design of partitions $\{\mathcal{P}_b\}_{b\in\{0,1\}^M}$ and decoders $\{\beta_b\}_{b\in\{0,1\}^M}$. In fact, all partitions $\{\mathcal{P}_b\}_{b\in\{0,1\}^M}$ can be derived from the M partitions $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$ associated with the individual descriptions, and thus optimal code design is equivalent to the optimal design of partitions $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$ and decoders $\{\beta_b\}_{b\in\{0,1\}^M}$. For any given collection of encoders or partitions optimal decoder design is trivial. Precisely, for the squared error distortion measure, the optimal decoder $\beta_b^\star$ satisfies

$$\begin{aligned} \beta_b^\star(\alpha_b(c)) &= \arg\min_{\mu\in\mathbb{R}} E[d(X, \mu)|\alpha_b(X) = \alpha_b(c)] \\ &= \arg\min_{\mu\in\mathbb{R}} E[d(X, \mu)|X \in c] = E[X|X \in c] \end{aligned}$$

for each $b \in \{0,1\}^M$ and $c \in \mathcal{P}_b$. As a result, the remainder of the discussion on fixed-rate multiple description scalar quantizer design focuses on the design of partitions $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$.

B. Multiresolution Scalar Quantizers

The definitions for M-resolution scalar quantizers are similar. Given a fixed vector $(R_1, \ldots, R_M)$ of incremental rates, an M-resolution scalar quantizer comprises a family of encoders $\{\alpha_m\}_{m=1}^M$ and a family of decoders $\{\beta_m\}_{m=1}^M$. Encoder α_m maps each source symbol to its mth description and decoder β_m maps each possible value for the first m descriptions to its corresponding reproduction. Formally,

$$\begin{aligned} \alpha_m : &\quad \mathcal{X} \to \{0, \ldots, K_m - 1\} \\ \beta_m : &\quad \prod_{i=1}^{m}\{0, \ldots, K_m - 1\} \to \mathbb{R} \end{aligned}$$

for each $m \in \{1, \ldots, M\}$. As with the multiple description problem, it is convenient to choose a notational shorthand to represent all possible complete descriptions that might be received by the decoder. In this case, that notation is α^m, where for each $m \in \{1, \ldots, M\}$,

$$\alpha^m(x) = (\alpha_1(x), \ldots, \alpha_m(x)).$$

The M encoders of a multiresolution code define M partitions denoted by $\{\mathcal{P}_m\}_{m=1}^M$. Partition $\mathcal{P}_m$ breaks the alphabet $\mathcal{X}$ into subsets such that x and x' are in the same subset if and only if $\alpha^m(x) = \alpha^m(x')$ – meaning that their first m descriptions are identical. Let $\alpha^m(c)$ denote this shared description; then $\alpha^m(c) \in \prod_{i=1}^m\{0, \ldots, K_i - 1\}$. Again, $c \in \mathcal{P}_m$ a *codecell* for $\mathcal{P}_m$ and its reproduction $\beta_m(\alpha^m(c))$ is a *codeword*.

Given a fixed vector of incremental rates $(R_1, \ldots, R_M)$, the goal of the multiresolution scalar quantizer design algorithm is to minimize the expected distortion with respect to a known distribution $\{q(m)\}_{m=1}^M$ on the M resolutions. In this case, an optimal code $(\alpha^\star, \beta^\star)$ satisfies $(\alpha^\star, \beta^\star) = \arg\min_{(\alpha,\beta)} J_R(p, q, \alpha, \beta)$, where

$$\begin{aligned} &J_R(p, q, \alpha, \beta) \\ &= \sum_{m=1}^{M} q(m) \sum_{n=1}^{N} p(n) d(x_n, \beta_m(\alpha^m(x_n))) \\ &= \sum_{m=1}^{M} q(m) \sum_{c\in\mathcal{P}_m} \sum_{x_n\in c} p(n) d(x_n, \beta_m(\alpha^m(c))). \end{aligned}$$

We can equivalently view multiresolution scalar quantizer design as either the design of encoders $\{\alpha_m\}_{m=1}^M$ and decoders $\{\beta_m\}_{m=1}^M$ or partitions $\{\mathcal{P}_m\}_{m=1}^M$ and decoders $\{\beta_m\}_{m=1}^M$.

