Non-independent Randomized Rounding and an Application to Digital Halftoning

Benjamin Doerr and Henning Schnieder

Mathematisches Seminar II, Christian–Albrechts–Universität zu Kiel,
Ludewig–Meyn–Str. 4, D–24098 Kiel, Germany,
`{bed,hes}@numerik.uni-kiel.de`,
WWW home page: `http://www.numerik.uni-kiel.de/~{bed,hes}/`

Abstract. We investigate the problem to round a given $[0,1]$–valued matrix to a $0,1$ matrix such that the rounding error with respect to 2×2 boxes is small. Such roundings yield good solutions for the digital halftoning problem as shown by Asano et al. (SODA 2002). We present a randomized algorithm computing roundings with expected error at most 0.6287 per box, improving the 0.75 non-constructive bound of Asano et al. Our algorithm is the first one solving this problem fast enough for practical application, namely in linear time.
Of a broader interest might be our rounding scheme, which is a modification of randomized rounding. Instead of independently rounding the variables (expected error 0.82944 per box in the worst case), we impose a number of suitable dependencies.
Experimental results show that roundings obtained by our approach look much less grainy than by independent randomized rounding, and only slightly more grainy than by error diffusion. On the other hand, the latter algorithm (like all known deterministic algorithms) tends to produce unwanted structures, a problem that randomized algorithms like ours are unlikely to encounter.
Key words: Randomized rounding, discrepancy, digital halftoning.

1 Introduction

In this paper we are concerned with rounding problems and, in particular, the digital halftoning problem. In general form, rounding problems are of the following type: Given some numbers $x_1, \dots, x_n$, one is looking for roundings $y_1, \dots, y_n$ such that some given error measures are small. By rounding we always mean that $y_i = \lfloor x_i \rfloor$ or $y_i = \lceil x_i \rceil$. Since there are 2^n possibilities, such rounding problems are good candidates for hard problems. In fact, even several restricted versions like the combinatorial discrepancy problem are known to be NP–hard.

On the other hand, there are cases that can be solved optimally in polynomial time. Knuth [9] for example has shown that there exists a rounding such that all errors $|\sum_{i=1}^{k}(y_i - x_i)|$ and $|\sum_{i=1}^{k}(y_{\pi(i)} - x_{\pi(i)})|$ for a fixed permutation π are at most $\frac{n}{n+1}$. Such roundings can be obtained by computing a maximum flow in a network. A recent generalization to arbitrary totally unimodular rounding problems can be found in [5].

1.1 The Digital Halftoning Problem: A Matrix Rounding Problem

Our study of rounding problems is motivated by an application to digital halftoning. The digital halftoning problem is to convert a continuous-tone intensity image (each pixel may have an arbitrary 'color' on the white-to-black scale) into a binary image (only black and white dots are allowed). An intensity image can be represented by a $[0,1]$–valued $m \times n$ matrix A. Each entry a_{ij} corresponds to the brightness level of the pixel with coordinates (i,j). Since many devices, e.g., laser printers, can only output white and black dots, we have to round A towards a $0,1$ matrix. Naturally, this has to be done in a way that the resulting image looks similar to the original one.

This notion of similarity is a crucial point. From the viewpoint of application, similarity is defined via the human visual system: A rounding is good, if an average human being can retrieve the same information from both the original image and the rounded one. Using this criterion, several algorithms turned out to be useful: Floyd and Steinberg [7] proposed the error diffusion algorithm that rounds the entries one by one and distributes the rounding error over neighboring not yet rounded entries. Lippel and Kurland [10] and Bayer [3] investigated the ordered dither algorithm, which partitions the image into small submatrices and rounds each submatrix by comparing its entries with a threshold matrix of same dimension. One advantage of this approach is that this algorithm can be parallelized easily. Knuth [8] combined ideas of both approaches to get an algorithm called dot diffusion.

