RJ 10249 (A0207-040) July 29, 2002
Computer Science

IBM Research Report

Error Diffusion Method for Resampling in the Transform Domain

Tomasz Nowicki[1], J. Q. Trelewicz[2], Marco Martens[1], Timothy Trenary[3]

[1]IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

[2]IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099

[3]IBM Printing Systems Division
6300 Diagonal Highway
Boulder, CO 80301

Error Diffusion Method for Resampling in the Transform Domain

Tomasz Nowicki, J. Q. Trelewicz*, Marco Martens, Timothy Trenary

Abstract— **The algorithm presented here addresses the problem of resampling images or signals with computational efficiency in real time, with an algorithm capable of providing arbitrary resampling ratios that are determined on-the-fly. The algorithm uses a set of resampling algorithms with predetermined, fixed resampling ratios, utilizing Error Diffusion (ED) to obtain an arbitrary resampling factor between two fixed values. Specifically, optimized implementations for the collection of fixed resampling values may be pre-calculated offline, and our algorithm can achieve resampling ratios that lie between any two of these fixed resampling ratios. Because the fixed resampling values can be pre-calculated based on any number of prior art fast resampling algorithms, resampling can be performed either in the transform domain or in the sample domain. The algorithm is appropriate for implementation in software or hardware.**

Keywords—

I. Introduction

In real-time, high-throughput, image processing systems, resampling compressed images is a common operation. For example,

- A server provides a wavelet-compressed image at one resolution, but a client Personal Digital Assistant (PDA) needs the image at a different resolution. This PDA resolution is only known at the time of download request, and the resampling must be performed efficiently, either at the server or at the PDA. Thus, resampling must be performed in real time. If the resampling is performed at the PDA, limited processing resource may be available for the operation because of space/power/weight/cost concerns, potentially making the task even more challenging.
- A high-speed production printer receives a customer job to print in real time, in an encoded print data stream. A JPEG image, embedded in the data stream, must be resampled prior to printing, and the resampling factor is only known at the time of decoding of the data stream.

The resampling factor $b > 0$ is defined such that the ratio of the number of samples in the d-dimensional signal ($d \in \mathbb{N}$) before and after the resampling is b^d.

Much work has been done on high quality image resampling. For example, in [1], the authors present a resampling algorithm for image zooming, preserving edge locations and contrast and avoiding aliasing. In [2], the authors discuss a method for B-spline image resampling that has better computational complexity than some other high-quality image resampling techniques. However, this body of literature focuses primarily on image quality, for applications such as scientific and medical imaging, where preservation of detail content is critical. In the commercial display applications that we target with our algorithm, these high-quality resampling techniques still require too much computation to be practical in real time.

To address computational complexity, the resampling matrices may be calculated offline and loaded into the image processing system. However, such an approach is infeasible when the b is only known at the time of image access, and access must be provided in real time. We describe in this paper an algorithm to address this challenge, where algorithms for fixed resampling factors are optimized in advance, and used to approximate the desired b value.

Transform coding is the name given to a wide family of techniques for data coding, in which each block of data to be coded is transformed by some mathematical function prior to further processing. A block of data may be a part of a data object being coded, or may be the entire object. The data generally represent some phenomenon, for example an image, a video clip, or an audio segment. The transform function is usually chosen to reflect some quality of the phenomenon being coded; for example, in coding of still images and motion pictures, the Fourier transform or DCT can be used to analyze the data into frequency terms or coefficients. Often, the objective of transforming the data is that the transformed data may be represented more compactly with high quality. For example, if the information is concentrated into a few frequency coefficients, the other coefficients may be represented in a lossy manner. Then the end-to-end transform is not invertible, even though A is.

For example, the JPEG standard [3] allows the interchange of images between diverse applications and opens up the capability to provide digital continuous-tone color images in multimedia applications. For this reason, our examples utilize the DCT, the invertible transform at the core of JPEG. JPEG achieves compression through the following steps: obtain 8×8 blocks AF corresponding to the DCT-domain signal, quantize the elements of AF using empirically-determined quantization steps (the non-invertible step in encoding), apply entropy coding to the quantized elements. Our algorithms operate on transform-domain signals; thus, a JPEG-encoded image will have its entropy coding removed prior to resampling. Entropy coding is restored so that the image is again in JPEG format, and the image is displayed or printed. These methods can also be extended to other image representation formats, such as MPEG.