For any given collection of encoders $\{\alpha_m\}_{m=1}^M$ or partitions $\{\mathcal{P}_m\}_{m=1}^M$, optimal decoder design is again trivial, giving

$$\begin{aligned}\beta_m^\star(\alpha^m(c)) &= \arg\min_{\mu\in\mathbb{R}} E[d(X,\mu)|\alpha^m(X)=\alpha^m(c)] \\ &= \arg\min_{\mu\in\mathbb{R}} E[d(X,\mu)|X\in c] = E[X|X\in c]\end{aligned}$$

for each $m\in\mathcal{M}$ and $c\in\mathcal{P}_m$ using the squared error distortion measure. Thus, the focus of the discussion going forward is again on the design of M partitions – in this case $\{\mathcal{P}_m\}_{m=1}^M$.

III. Codecell Convexity

The algorithms presented in [2], [10], [3] design optimal multiresolution and multiple description scalar quantizers subject to the constraint that all codecells of all of the partitions defined in Section II are convex. Thus, the *convex codecell constraint* requires that c is a convex subset of $\mathcal{X}$ for all $c\in\mathcal{P}_b$, $b\in\{0,1\}^M$, in multiple description coding and for all $c\in\mathcal{P}_m$, $m\in\mathcal{M}$, in multiresolution coding. Given the focus on *scalar* quantizers, partition $\mathcal{P}$ of $\mathcal{X}$ has convex codecells if and only if there exists an increasing sequence of thresholds $T=\{t_k\}_{k=0}^{|\mathcal{P}|}\subseteq\{x_0\}\cup\mathcal{X}$ with $t_0=x_0$, $t_{|\mathcal{P}|}=x_N$ and $\mathcal{P}=\{(t_{i-1},t_i]\}_{i=1}^{|\mathcal{P}|}$, where x_0 is an arbitrary real value less than x_1.

The restriction to the design of scalar quantizers with convex codecells is practically motivated. With this constraint, optimal code design is equivalent to designing the optimal threshold sequences T (and the corresponding codewords) for all relevant partitions. The fact that the number of distinct threshold sequences is polynomial in N enables fast algorithms for designing codes with convex codecells.

A. Convex Codecells in All Partitions Preclude Optimality

While there always exist optimal fixed-rate conventional scalar quantizers with convex codecells (see [11] for the squared-error distortion measure and [12, Lemma 6.2.1] and [2, Theorem 3] for more general distortion measures), the following theorem from [2] proves that the convex codecell constraint sometimes precludes optimality in multiresolution and multiple description code design. The theorem's proof is by construction of an example for which optimal performance cannot be obtained with convex codecells.

Theorem 1 ([2, Theorem 5]): Requiring codecell convexity in partitions $\mathcal{P}_{0^{m-1}10^{M-m}}$, $m\in\{1,\ldots,M\}$, of a fixed-rate multiple description scalar quantizer or partition $\mathcal{P}_1$ of a fixed-rate multiple description scalar quantizer precludes optimality for some finite-alphabet sources.

The following example, discussed briefly in [2, Section VII], lends some insight into the shortcomings of the convex codecell constraint in multiple description code design. The example shows that a multiple description scalar quantizer with convex codecells has a maximal number of codecells in $\mathcal{P}_{1^M}$ far smaller than the maximal number of codecells in that partition when $\mathcal{P}_{0^{m-1}10^{M-m}}$ are not constrained to be convex. The example also demonstrates that this restriction can cause severe performance degradation for some distributions q, especially at high rates.

Example 1: Consider an M description scalar quantizer with convex codecells. Since the codecells are convex, each partition $\mathcal{P}_{0^{m-1}10^{M-m}}$ is defined by K_m-1 threshold values $t_1,\ldots,t_{K_m-1}$. (Since t_0 and t_{K_m} are fixed, they do not play an active role in the definition.) Recall that partition $\mathcal{P}_{1^M}$ breaks $\mathcal{X}$ into subsets c such that $x,x'\in c$ implies that $\alpha_{1^M}(x)=\alpha_{1^M}(x')$, which implies that $\alpha_m(x)=\alpha_m(x')$ for all $m\in\mathcal{M}$. Thus each codecell in partition $\mathcal{P}_{1^M}$ is the intersection of some collection of M codecells $c_1,\ldots,c_M$ where $c_m\in\mathcal{P}_{0^{m-1}10^{M-m}}$ for each $m\in\mathcal{M}$. Further, since each $\mathcal{P}_{0^{m-1}10^{M-m}}$ is defined by K_m-1 threshold values, the intersection of these partitions is defined by at most $\sum_{m=1}^M(K_m-1)$ threshold values. Thus $\mathcal{P}_{1^M}$ contains at most $(\sum_{m=1}^M K_m)-M+1$ codecells.