What was missing so far from the theoretical point of view is a good mathematical formulation of similarity. Such a similarity measure is desirable for two reasons: Firstly, it allows to compare algorithms without extensive experimental testing. This is particularly interesting, since comparing different halftonings is a delicate issue. For example, it makes a huge difference whether the images are viewed on a computer screen or are printed on a laser printer. Even different printers can give different impressions. Therefore, a more objective criterion would be very helpful. A second reason is that having a good criterion, one would have a clearer indication of how a digital halftoning algorithm should work. Thus developing good algorithms would be easier.

On this year's SODA conference, Asano et al. [2] reported that they made some progress into this direction. Experimental results indicate that good digital halftonings have small error with respect to all 2×2 subregions. More formally, we end up with this problem:

1.2 Problem Statement

Let $A \in [0,1]^{m \times n}$ denote our input matrix. A set $R_{ij} := \{(i,j),(i+1,j),(i,j+1),(i+1,j+1)\}$ for some $i \in [m-1], j \in [n-1]$ is called a 2×2 subregion (or box) in $[m] \times [n]$.[1] Denote by $\mathcal{H}$ the set of all these boxes. We write $A_{R_{ij}}$

---

[1] For a number r we denote by $[r]$ the set of positive integers not exceeding r.

for the 2×2 matrix $\begin{pmatrix} a_{i,j} & a_{i,j+1} \\ a_{i+1,j} & a_{i+1,j+1} \end{pmatrix}$ induced by R_{ij}. For any matrix A put $\Sigma A := \sum_{i,j} a_{ij}$.

For a matrix $B \in \{0,1\}^{m \times n}$ — which by definition is a rounding of A — we define the rounding error of A with respect to B by

$$d_{\mathcal{H}}(A,B) := \sum_{R \in \mathcal{H}} |\Sigma A_R - \Sigma B_R| .$$

We usually omit the subscript $\mathcal{H}$ when there is no danger of confusion.

From a broader perspective, this rounding problem has the interesting aspect that the errors regarded depend on few variables (here four) only. Thus the common approach of randomized rounding is not very effective, since it depends on high concentration results for sums of many independent random variables.

1.3 Theoretical Results

Asano et al. [2] exhibited that roundings B such that $d(A,B)$ is small, yield good digital halftonings. They showed that for any A an optimal rounding B^* satisfies $d(A,B^*) \leq \frac{3}{4}|\mathcal{H}|$. They also gave a polynomial time algorithm computing a rounding B such that $d(A,B) - d(A,B^*) \leq \frac{9}{16}|\mathcal{H}| = 0.5625|\mathcal{H}|$. It is easy to see that there are matrices A such that all roundings (and in fact all integral matrices) B have $d(A,B) \geq \frac{1}{2}|\mathcal{H}|$.

A major draw-back of the algorithm given in [2] is that it is not very practical, as it requires the solution of an integer linear program with totally unimodular constraint matrix. As pointed out by the authors, this is too slow for digital halftoning applications.

In this paper we present a randomized algorithm that runs in linear time. It may be used in parallel without problems. The roundings computed by our algorithm have an expected error that exceeds the optimal one by at most $\frac{15}{32}|\mathcal{H}| \leq 0.4688|\mathcal{H}|$ (instead of $0.5625|\mathcal{H}|$). In addition, our roundings satisfy the absolute bound $E(d(A,B)) \leq 0.6287|\mathcal{H}|$ beating the non-algorithmic bound of $\frac{3}{4}|\mathcal{H}|$ of [2].

The distribution of the resulting error is highly concentrated around the expected value. The probability that the error exceeds $(0.6287+\varepsilon)|\mathcal{H}|$ is bounded by $\exp(-\Omega(\varepsilon^2|\mathcal{H}|))$. More concrete, when processing a 4096×3072 pixel image, the probability that the error exceeds the expectation by more than $0.01|\mathcal{H}|$ is less than $2.5 \cdot 10^{-7}$. Hence a feasible approach in practice is to compute the random rounding without checking its error.