Tomasz Nowicki and Marco Martens are with IBM T. J. Watson Research Ctr., 1101 Kitchawan Rd., Rte. 134, Yorktown Heights, NY 10598 USA. J. Q. Trelewicz is with IBM Almaden Research Ctr., 650 Harry Rd., San Jose, CA 95120 USA (e-mail: trelewicz@us.ibm.com). Timothy Trenary is with IBM Printing Systems Div., 6300 Diagonal Hwy., Boulder, CO 80301 USA.

This paper is organized as follows. In Sec. II we develop notation and terminology for the paper, as well as discussing foundation technologies for our resampling algorithm. In Sec. III we describe our resampling algorithm, giving properties of the algorithm and its convergence to conventional resampling algorithms. Finally, in Sec. IV, we give a specific example of application of the algorithm, showing resampling results for an image.

II. Background

The following notation is used in this paper. Let $\mathbb{R}^+ = \mathbb{R} \cap [0, \infty)$ and for $m_1, m_2 \in \mathbb{Z}$, $m_1..m_2 = [m_1, m_2] \cap \mathbb{Z}$.

A. Sampled and Continuous Images

It is assumed in this work for purposes of rigor that the sampled image is a representative of a continuous image. Specifically, let

$$\varphi(t) \quad = \quad \begin{cases} \frac{\sin(t)}{t}, & t \neq 0 \\ 1, & t = 0 \end{cases}$$

Let $f : \mathbb{R} \to \mathbb{C}$ be bandlimited to the band $[-\Omega_f, \Omega_f]$ in the sense that its Fourier transform $\hat{f}$ satisfies $\text{supp}\hat{f} \subset [-\Omega_f, \Omega_f]$. The Whittaker-Kotel'nikov-Shannon (WKS) sampling theorem [4] implies that

$$f(t) = \sum_{k \in \mathbb{Z}} f(kT)\varphi(\Omega_f(t - kT)) \tag{1}$$

where $T = \pi/\Omega_f$. Knowledge of the samples

$$\{f(kT) : k \in \mathbb{Z}\}$$

is thus equivalent to knowledge of f. We let $F(k) = f(kT)$. Such criteria for f are assumed throughout this paper. The specific units of Ω_f are not of particular importance, since we are interested in resampling ratios rather than absolute signal measure in the time or space domain. Furthermore, it is assumed that a function F scaled by b is also bandlimited such that $\text{supp}\hat{f} \subset [-b\Omega_f, b\Omega_f]$. In the case where f is of finite support (e.g., if f represents a digital photograph), f is taken to be periodically extended with period NT for purposes of analysis.

B. Transforms on Sampled Signals

A transform on higher dimensional data is called *separable* if there is a orthogonal transform on one-dimensional (1-D) data such that the multidimensional transform can be performed by applying the 1-D transform to each dimension of the signal. In sequel we will only discuss 1-D transforms, since the corresponding results hold automatically for all separable transforms generated by such a 1-D transform. Fix such a transform A. The samples on which this transform acts are vectors in $\mathcal{S} \subset \mathbb{C}^{\ltimes}$ for some fixed $n \in \mathbf{N}$, n is called the block size and $\mathcal{S}$ the domain of the transform, the "sample domain". In practice, samples are taken with finite precision; thus, $\mathcal{S}$ is point-wise bounded by M. For example, $\mathcal{S} = [0, 1]^n$ may be the sample space of a digital photograph, with sample values in $[0, 1]$. The sample-domain signal is denoted $F : \mathbb{N} \to \mathcal{S}$, which is the unique representation of $f : \mathbb{R} \to \mathcal{S}$ discussed above. Let $\mathcal{T} \subset \mathbb{C}^n$ denote the range of the transform, comprising the "transform domain", $\mathcal{T} = \mathcal{A}\mathcal{S}$. Thus let $A : \mathcal{S} \to \mathcal{T}$ be the linear, orthogonal transform, which will, in our examples, be the 1-D transform.