The given bound on the maximal number of codecells in $\mathcal{P}_{1^M}$ is far smaller than we would expect based on the rates. In particular, decoder β_{1^M} receives all M descriptions at a total rate of $\sum_{m=1}^M R_m$. If all combinations of these descriptions could occur, then partition $\mathcal{P}_{1^M}$ would have $K=\prod_{m=1}^M K_m$ codecells. When M is large and $R_m>0$, $K>>(\sum_{m=1}^M K_m)-M+1$. For example, when $R_m=R$ for all m, $K=2^{MR}$ grows exponentially with M while $(\sum_{m=1}^M K_m)-M+1=M(2^R-1)+1$ grows only linearly with M. Thus, when all codecells are convex, many combinations of descriptions simply cannot occur, and multiple description scalar quantizers cannot take full advantage of the diversity available through their distinct descriptions.

While the optimal number of codecells in $\mathcal{P}_{1^M}$ is unknown in general, the given restriction can cause severe performance degradation for some description distributions q. For example, let $q(1^M)=1$ and $R=1$. Then the optimal fixed-rate multiple description scalar quantizer that violates the codecell convexity constraint requires $M=\log|\mathcal{X}|$ descriptions to achieve $J_D(p,q,\alpha,\beta)=0$ while the optimal fixed-rate multiple description scalar quantizer with convex codecells requires $M=|\mathcal{X}|-1$ descriptions to achieve the same performance. The difference between these two rates is unbounded for large M. □

The sub-optimality of non-convex codecells for multiresolution scalar quantization is more subtle. Here the partitions under consideration are not the partitions associated with each individual description, as in the multiple description case. Instead, for $m\in\{1,\ldots,M\}$, each $\mathcal{P}_m$ is the partition associated with receiving the first m descriptions. As a result, a multiresolution scalar quantizer with convex codecells in $\{\mathcal{P}_1,\ldots,\mathcal{P}_M\}$ can achieve the full $\prod_{i=1}^m K_i$ codecells in partition $\mathcal{P}_m$ for each $m\in\{1,\ldots,M\}$. Instead, the problem with codecell convexity in multiresolution scalar quantization arises when the optimal partition of data for one resolution causes suboptimal performance for another. In such cases, the need to compromise between resolutions sometimes results in a solution that uses non-convex codecells. The following example, an expansion on [2, Example 5], illustrates this point.

Example 2: Let $\mathcal{X}=\{20,40,60,140\}$ and $p(1)=p(2)=1/8$, $p(3)=p(4)=3/8$, and consider designing a fixed-rate 2-resolution scalar quantizer for the given source with

$R_1 = R_2 = 1$. The partition $\mathcal{P}_1$ has two codecells, while the partition $\mathcal{P}_2$ has four codecells that refine the codecells from $\mathcal{P}_1$.

When $q(1) = 1$, the goal of fixed-rate multiresolution scalar quantizer design is to minimize the expected distortion achieved by partition $\mathcal{P}_1$ using its corresponding optimal codewords. Since $q(1) = 1, q(2) = 0$ and this is a conventional scalar quantizer design problem. Therefore convex codecells in $\mathcal{P}_1$ are optimal. The optimal choice for partition $\mathcal{P}_1$ is $\mathcal{P}_1^\star = \{\{20, 40, 60\}, \{140\}\}$ (with codewords 48 and 140).

When $q(2) = 1$, the goal of fixed-rate multiresolution scalar quantizer design is to minimize the expected distortion achieved by partition $\mathcal{P}_2$ using its corresponding optimal codewords. This is another conventional scalar quantizer design problem, so convex codecells in $\mathcal{P}_2$ are optimal. The optimal partition $\mathcal{P}_2$ is $\mathcal{P}_2^\star = \{\{20\}, \{40\}, \{60\}, \{140\}\}$.