Another nice feature from the viewpoint of application to digital halftoning is that the roundings computed by our algorithm have small rounding error also with respect to other geometric structures that might indicate good (according to humans' eyes) halftonings. In particular, the expected error with respect to the set $\mathcal{H}_3$ of all 3×3 boxes is bounded by $0.82944|\mathcal{H}_3|$, and the error with respect to single entries (1×1 boxes) is bounded by $0.5|\mathcal{H}_1|$, which is optimal in the worst-case.

1.4 Non-independent Randomized Rounding

The key idea of our algorithm might also be of a broader interest. We develop a randomized rounding scheme where the individual roundings are not independent. The classical approach of randomized rounding due to Raghavan and Thompson [11, 12] is to round each variable independently with probability depending on the fractional part of its value. This allows to use Chernoff-type large deviation inequalities showing that a sum of independent random variables is highly concentrated around its expectation. Randomized rounding has been applied to numerous combinatorial optimization problems that can be formulated as integer linear programs. Though being very effective in the general case, one difficulty with randomized rounding is to use structural information of the underlying problem. One way to overcome this is to use correlation among the events. This allows to strengthen the classical bounds as shown by Srinivasan [13].

In this paper we try so use the structure in an earlier phase, namely in the design of the random experiment. This leads to randomized roundings where the variables are not rounded independently. There have been a few attempts in this direction (cf. [11, 12, 4]), but they do not go further than translating constraints into dependencies or imposing restrictions on the number of variables rounded up or down. Therefore we feel that the option to design the random experiment in a way that it reflects the structure of the underlying problem has not been exploited sufficiently. In this paper we try to move a step forward into this direction. We impose dependencies that are not necessary in the sense of feasibility, but helpful since they reduce the expected rounding error. For the rounding problem arising from digital halftoning, this improves the bound of $0.82944|\mathcal{H}|$ obtained by independent randomized rounding down to $0.6287|\mathcal{H}|$.

In some sense this result can be seen as an extension of ideas of [1, 6]. These results, however, apply only to a very restricted class of rounding problems: The so-called combinatorial discrepancy problem is to round the vector $x = \frac{1}{2}\mathbf{1}_n$ to a $0, 1$ vector y such that $\|A(x - y)\|_\infty$ is small for some $0, 1$ matrix A.

Though a general result on non-independent randomized rounding is hard to imagine due to the high dependence on the structure of the target problem, we believe that this work shows the power of non-independent randomized rounding and motivates further research in this area.

1.5 Experimental Results

To estimate the visual quality of our algorithm, we generated roundings of several both real-world and artificial images with existing algorithms and the new one. It turns out that the use of dependencies in the random experiment greatly reduces the graininess of the output images (compared to independent randomized rounding). Still, the error diffusion algorithm remains unbeaten in this category.

On the other hand, since our algorithms is randomized, we have no problems with unwanted structures or textures, a weakness of error diffusion and other deterministic algorithms. Summarizing our experimental results, we feel that our work brings randomized rounding back into a group of feasible approaches for the digital halftoning problem in which neither can claim himself superior to the others.

2 Randomized Rounding

For a number x we write $\lfloor x \rfloor$ for the largest integer not exceeding x, $\lceil x \rceil$ for the smallest being not less that x, and $\{x\} := x - \lfloor x \rfloor$ for the fractional part of x. We say that some random variable X is a *randomized rounding* of x if $P(X = \lfloor x \rfloor + 1) = \{x\}$ and $P(X = \lfloor x \rfloor) = 1 - \{x\}$. In particular, if $x \in [0, 1]$, we have $P(X = 1) = x$ and $P(X = 0) = 1 - x$.

We first analyze what can be achieved with independent randomized rounding. The result below is needed not only to estimate the superiority of non-independent randomized rounding, but also in the proofs of two of our results. Note that the proof (omitted for reasons of space) has to be different from typical randomized rounding applications: Since the boxes are small, using a large deviation bound makes no sense and one has to compute the expected error exactly.