The whole signal, potentially consisting of n-blocks $\{B_k\}$, will be assumed to have dimension N, where n divides N. Let $\mathcal{R} \subset \mathbb{C}^N$ denote the set of possible signals in the transform domain. When AF is resampled by factor b, the resulting signal has dimension $N' = \lfloor bN \rfloor$, and lies in $\mathcal{R}' \subset \mathbb{C}^{N'}$. When the transform is the Fourier transform, we denote AF by $\hat{F}$.

Let $D : \mathcal{S} \times \mathcal{S} \to \mathbb{R}^+$ be the distance between two elements given by $D(G, H) = \|G - H\|$ (the ℓ_2 norm). We use here ℓ_2 norm, because Signal-to-Noise Ratio (SNR) is based on the ℓ_2 or L_2 norms. Let $D(G, \mathcal{S})$ be the "distance" from G to $\mathcal{S}$, given by

$$D(G, \mathcal{S}) = \min_{H \in \mathcal{S}} D(G, H)$$

The usage of D will be clear from the context. Similar definitions are used for the distance to other spaces.

C. Error Diffusion for Image Processing

Error Diffusion (ED) methods have been used in screening for years. Screening by ED reduces quantization error locally in the image by filtering the quantization error in a feedback loop. Furthermore, the local screen is required to be "noisy" (i.e., with local high-frequency content), avoiding the creation of regular patterns in the screened image to which the human visual system (HVS) will be sensitive. For example, in [5] the authors describe ED screening which is performed so that the image can be displayed at several resolutions after screening without creating moiré effects from sampling the screened image. However, true ED methods are inherently serial, leading to high levels of computational complexity. In [6], a lookup-table (LUT) method of approximated ED screening is discussed, which has lower computational complexity than conventional serial ED methods.

In Sec. III, we extend these ED techniques to resampling, evaluating the relative quality of the resulting image.

III. ED Resampling

According to the algorithm, the solution provided does 1-D resampling, with an arbitrary resampling factor $b > 0$. It can be seen from the description below that the principles of the algorithm can be applied to MPEG compressed data or data compressed by other compression techniques that utilize DCT transforms or wavelet transforms that separate into orthogonal 1-D transforms.

In this algorithm the 1-D resampling operations are described by consecutive resampling algorithms of the blocks by one of two chosen resampling algorithms with appropriate resampling factors. It is assumed that there is a given collection $\{\Psi_k\}$ of fixed, pre-calculated resampling matrices, with corresponding resampling ratios $\{a_k\}$, assumed to be numbered sequentially according to the canonical ordering on $\mathbb{R}$. The matrices $\{\Psi_k\}$ all operate on an integer multiple of n samples of data at a time; i.e., all matrices operate on collections of the same block size and aspect ratio. This section limits the discussion to 1-D resampling, since multidimensional resampling is a separable operation into the product of 1-D resampling, and therefore the algorithm can be trivially generalized to multidimensional resampling. The resampling of images with more colors (such as RGB or CMYK) is done by repetition of the resampling on each of the colors separately.

We utilize ED methods for resampling with Φ_b in the following manner. Resampling blocks with Ψ_k and Ψ_{k-1} causes local resampling quantization error, in that the local average resampling factor is different from b. Filtering this quantization error locally, so that the average local resampling factor is kept closer to b, reduces the local quantization error. As in ED screening, it is also desirable that the resampling be noisy, so that regular patterns at block-level boundaries are not created. Specifically, Ψ_k and Ψ_{k-1} are deployed on a block-by-block basis according to a pseudorandom distribution.

The LUT ED screening method employs the LUT to determine whether a given pel is painted on or off in a screened image to maintain a target local average intensity level. In a similar manner, LUT ED resampling utilizes the LUT to determine whether a given collection of transform-domain coefficients is resampled with Ψ_k or with Ψ_{k-1}, maintaining a target local average resampling value.