The key here is that $\mathcal{P}_2^\star$ for $(q(1), q(2)) = (0, 1)$ is not a rate-(1,1) refinement of $\mathcal{P}_1^\star$ for $(q(1), q(2)) = (1, 0)$; that is, we cannot obtain $\mathcal{P}_2^\star$ by dividing each codecell of $\mathcal{P}_1^\star$ into two subsets. As a result, for any $(q(1), q(2)) = (q, 1-q)$ with $q \notin \{0, 1\}$, the optimal multiresolution code design requires a compromise between the partitions $\mathcal{P}_1^\star$ and $\mathcal{P}_2^\star$ that give the best possible performance in resolutions 1 and 2, respectively. When $q = .01$, the partitions $\mathcal{P}_1 = \{\{20, 60\}, \{40, 140\}\}$ (with optimal codewords 50 and 115) and $\mathcal{P}_2 = \{\{20\}, \{60\}, \{40\}, \{140\}\}$ (with optimal codewords 20, 60, 40, and 140) have the required refinement property and yield the optimal performance

$$\begin{aligned} &J_R(p, q, \alpha, \beta) \\ &= .01 \left(\frac{1}{8}(20-50)^2 + \frac{3}{8}(60-50)^2 \right. \\ &\quad \left. + \frac{1}{8}(40-115)^2 + \frac{3}{8}(140-115)^2 \right) + .99 \cdot 0 \\ &= 10.875. \end{aligned}$$

The large value of $q(2)$ weights the distortion of resolution 2 heavily, which therefore favors a solution in which $\mathcal{P}_2$ has only a single symbol per codecell (and thus 0 distortion). This forces $\mathcal{P}_1$ to place two symbols in each codecell. Considering partitions that meet this constraint, $\mathcal{P}_1 = \{\{20, 60\}, \{40, 140\}\}$ (with optimal codewords 50 and 115) yields lower distortion than the convex codecell partition $\mathcal{P}_1' = \{\{20, 40\}, \{60, 140\}\}$ (with optimal codewords 30 and 100) since $\mathcal{P}_1$ achieves more accurate reproduction of the more probable symbols. □

B. Convex Codecells in Some Partitions are Optimal

While Theorem 1 proves that requiring *all* partitions of a multiple description or multiresolution scalar quantizer to have convex codecells precludes optimality for some sources, Theorem 2 demonstrates that requiring codecell convexity in only the finest partition ($\mathcal{P}_{1^M}$ in multiple description scalar quantization and $\mathcal{P}_M$ in multiresolution scalar quantization) has no adverse effects. That is, there always exists an optimal code of the desired type with convex codecells. This property is critical for the proposed code design. We therefore include its proof for completeness.

Theorem 2 ([2, Theorem 6]): For any description distribution q, there exists an optimal fixed-rate multiple description or multiresolution scalar quantizer with convex codecells.

Proof: Since the argument in [2, Theorem 6] focuses on the multiresolution case, this proof focuses on the multiple description case. The proofs are conceptually the same. Let the description distribution q be fixed. We show that given any fixed-rate multiple description scalar quantizer there exists another code (of the same rate and type) with convex codecells that does at least as well. This gives the desired result since it can be applied, for example, to any optimal code. The argument is by construction. The construction depends on the description distribution q but not on the source distribution p.

Let (α, β) be a fixed-rate multiple description scalar quantizer with rate vector $(R_1, \ldots, R_M)$. Let $\{\mathcal{P}\}_{b \in \{0,1\}^M}$ be the partitions defined by this code. For any $r = (r_1, \ldots, r_M) \in \prod_{m=1}^M \{0, \ldots, K_m - 1\}$, let $r_b = (r_m : m \in \mathcal{M}(b))$; then $\beta_b(r_b)$ gives the reproduction associated with description r when only the descriptions of b are received. The argument that follows constructs a new multiple description scalar quantizer (α', β') that uses the given decoder ($\beta' = \beta$) but defines a new encoder ($\alpha' \neq \alpha$). Let $\{\mathcal{P}_b'\}_{b \in \{0,1\}^M}$ be the partitions defined by encoder α'; α' is chosen to ensure that $\mathcal{P}_{1^M}'$ has convex codecells (although the other partitions may not) and achieves performance $J_D(p, q, \alpha', \beta') \leq J_D(p, q, \alpha, \beta)$ for every source distribution p on $\mathcal{X}$.