We say that B is an *independent randomized rounding* of A if each entry b_{ij} is a randomized rounding of a_{ij} and all these roundings are mutually independent.

Theorem 1. *Let $A \in [0,1]^{m \times n}$, B an independent randomized rounding of A and B^* an optimal rounding of A. Then*

$$\begin{aligned} E(d(A,B)) &\le 0.82944|\mathcal{H}|, \\ E(d(A,B)) &\le 0.75|\mathcal{H}| + d(A,B^*). \end{aligned}$$

□

The two bounds of Theorem 1 are tight as shown by matrices with all entries 0.4 and 0.5 respectively.

3 Non-independent Randomized Rounding

In this section we improve the previous bounds by adding some dependencies to the rounding process. We start with an elementary approach called *joint randomized rounding*, which reduces the chance that neighboring matrix entries are both rounded in the wrong way. This leads to a first improvement in Corollary 1. We then add further dependencies leading to what we call *block randomized rounding.*

Definition 1. *Let $a_1, a_2 \in [0,1]$. We say that (b_1, b_2) is a* joint randomized rounding *of (a_1, a_2) if $P(b_1 = 1 \wedge b_2 = 0) = a_1$, $P(b_1 = 0 \wedge b_2 = 1) = a_2$ and $P(b_1 = 0 \wedge b_2 = 0) = 1 - a_1 - a_2$ holds if $a_1 + a_2 \le 1$, and if $P(b_1 = 1 \wedge b_2 = 0) = 1 - a_2$, $P(b_1 = 0 \wedge b_2 = 1) = 1 - a_1$ and $P(b_1 = 1 \wedge b_2 = 1) = a_1 + a_2 - 1$ holds in the case $a_1 + a_2 > 1$.*

Immediately, we have

Lemma 1. *Let* $a_1, a_2 \in [0,1]$*. Then* (b_1, b_2) *is a joint randomized rounding of* (a_1, a_2) *if and only if the following two properties are valid:*

(i) For all $i \in [2]$*, we have* $P(b_i = 1) = a_i$ *and* $P(b_i = 0) = 1 - a_i$*. Hence* b_i *is a randomized rounding of* a_i*.*

(ii) $b_1 + b_2$ *is a randomized rounding of* $a_1 + a_2$*.*

The next lemma shows that two joint randomized roundings are superior to four independent randomized roundings in terms of the rounding error.

Lemma 2. *Let* $A = (a_{11}, a_{12}, a_{21}, a_{22})$ *be a box. Let* (b_{11}, b_{12}) *be a joint randomized rounding of* (a_{11}, a_{12})*, and* (b_{21}, b_{22}) *one of* (a_{21}, a_{22}) *independent of the first. Then the expected error of* A *with respect to* $B = (b_{11}, b_{12}, b_{21}, b_{22})$ *is at most* $E(d(A,B)) = \frac{16}{27}|\mathcal{H}| \le 0.5926|\mathcal{H}|$*. In comparison with an optimal rounding* B^* *we have* $E(d(A,B)) - d(A,B^*) \le 0.5|\mathcal{H}|$*.*

Proof (sketched). $b_{11} + b_{12}$ behaves like a randomized rounding of $a_{11} + a_{12}$, and the same holds for the second row. Hence all we have to do is to bound the expected deviation of a sum of two independent randomized roundings from the sum of the original values. ☐

The bounds Lemma 2 are sharp as shown by matrices having $a_{11} + a_{12} = a_{21} + a_{22} = \frac{1}{3}$ and $a_{11} + a_{12} = a_{21} + a_{22} = \frac{1}{2}$ respectively. Plastering the grid with joint randomized roundings already yields a first improvement over independent randomized rounding:

Corollary 1. *Let* $A \in [0,1]^{m \times n}$*. Compute* $B \in \{0,1\}^{m \times n}$ *by independently obtaining* $(b_{i,2j-1}, b_{i,2j})$ *from* $(a_{i,2j-1}, a_{i,2j})$ *by joint randomized rounding for all* $i \in [m], j \in [\frac{n}{2}]$*, and also independently obtaining* b_{in} *from* a_{in}*,* $i \in [m]$*, by usual randomized rounding, if* n *is odd. Then*

$$E(d(A,B)) \le 0.7111|\mathcal{H}|.$$

Proof. At least half of the boxes, namely all R_{ij} such that j is odd, are rounded in the manner of Lemma 2. The remaining ones contain four independent randomized roundings. Thus $E(d(A,B)) \le \frac{1}{2}0.5926|\mathcal{H}| + \frac{1}{2}0.82944|\mathcal{H}| \le 0.7111|\mathcal{H}|$ by Theorem 1 and Lemma 2. ☐

Block Randomized Rounding

Definition 2. *Let* $A = (a_{11}, a_{12}, a_{21}, a_{22})$ *be a box. We call* $B = (b_{11}, b_{12}, b_{21}, b_{22})$ *a* block randomized rounding *of* A *if*

(i) each single entry of B *is a randomized rounding of the corresponding one of* A*, i. e.,* $P(b_{ij} = 1) = a_{ij}$ *and* $P(b_{ij} = 0) = 1 - a_{ij}$ *for all* $i, j \in [2]$*.*

(ii) each pair of neighboring entries has the distribution of the corresponding joint randomized rounding, i. e., in addition to (i) we have for all $(i,j), (i',j') \in [2] \times [2]$ *such that either* $i \ne i'$ *or* $j \ne j'$*,*

$$P(b_{ij} + b_{i'j'} = \lfloor a_{ij} + a_{i'j'} \rfloor + 1) = \{a_{ij} + a_{i'j'}\}$$
$$P(b_{ij} + b_{i'j'} = \lfloor a_{ij} + a_{i'j'} \rfloor) = 1 - \{a_{ij} + a_{i'j'}\}.$$

(iii) the box in total behaves like a randomized rounding, i. e., we have

$$P(\Sigma B = \lfloor \Sigma A \rfloor + 1) = \{\Sigma A\}$$
$$P(\Sigma B = \lfloor \Sigma A \rfloor) = 1 - \{\Sigma A\}.$$

By item *(iii)* of this definition, block randomized roundings have low rounding errors. If B is a block randomized rounding of a 2×2 matrix A, then $E(d(A, B)) \leq \frac{1}{2}$ and $E(d(A, B)) - \min_{B^*} d(A, B^*) \leq \frac{1}{8}$. The interesting point is that block randomized roundings always exist:

Lemma 3. *For all boxes $A = (a_{11}, a_{12}, a_{21}, a_{22})$, a block randomized rounding exists and can be generated efficiently.*

Proof. We first show that it is enough to consider boxes A such that $s := \Sigma A = a_{11} + a_{12} + a_{21} + a_{22} \leq 2$. Assume $s > 2$. Define $A' = (a'_{ij})$ by $a'_{ij} := 1 - a_{ij}$. Then $\Sigma A' = 4 - \Sigma A < 2$. Let $B' = (b'_{ij})$ be a block randomized rounding of A'. Define $B = (b_{ij})$ by $b_{ij} := 1 - b'_{ij}$. Now it remains to show that the pair (A, B) fulfills (i) to (iii) of Definition 2, which we leave to the reader.

For the remainder of this proof let us assume that $s \leq 2$. Let $i, i', j, j' \in [2]$. We call $e = \{(i, j), (i', j')\}$ an *edge*, if either $i \neq i'$ or $j \neq j'$, and write $a_e := a_{ij} + a_{i'j'}$. An edge e is *heavy*, if $a_e > 1$. We first note that there cannot be disjoint heavy edges. If e and f were disjoint and heavy, then $s = a_e + a_f > 2$ contradicting our previous assumption. In particular, there can be at most two different heavy edges. We will treat these three cases separately after introducing some more notation. We denote by E_{ij} the 2×2 matrix having all entries zero except a single entry of one on position (i, j). For an edge $e = \{(i, j), (i', j')\}$ let $E_e := E_{ij} + E_{i'j'}$. Finally, put $E_\backslash := E_{11} + E_{22}$ and $E_/ := E_{12} + E_{21}$.