A. Details of the Algorithm

Suppose that we want to resample by $b \in (a_m, a_k)$, where $b \notin \{a_k\}$. It is shown in Lemma III.2 that the "best results" are obtained (i.e., the approximation errors at the block boundaries are minimized) when $(a_k, a_m) \cap \{a_k\} = \emptyset$.

Define $\nu = mesh\{a_k\}$; i.e., largest distance between adjacent sampling ratios. If ν is large or b is far from the closest a_k, severe resampling artifacts can be realized in the image. Specifically, the ED process resamples different image blocks by different amounts, so that the larger $\delta_k = a_k - a_{k-1}$, the more marked the local resampling differences between blocks. However, this is not a limiting factor of the algorithm, since the required mesh (i.e., largest distance between adjacent sampling ratios) of $\{a_k\}$ can be determined by analysis of sample images offline.

In this algorithm, ED is used so that the average resampling factor is perceived to be b although the actual resampling is done locally with the factors a_k and a_{k-1}. This resampling transform is denoted $\Phi_b : \mathcal{R} \to \mathcal{R}'$. The precise resampling transform is denoted $\Psi_b : \mathcal{R} \to \mathcal{R}'$.

This algorithm can be combined with any number of other resampling algorithms of the prior art to achieve arbitrary resampling with computational efficiency. In [7] and [8], the authors discuss a method for computing the DCT, which may be combined with our resampling algorithm for computational efficiency. In [9], the authors discuss a floating-point method for doing block-domain image processing, including resampling, rotation, and shearing. These algorithms may be used to calculate $\{\Psi_k\}$ for Φ_b.

B. Computational Complexity

This section outlines the computational complexity of resampling with Φ_b as compared to Ψ_b. In both applications, the signal is resampled in the transform domain, so that any operations required for the removal and reapplication of entropy coding are required for both. The number of operations required for resampling the image per sample is denoted α, and is a function of the Ψ employed for both. Precise computational complexity depends on the underlying Ψ and A, but can be discussed here for specific implementation examples.

For example, consider resampling DCT-domain data for JPEG-compressed images. When resampling with Ψ_b, because b may be arbitrary, and N may be unknown until the signal is received, optimized resampling configurations may not be developed in advance for all specific values of b. As a result, the modified DCT blocks, which are no longer 8×8, must be transformed back to 8×8 blocks to be compatible with JPEG decoders. Thus, Ψ_b requires

1. Resampling the blocks, modifying their dimension, taking O_R operations per sample,
2. Calculating A^{-1} of the resampled blocks, taking O_A operations per sample,
3. Applying A to recreate the 8×8 blocks, taking O'_A operations per sample. If using an algorithm such as binDCT-C [10], which has been optimized specifically for 8-sample blocks in JPEG, the 8-sample block takes 25 additions and 8 shifts, or 4.125 operations per sample.

Using Φ_b, because the $\{\Psi_k\}$ are known in advance, steps 1-3 are optimized into a one-step resampling transform $\Psi_k : \mathcal{R} \to \mathcal{R}'$, taking O'_R operations per sample. Because b is not known in advance, step 2 above is assumed to not be optimized for operations, so that $O_A \approx O'_R$. ED is used to distribute the Ψ_k over the image area, which can be well-approximated with a lookup table [6]. When approximated by lookup table, ED adds negligible operations to the calculation. Thus, our ED resampling takes about O_A operations, compared to $O_R + O_A + O'_A$ for conventional resampling.

C. Distance

For analyzing the closeness of the approximation formed by $\Phi_b AF$ to $\Psi_b AF$, we define in this section a distance function d appropriate for such images or signals as may be resampled by this algorithm.

Let $d : \mathcal{R}' \times \mathcal{R}' \to \mathbb{R}^+$ be given by

$$d(F, G) = \|\hat{F} - \hat{G}\|_2 \equiv \|F - G\|_2$$

(i.e., the ℓ_2-norm of the difference between Fourier spectra.)