The following definitions are used in the construction of encoder α', whose description is given in terms of its corresponding partitions. For any $x \in \mathcal{X}$ and any $r \in \prod_{m=1}^M \{0, \ldots, K_m - 1\}$, let

$$j_D(q, x, r) = \sum_{b \in \{0,1\}^M} q(b) d(x, \beta_b(r_b)).$$

Then for each $r, s \in \prod_{m=1}^M \{0, \ldots, K_m - 1\}$, let

$$c''(r, s) = \begin{cases} \{x : j_D(q, x, r) < j_D(q, x, s)\} & \text{if } r > s \\ \{x : j_D(q, x, r) \leq j_D(q, x, s)\} & \text{if } r \leq s. \end{cases}$$

Finally, for each $r \in \prod_{m=1}^M \{0, \ldots, K_m - 1\}$ let

$$c'(1^M, r) = \cap_{s \in \prod_{m=1}^M \{0, \ldots, K_m - 1\}} c''(r, s)$$

and for each $b \in \{0,1\}^M$ define

$$\begin{aligned} c'(b, r_b) &= \cup_{s \in \prod_{m=1}^M \{0, \ldots, K_m - 1\} : s_b = r_b} c'(1^M, s) \\ \mathcal{P}_b' &= \{c'(b, r_b)\}_{r_b \in \prod_{m \in \mathcal{M}(b)} \{0, \ldots, K_m - 1\}}. \end{aligned}$$

For the squared error distortion measure, the difference $j_D(q, x, r) - j_D(q, x, s)$ is linear (the quadratic terms cancel), and thus each non-empty $c''(r, s)$ is a half line. Since any non-empty intersection of half lines is an interval, $\mathcal{P}_{1^M}'$ has convex codecells.

It remains to show that the newly designed partitions achieve good expected distortion. This follows since

$$\begin{aligned} & J_D(p,q,\alpha',\beta') \\ &= \sum_{b\in\{0,1\}^M} q(b) \sum_{c\in\mathcal{P}_b} \sum_{x_n\in c} p(n) d(x_n, \beta_b(\alpha_b(c))) \\ &= \sum_{x_n\in\mathcal{X}} p(n) j_D(q, x_n, \alpha'(x_n)) \\ &\le \sum_{x_n\in\mathcal{X}} p(n) j_D(q, x_n, \alpha(x_n)) \\ &= J_D(p,q,\alpha,\beta), \end{aligned}$$

where the inequality holds point-wise for every x_n since, by construction, α' maps each source symbol to the description that minimizes its expected distortion. ■

C. The Implications of Codecell Convexity

The codecell convexity assumption is used in prior code design algorithms to restrict the family of possible codes and thereby enable fast design algorithms. Since Theorem 1 and the discussion following demonstrate that there can be an arbitrarily large penalty associated with this restriction, the convexity assumption must be dropped in the proposed design algorithm. The number of partitions $\mathcal{P}$ of $\mathcal{X}$ with at most $K \le N$ convex codecells is

$$\sum_{k=1}^{K} \begin{pmatrix} N-1 \\ k-1 \end{pmatrix};$$

without this restriction, the corresponding number is $K^N/K!$ (we don't distinguish between distinct labellings of the same partition). While there are no known bounds for the optimal number of codecells in variable-rate sources codes, the number of codecells in each partition for a fixed-rate code is a simple function of the rate vector. Thus for fixed-rate coding, K (the maximal number of cells in a partition) is fixed while N (the alphabet or training set size) is allowed to grow without bound. With the convex codecell constraint, the number of partitions under consideration is polynomial in N while the number of unrestricted partitions grows exponentially in N. This size difference is a critical hurdle to fast, optimal code design.