Case 1: No heavy edges. If $s \leq 1$, let B be such that $P(B = E_{ij}) = a_{ij}$ for all $i, j \in [2]$, and $P(B = 0) = 1 - s$. Here as well as in the remaining cases, we leave it to the reader to check that (i) to (iii) of Definition 2 are fulfilled. Assume now that $s > 1$, but there are still no heavy edges. By symmetry, we may assume that a_{11} is a minimal entry of A. Put $p_\backslash = \min\{s - 1, a_{11}\}$ and $p_/ = s - 1 - p_\backslash$. Define B by

$$\begin{aligned}
&P(B = E_\backslash) = p_\backslash, \\
&P(B = E_/) = p_/, \\
&P(B = E_{ii}) = a_{ii} - p_\backslash, i \in [2], \\
&P(B = E_{ij}) = a_{ij} - p_/, i, j \in [2], i \neq j.
\end{aligned}$$

Since a_{11} is minimal and there are no heavy edges, all probabilities defined above are non-negative.

Case 2: One heavy edge. Omitted.

Case 3: Two heavy edges. Omitted. □

Since block randomized roundings have low rounding errors and can be computed efficiently, the following algorithm suggests itself: Partition the input matrix A into 2×2 boxes and round each thereof as in the proof of Lemma 3. To be precise:

Definition 3. *Let* $A \in [0,1]^{m \times n}$*. We say that* B *is a* block randomized rounding *of* A*, if it is computed by the following rounding scheme:*

(i) for all $i \in [\frac{m}{2}], j \in [\frac{n}{2}]$*,* $R := \{2i-1, 2i\} \times \{2j-1, 2j\}$*,* $B_{|R}$ *is a block randomized rounding of* $A_{|R}$ *as in the proof of Lemma 3,*
(ii) if m *is odd, then for all* $j \in [\frac{n}{2}]$*,* $(b_{m,2j-1}, b_{m,2j})$ *is a joint randomized rounding of* $(a_{m,2j-1}, a_{m,2j})$ *as in Definition 1,*
(iii) if n *is odd, then for all* $i \in [\frac{m}{2}]$*,* $(b_{2i-1,n}, b_{2i,n})$ *is a joint randomized rounding of* $(a_{2i-1,n}, a_{2i,n})$ *as in Definition 1,*
(iv) if both m *and* n *are odd, then* b_{mn} *is a randomized rounding of* a_{mn}*.*
(v) All roundings in (i) to (iv) shall be independent.

We shall first analyze the error of block randomized roundings and then turn to the computational aspects of this approach.

Theorem 2. *Let* B *be a block randomized rounding of* $A \in [0,1]^{m \times n}$*. Then*

(i) The expected rounding error is $E(d(A,B)) \leq 0.6287|\mathcal{H}|$*.*
(ii) If B^* *is an optimal rounding, then*

$$E(d(A,B)) - d(A,B^*) \leq \tfrac{15}{32}|\mathcal{H}| \leq 0.4688|\mathcal{H}|.$$

(iii) $P(d(A,B) > E(d(A,B)) + \varepsilon|\mathcal{H}|) < 9\exp(-\frac{1}{72}\varepsilon^2(m-2)(n-2))$*.*

Proof (sketched). At least a quarter of all boxes (those R_{ij} where both i and j are odd) are block randomized roundings. Less than half of the boxes contain two independent joint randomized roundings, namely those R_{ij} where either i or j is odd. The remaining boxes, which are a fraction of at most a quarter, have all their entries rounded independently. Thus Theorem 1, Lemma 2 and the remark following Definition 2 yield the bounds (i) and (ii).