This distance is well-suited to comparing Φ_b with Ψ_b in the following manner:

- Because Φ_b and Ψ_b differ primarily in their behavior at block boundaries $\{\partial(B_k)\}$, where Φ_b potentially introduces a collection of errors, the measurement of behavior at these boundaries is the primary function of d. The perceptual significance of the behavior of Φ_b at $\{\partial(B_k)\}$ is dependent on the signal content at these boundaries; e.g., errors in high-frequency content or high-contrast content are more significant perceptually. Discontinuities in contrast or content at the block edges introduce high-frequency components into the Fourier spectrum, increasing d.

- High-frequency content present in the original signal is subtracted, so that it does not affect d.
- A spatial or chronological shift in the signal may result from local resampling quantization from Φ_b, which is considered to degrade resampled signal quality. For example, local spatial shifts in the resampling of a straight line can "break" the edge of the line, making it appear crooked in the resampled image. The HVS is particularly sensitive to this type of distortion.

Lemma III.1: For $b > 0$, $\exists\Psi_b$ such that

$$(A^{-1}\Psi_b AF)(x) = \frac{1}{b}F\left(\frac{x}{b}\right);$$

i.e., the signal resampled with Ψ_b is equal to the signal resampled continuously by b.

Proof: By (1), with $T = \pi/\Omega_f$,

$$f(t) \;=\; \sum_{k\in\mathbb{Z}} F(k)\varphi(\Omega_f(t-kT)) = \sum_{n=0}^{N-1} F(n)\varphi'(n)$$

Where

$$\varphi'(n) = \sum_{k\in\mathbb{Z}} \varphi(\Omega_f t - Nk\pi - n\pi)$$

which is well-defined since it converges in L_2, and which is not dependent on F. Thus, φ' is known in advance of resampling by Ψ_b. As stated above, supp$\hat{f} \subset [-b\Omega_f, b\Omega_f]$. Thus, let Ψ_b' be defined such that $(\Psi_b' F)(k) = f(kT/b)$. With $A^{-1}\Psi_b AF = A\Psi_b'(A^{-1}AF)$, the claim is proven. ■

Therefore, in the proofs which follow, it is assumed that the Ψ_b meets the criterion of Ψ_b in Lemma III.1. If instead

$$\max_{c>0}\left\|(A^{-1}\Psi_c AF)(x) - \frac{1}{c}F\left(\frac{x}{c}\right)\right\|_2 \in (0,\epsilon)$$

for fixed small $\epsilon > 0$, this non-ideality must be considered in the convergence results.

Lemma III.2: Suppose that $b \in (a_{k-1}, a_k)$, $m \le k-1$, $k \le m'$ and Φ_b utilizes Ψ_m and $\Psi_{m'}$. Then $\max_F d(A^{-1}\Phi_b AF, A^{-1}\Psi_b AF)$ is minimized for $m = k-1$ and $m' = k$.

Proof: It suffices to consider 1-D resampling, since multi-dimensional resampling follows by separability. Let m' and m be the indices that minimize $d(A^{-1}\Phi_b AF, A^{-1}\Psi_b AF)$.

Let block B_j be given and let B_{j,a_m} be the block scaled continuously by $a_m < b$. On B_j, assume that Ψ_m is used, and denote by η the local shift in x from resampling quantization. Then by Lemma III.1, on block B_{j,a_m},

$$(A^{-1}\Psi_{a_m} AF)(x) = \frac{1}{a_m}F\left(\frac{x}{a_m} - \eta\right)$$

Let μ be Lebesgue measure, $\alpha_1 = 2\pi iy/a_m$, $\alpha_2 = \exp(-2\pi iy\eta)/a_m$, $\beta_1 = 2\pi iy/b$, $\beta_2 = 1/b$, and note that $\mu(B_{j,a_m}) < \infty$ and the integrand is bounded. Then after a change of variables,

$$\begin{aligned}
&\|A^{-1}\Phi_b AF - A^{-1}\Psi_b AF\|_2^2\\
&= \int_{B_{j,a_m}} |F(x)\,(\alpha_2\exp(\alpha_1 x) - \beta_2\exp(\beta_1 x))|^2\,dx\\
&\le M^2\int_{B_{j,a_m}} |\alpha_2\exp(\alpha_1 x) - \beta_2\exp(\beta_1 x)|^2\,dx\\
&= M^2\int_{B_{j,a_m}} \left(\alpha_2^2+\beta_2^2 - 2\alpha_2\beta_2\cos((\alpha_1-\beta_1)x)\right)dx
\end{aligned}$$