Theorem 2 proves that while the coarser partitions of a multiple description or multiresolution scalar quantizer may require non-convex codecells, the finest partition never does; this turns out to be the key insight in overcoming the issue of code design complexity. Since every codecell of every partition in a multiple description or multiresolution scalar quantizer is a union of the codecells of the finest partition, the number of possible partitions that meet the constraint of Theorem 2 is far smaller than the number in the unrestricted class. If the finest partition has K codecells while the partition associated with description m ($\mathcal{P}_{0^{m-1}10^{M-m}}$ in multiple description scalar quantization or $\mathcal{P}_m$ in multiresolution scalar quantization) has K_m codecells, a loose upper bound on the number of possible values for this partition is

$$\begin{pmatrix} N-1 \\ K-1 \end{pmatrix} \frac{(K_m)^K}{K_m!},$$

which is again polynomial in N.

IV. Optimal Code Design

The proposed algorithm finds optimal multiple description and multiresolution scalar quantizers by efficiently comparing all alternatives.

A bit of background is required. For any $A \subseteq \mathcal{X}$, let

$$\begin{aligned} p(A) &= \Pr(X\in A) \\ \mu(A) &= E[X|X\in A] \\ d(A,\mu(A)) &= \sum_{n:x_n\in A} p(n) d(x_n,\mu(A)) \\ d(\mathcal{P}) &= \sum_{c\in\mathcal{P}} d(c,\mu(c)). \end{aligned}$$

The following properties of the squared error distortion measure and fixed-rate multiple description and multiresolution scalar quantizers are useful in the code design.

1) For any $A \subseteq \mathcal{X}$,
$$\mu(A) = \arg\min_{\mu\in\mathbb{R}} \sum_{n:x_n\in A} p(n) d(x_n,\mu).$$
2) For any partition $\{B, C\}$ of $A \subseteq \mathcal{X}$,
$$\begin{aligned} p(A) &= p(B)+p(C) \\ \mu(A) &= \frac{p(B)\mu(B)+p(C)\mu(C)}{p(A)} \\ d(A,\mu(A)) &= d(B,\mu(B)) + p(B) d(\mu(B),\mu(A)) \\ &\quad + d(C,\mu(C)) + p(C) d(\mu(C),\mu(A)) \end{aligned}$$
3) For any fixed-rate multiple description scalar quantizer with rate vector $(R_1,\ldots,R_M)$, there exists an optimal collection of partitions $\{\mathcal{P}_b\}_{b\in\{0,1\}^M}$ with
$$|\mathcal{P}_{0^{m-1}10^{M-m}}| = K_m \quad \forall\, m\in\{1,\ldots,M\}.$$
4) For any fixed-rate multiresolution scalar quantizer with incremental rate vector $(R_1,\ldots,R_M)$, there exists an optimal collection of partitions $\{\mathcal{P}_m\}_{m=1}^M$ with
$$|\mathcal{P}_1| \ge K_1.$$

Properties 1 and 2 are straight forward and well known. Property 3 follows since in a fixed-rate multiple description code splitting a codecell of $\mathcal{P}_{0^{m-1}10^{M-m}}$ decreases (or leaves unchanged) the distortion in all resolutions simultaneously and therefore decreases (or leaves unchanged) the expected distortion. Taking the intersection of the codecells of two partitions can only increase the number of codecells, so $|\mathcal{P}_b| \ge |\mathcal{P}_{0^{m-1}10^{M-m}}|$ for all $m\in\mathcal{M}(b)$. Property 4 follows since in a fixed-rate multiresolution code splitting a codecell of $\mathcal{P}_1$ decreases or leaves unchanged the distortion for all partitions and therefore decreases or leaves unchanged the expected distortion. Further, the codecells of $\mathcal{P}_m$ refine the codecells of $\mathcal{P}_{m-1}$, so $|\mathcal{P}_m| \ge |\mathcal{P}_{m-1}|$ for all $m \ge 2$.

For each interval $A \subseteq \mathcal{X}$, we calculate and store $p(A)$, $\mu(A)$, and $d(A, \mu(A))$. Each such calculation can be done in constant time using property 1 provided that the calculations are carefully ordered. Specifically, we work with intervals of increasing size and calculate the distortion for each interval of size greater than 1 using the distortions of two intervals of roughly half that size.

From single intervals, we move to unions of intervals, now calculating $p(A)$, $\mu(A)$, and $d(A, \mu(A))$ for each such union A. These calculations represent the distortions of individual codecells. Again, careful ordering of the operations allows us to perform each calculation in constant time, building each value from two previously calculated terms. We restrict our attention to codecells comprised of at most $K/2$ non-consecutive intervals, where intervals $(a, b]$ and $(c, d]$ are consecutive if $b = c$. (A pair of consecutive intervals can be described as a single interval.) This choice is sufficient to achieve optimality by the following argument.