The proof of the large deviation bound is an application of the Chernoff inequality. Note that $\mathcal{H}$ can be partitioned into nine sets of boxes such that the errors within the boxes of each partition class are mutually independent random variables. □

We remark that the large deviation bound is not optimal, but by far sufficient for our purposes. The matrix A with all entries 0.3588 has $E(d(A,B)) \geq 0.6173|\mathcal{H}|$ for a block randomized rounding B. Thus our analysis is quite tight.

Algorithmic Properties of Block Randomized Rounding

Computing matrix roundings with the randomized approach above is fast. Each single block randomized rounding of a 2×2 block takes constant time, therefore the whole rounding can be done in linear time. From the view-point of application to digital halftoning, this is a crucial improvement over the algorithm of [2], that roughly has quadratic time complexity (ignoring a polylogarithmic factor). As stated in [2], this is too slow for high-resolution images.

The problem of computing time can be further addressed with parallel computing (and this is actually an issue when discussing digital halftoning algorithms). Since the block randomized roundings of 2×2 blocks are done independently, it is no problem to assign them to different processors.

Further Aspects of Block Randomized Roundings

Let us sketch another advantage of our approach. Suppose that we do not want to find a good rounding with respect to the 2×2 boxes, but with respect to larger structures, say 3×3 boxes. We currently have no hint whether the error with respect to these sets is a better measure for the visual quality of the resulting digital halftoning, but it seems plausible to try this experimentally. Hence we need an algorithm computing such roundings.

Whereas it seem difficult to extend the approach of Asano et al. to larger structures, non-independent rounding does very well: For 3×3 boxes, we may even use the same rounding scheme as before. Doing so, each 3×3 box contains exactly one block randomized rounding, two joint randomized roundings and one single randomized rounding. Since the four values of the block randomized rounding in total behave like a single randomized rounding, and so do the two values of each joint randomized rounding, the expected error of a 3×3 box is just given by Theorem 1. We have

Theorem 3. *With respect to the family* $\mathcal{H}_3$ *of* 3×3 *boxes, a block randomized rounding* B *of* A *has expected error* $d_{\mathcal{H}_3}(A,B) \leq 0.82944|\mathcal{H}_3|$.

We thus may claim that our approach has a broader range of application than previous results on this problem. This, by the way, works also in the other direction: From the view-point of digital halftoning, it is of course desirable that also the single entries are not rounded to badly, i.e., the error with respect to 1×1 boxes should not be too large. Our algorithm solves this in a very simple manner: If B is a block randomized rounding of A, then each single entry b_{ij} is a randomized rounding of a_{ij}. Hence the expected error of the single entries is at most $\frac{1}{2}mn$, which is, of course, optimal in the worst-case.

4 Experimental Results

We applied the three classical algorithms mentioned in the introduction together with independent randomized rounding and our algorithm to several images. For reasons of space we are not able do give any details concerning the three algorithms except for the remark that all three are deterministic. All image data used 1 byte per pixel resulting in an integer value between 0 and 255. We used two types of input data: Real-world images taken with a digital camera, and artificial images produced with a commercial imaging software. Naturally, the first type is more suitable to estimate how well the algorithm performs in real-world applications, whereas the second is better suited to demonstrate the particular strengths and weaknesses of an algorithm.

For reasons of space the images displayed in this paper are only small parts of the images we processed.[2] These parts have a size of 160×160 pixels, and are displayed in 72 dpi. All printers nowadays can handle higher resolutions, of course, but the single pixels would be harder to recognize, and some unwanted

---

[2] The full size images can be found at http://www.numerik.uni-kiel.de/~hes/NIRR.htm.