The integrand is equal to

$$\begin{aligned}
&\frac{(a_m - b\cos 2\pi y\eta)^2 + b^2\sin^2 2\pi y\eta}{a_m^2 b^2} \\
&+\frac{2\exp(-2\pi iy\eta)}{a_m b}\left(1-\cos((\alpha_1-\beta_1)x)\right)
\end{aligned} \tag{2}$$

It is assumed that η is linear in $a_m - b$. For $a_m - b$ not small, the term $(a_m - b\cos 2\pi y\eta)^2$ of (3) will be significant. Thus, smaller distance is achieved for $a_m - b$ small. Then $\sin 2\pi y\eta$ is bounded by a linear in $a_m - b$. Also in $a_m - b$, $\alpha_1 - \beta_1$ is linear, so that $\cos((\alpha_1-\beta_1)x) \to 1$ is bounded by a quadratic. Thus (3) is bounded by a quadratic in $a_m - b$.

In a similar manner, considering $F \equiv 1$,

$$\begin{aligned}
&\max_F \|A^{-1}\Phi_b AF - A^{-1}\Psi_b AF\|_2^2\\
&\ge \int_{B_{j,a_m}} \left(\alpha_2^2+\beta_2^2-2\alpha_2\beta_2\cos((\alpha_1-\beta_1)x)\right)dx
\end{aligned}$$

Thus, the smaller $|a_m - b|$, the smaller the lower bound.

A similar argument for blocks scaled with $\Psi_{m'}$ gives $m' = k$. ■

The upper bound in Lemma III.2 also gives the following result:

Lemma III.3: Let mesh$\{\Psi_k\} = \max_k \delta_k$. Then

$$\lim_{\text{mesh}\{\Psi_k\}\to 0}\ \max_b d(A^{-1}\Phi_b AF, A^{-1}\Psi_b AF) = 0$$

with quadratic approach to zero.

IV. Results

Results are shown in this section in the context of resampling JPEG-encoded images purely in the transform domain, for high-throughput production printing or fast image display. However, the algorithm is easily extendable to sample-domain, wavelet-domain, and interdomain (e.g., JPEG-to-wavelet) resampling, since the algorithm can use any pre-calculated resampling algorithm. It is another simple extension to use this algorithm to calculate arbitrary resampling factors; e.g., 0.562.

The following example shows a specific use of the algorithm.

Suppose that $a_{k-1} = 4/8$ and $a_k = 5/8$, and that Ψ_{k-1} is designed to transform 2 DCT blocks into 1 DCT block, and Ψ_k is designed to transform 8 DCT blocks into 5 DCT blocks. Take $b = 9/16$. The ED algorithm partitions the image into groups of 8 and 2 DCT blocks, where the 2-block groups are resampled with Ψ_{k-1} and the 8-block groups are resampled with Ψ_k.

For a given ratio $a_{k-1} < b < a_k$, the block indices $1..N/n$ are partitioned into superblocks $\tilde{B}_i$ of lengths γ_k and γ_{k-1} such that $a_k\gamma_k$ and $a_{k-1}\gamma_{k-1}$ are both divisible by n. In this example, $\gamma_k = 2n$ and $\gamma_{k-1} = 5n$. These integer γ_j exist in this case because b is rational, with small numerator and denominator. These superblocks are resampled with Ψ_{k-1} and Ψ_k), respectively. Then for each such superblock $\tilde{B}_i$, one will make an error of size $e_i = a_j\gamma_j - b\gamma_j$, with $j \in \{k-1, k\}$. Let $E_j = \sum_{i \le j} e_i$, be the accumulated error of the first j consecutive superblocks $\tilde{B}_i$. The ED algorithm prescribes the choice of a_k or a_{k-1} for the superblock $\tilde{B}_{j+1}$, such that $|E_{j+1}|$ is minimized. This rule suggests the iterative definition

$$E_{j+1} = E_j + \begin{cases} \Delta_{k-1}, E_j > E \\ \Delta_k, E_j < E \end{cases}$$

where $E = -\frac{\Delta_{k-1}+\Delta_k}{2}$. This algorithm for E_j may be represented as a graph of a function $f : \mathbb{R} \to \mathbb{R}$. Using this graph and the fact that $E_0 = 0$ one proves