Relying on Theorem 2, we need consider only codecells that can occur in a code for which the finest partition is comprised of at most K codecells, all of which are convex. Recall that the codecell of any other partition is the union of codecells from the finest partition. Notice further that a codecell comprised of k non-consecutive intervals defines in total at least $2k - 1$ intervals – the intervals in the codecell plus the intervals between those intervals (and possibly an interval before the first interval in the codecell or an interval after the final interval). Therefore any codecell in a code with K convex codecells in its finest partition has at most $K/2$ non-consecutive intervals in any single codecell.

From codecells, we move to partitions, calculating the expected distortion $d(\mathcal{P})$ for each partition $\mathcal{P}$ with $2 \leq |\mathcal{P}| \leq K$. Again, careful ordering of the operations allows us to perform each calculation in constant time. In this case, we work from smaller to larger partitions Partitions of size 2 require the addition of two codecell distortions. Partitions of size greater than 2 require 3 additions to subtract off one codecell's distortion and add in the distortions of two codecells that refine the codecell that was removed.[4]

Finally, we calculate the expected performance of each code. Each code is described by M partitions. For multiple description scalar quantization, the partitions required are $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$, which uniquely determine $\{\mathcal{P}_b\}_{b \in \{0,1\}^M}$. The expected distortion for $\{\mathcal{P}_b\}_{b \in \{0,1\}^M}$ can be calculated by the weighted summation

$$J_D(p, q, \alpha, \beta) = \sum_{b \in \{0,1\}^M} q(b) d(\mathcal{P}_b)$$

of previously calculated distortions. For multiresolution scalar quantization, the partitions required are $\{\mathcal{P}_m\}_{m=1}^M$. The expected distortion is again a weighted sum of previously calculated distortions, here given by

$$J_R(p, q, \alpha, \beta) = \sum_{m=1}^{M} q(m) d(\mathcal{P}_m).$$

Theorem 3: The proposed algorithm finds optimal multiple description scalar quantizers in time $O(N^K + 2^M (NK)^{K-1})$. The proposed algorithm finds optimal multiresolution scalar quantizers in time $O(N^K + M(NK)^{K-1})$.

Proof: The algorithm compares all codes with no more than K codecells in its finest partition. This guarantees optimality by Theorem 2. The complexity of the proposed code is broken into two components. The first is the cost of calculating distortions for individual partitions; this cost is identical for multiple description and multiresolution codes. The second is the portion that is specific to a each type of code design.

The number of possible codecells with at least one and no more than $K/2$ segments is

$$\sum_{k=1}^{K/2} \binom{N+1}{2k}.$$

To find a simple upper bound on the number of partitions, note that there are

$$\binom{N-1}{K-1}.$$

ways to partition $\mathcal{X}$ into K intervals and $K^K/K!$ ways to assign those intervals to at most K codecells. (The denominator arises because two partitions that use the same codecells but assign different indices to those codecells are identical for our purposes.) Since the order of operations allows us to calculate each interval, codecell, and partition distortion in constant time, we can bound the total time required for calculating the partition distortions by

$$\sum_{k=1}^{K/2} \binom{N+1}{2k} + \binom{N-1}{K-1} \left(\frac{K^K}{K!}\right).$$

We bound the number of partition combinations to be considered in multiple description coding using the number of ways to divide $\mathcal{X}$ into K segments and the number of ways to choose each partition assuming that those K segments are fixed. The resulting bound on the number of choices for $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$ is

$$\binom{N-1}{K-1} \prod_{m=1}^{M} \left(\frac{K^{K_m}}{K_m!}\right).$$

The full collection of partitions $\{\mathcal{P}_b\}_{b \in \{0,1\}^M}$ is fully determined by $\{\mathcal{P}_{0^{m-1}10^{M-m}}\}_{m=1}^M$. Given the partition distortion calculations, calculating the resulting value $J_D(p, q, \alpha, \beta)$ requires the summation of 2^M previously calculated values. The resulting bound on the complexity of the full code design is

$$\sum_{k=1}^{K/2} \binom{N+1}{2k} + \binom{N-1}{K-1} \left(\frac{K^K}{K!}\right) + 2^M \binom{N-1}{K-1} \prod_{m=1}^{M} \left(\frac{K^{K_m}}{K_m!}\right),$$

[4]From properties 3 and 4, the minimal number of codecells per partition is at least $\min_{m \in \mathcal{M}} K_m$ for multiple description scalar quantization and K_1 for multiresolution scalar quantization. We neglect this observation and consider partitions with as few as 2 codecells in order to keep the algorithmic description and our bound on its complexity as simple as possible.

which give the desired result.

We bound the number of partition combinations to be considered in multiresolution coding using the number of ways to divide $\mathcal{X}$ into the K intervals of $\mathcal{P}_M$ and the number of ways to choose $\mathcal{P}_{m-1}$ given $\mathcal{P}_m$ for m decreasing from M to 2. For each $m \in \{1, \ldots, M\}$, let $K'_m = \prod_{i=1}^{m} K_i$. Then the resulting bound on the number of choices for $\{\mathcal{P}_m\}_{m=1}^{M}$ is

$$\left(\begin{array}{c} N-1 \\ K-1 \end{array} \right) \prod_{m=2}^{M} \left(\begin{array}{cccc} & & K'_m & \\ K'_{m-1} & K'_{m-1} & \cdots & K'_{m-1} \end{array} \right).$$

Given the partition distortion calculations, calculating the resulting value $J_R(p, q, \alpha, \beta)$ requires the summation of M previously calculated values. The resulting bound on the complexity of the full code design is

$$\sum_{k=1}^{K/2} \left(\begin{array}{c} N+1 \\ 2k \end{array} \right) + \left(\begin{array}{c} N-1 \\ K-1 \end{array} \right) \left(\frac{K^K}{K!} \right)$$
$$+M \left(\begin{array}{c} N-1 \\ K-1 \end{array} \right) \prod_{m=2}^{M} \left(\frac{K'_m!}{(K_{m-1}!)^{K_m}} \right),$$

which give the desired result.

■

V. Conclusions

The proposed algorithm finds optimal multiple description scalar quantizers and optimal multiresolution scalar quantizers in time polynomial in the size of the source alphabet N. The key insight used to simultaneously achieve computational feasibility and optimal code design is the observation that the codecell convexity is sufficient to achieve optimality in the finest partition of an optimal multiple description or multiresolution scalar quantizer.

References

[1] V. A. Vaishampayan. Design of multiple description scalar quantizers. *IEEE Transactions on Information Theory*, 39(3):821–834, May 1993.

[2] D. Muresan and M. Effros. Quantization as histogram segmentation: optimal scalar quantizer design in network systems. *IEEE Transactions on Information Theory*, 54(1):344–366, January 2008.

[3] S. Dumitrescu and X. Wu. Lagrangian optimization of two-description scalar quantizer design. *IEEE Transactions on Information Theory*, 53(11):3990–4012, nov 2007.

[4] D. Muresan and M. Effros. Quantization as histogram segmentation: globally optimal scalar quantizer design in network systems. In *Proceedings of the Data Compression Conference*, pages 302–311, Snowbird, Utah, April 2002.

[5] M. Fleming and M. Effros. Generalized multiple description vector quantization. In *Proceedings of the Data Compression Conference*, pages 3–12, Snowbird, UT, March 1999. IEEE.

[6] M. Fleming, Q. Zhao, and M. Effros. Network vector quantization. *IEEE Transactions on Information Theory*, 50(8):1584–1604, August 2004.

[7] P. A. Chou. Optimal partitioning for classification and regression trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4), April 1991.

[8] H. Brunk and N. Farvardin. Fixed-rate successively refinable scalar quantizers. In *Proceedings of the Data Compression Conference*, pages 250–259, Snowbird, UT, April 1996. IEEE.

[9] M. Effros and D. Dugatkin. Multiresolution vector quantization. *IEEE Transactions on Information Theory*, 50(12):3130–3145, December 2004.

[10] S. Dumitrescu and X. Wu. Algorithms for optimal multi-resolution quantization. *Journal on Algorithms*, 50:1–22, 2004.

[11] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–136, March 1982. Previously an unpublished Bell Laboratories Technical Note (1957).

[12] A. Gersho and R. M. Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston, 1992.