Lemma IV.1: For all $i_1 < i_2$

$$E + \Delta_{k-1} < \sum_{i=i_1}^{i_2} e_i < E + \Delta_k$$

The meaning of this lemma is that any consecutive sequence in the input data, say with length L, will be resampled to become a sequence of length L' with

$$|L' - bL| = |\sum_{i=i_1}^{i_2} e_i|,$$

with the error in the size of the resampled data being uniformly bounded, and the lemma giving the precise bound.

The resulting 9/16 resampled image is shown in Fig. 2, with the original image in Fig. 1 for comparison. It should be noted that ED artifacts are not evident in Fig. 2.

V. Conclusions and Summary

We have presented here an approach to resampling images or signals with arbitrary resampling ratio. The approach is well-suited to software or hardware implementation, allowing use of computation-optimized resampling operations for a finite, fixed set of resampling ratios. These fixed ratios are used to approximate the arbitrary ratio. When the ratios are chosen with reasonably fine mesh, resampling quantization artifacts in the resulting resampled signal are not perceptible. Furthermore, it was shown that refinement of the mesh reduces the resampling quantization error quadratically.

Fig. 1. The original image F, reduced with whole-image biorthogonal resampling for column-width display with no block-edge artifacts.

Fig. 2. The resampled image resulting from Φ_b with $b = 9/16$.

References

[1] A M Darwish, M S Bedair, and S I Shaheen, "Adaptive resampling algorithm for image zooming," *IEE Proc., Vis. Image Sig. Process.*, vol. 144, no. 4, pp. 207–212, Aug. 1997.

[2] A Gotchev, J Vesma, T Saramaki, and K Egiazarian, "Digital image resampling by modified B spline functions," in *Proc. of NORSIG2000*, (Vildmarkshotellet Kolmarden, Sweden), 13-15 June 2000, pp. 259–262.

[3] ITU-T Rec. T8I/10918-1, *Information technology: Digital compression and coding of continuous-tone still images*, IISO/IEC, 1994.

[4] R J Marks, *Introduction to Shannon Sampling and Interpolation Theory*, Springer-Verlag, New York, 1991.

[5] P-W Wong, "Adaptive error diffusion and its application in multiresolution rendering," *IEEE Trans. Image Proc.*, vol. 5, no. 7, pp. 1184–1196, July 1996.

[6] M Meşe and P P Vaidyanathan, "Look up table (LUT) method for image halftoning," in *Proc. of IEEE ICIP 2000, v.III*, (Vancouver, BC, Canada), 10-13 Sept. 2000, pp. 993–996.

[7] A. Elnaggar and H. M. Alnuweiri, "A new multidimensional recursive architecture for computing the discrete cosine transform," *IEEE Trans. Circuits Sys. Video Tech.*, vol. 10, no. 1, pp. 113–119, Feb. 2000.

[8] N-P Khanh, A Docef, and F Kossentini, "Quantized discrete cosine transform: a combination of DCT and scalar quantization," in *Proc. of IEEE ICASSP 1999*, (Phoenix, AZ, USA), 15-19 Mar 1999, pp. 3197–3200.

[9] B Shen and I K Sethi, "Block-based manipulations on transform-compressed images and videos," *Multimedia Sys.*, vol. 6, pp. 113–124, 1998.

[10] T D Tran, "The binDCT: fast multiplierless approximation of the DCT," *IEEE Sig. Proc. Let.*, vol. 7, no. 6, pp. 141–144, June 2000.