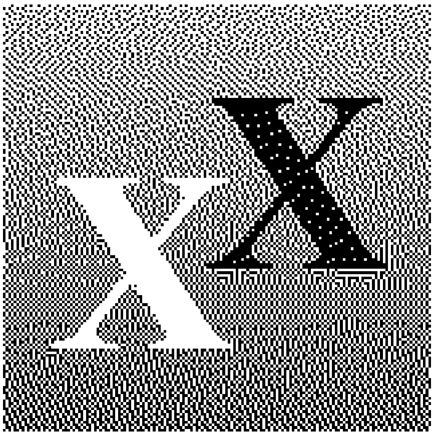

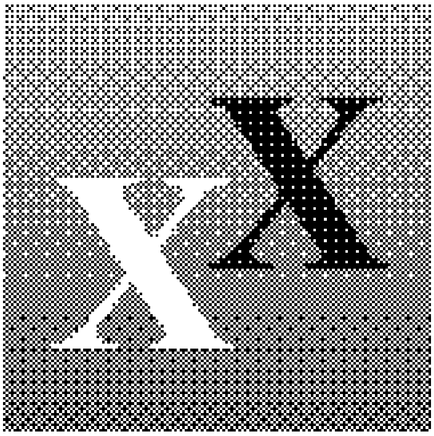

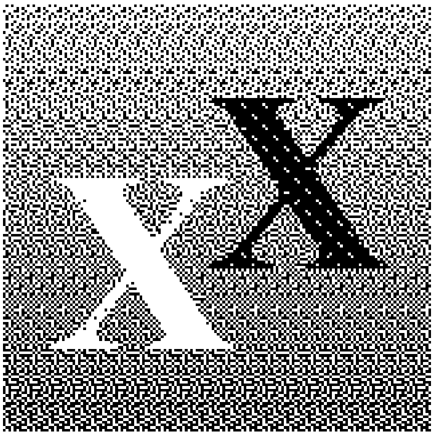

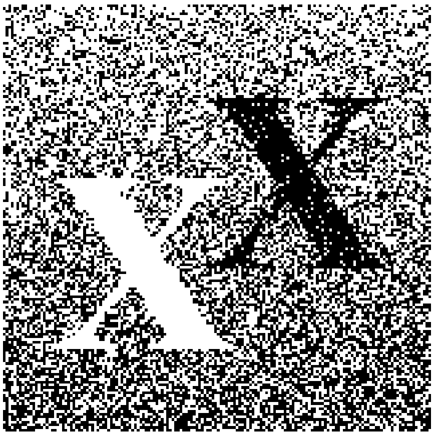

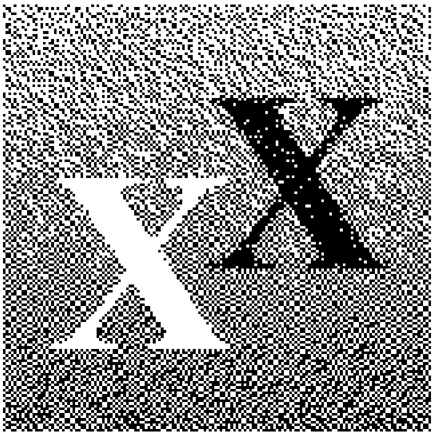

Fig. 1. Error Diffusion (top left).

Fig. 2. Ordered Dither (top right).

Fig. 3. Dot Diffusion (middle left).

Fig. 4. Randomized Rounding (middle right).

Fig. 5. Non-independent Randomized Rounding (bottom).

effects like small white dots disappearing in a large black area would spoil the result.

All known algorithms for the digital halftoning problem tend to produce some kind of structures or textures, which draw unwanted attention. Generally two kinds of textures can be observed. First there are regular patterns like snakes, crosses or labyrinths. In particular error diffusion and ordered dither algorithm tend to produce those, as can be seen in Fig. 1 and 2.

The second form of unwanted structures are grains. Grains emerge, if in dark (respectively light) parts of the picture two or more white (respectively black) pixels touch each other and thus build a recognizable block. As observed already in the seventies, randomized rounding is very vulnerable to this problem, which is why it is not used in practice for digital halftoning. On the other end we find error diffusion, which hardly produces any grains. It seems that algorithms that are good concerning graininess tend to produce unwanted structures and vice versa. In this sense, non-independent randomized rounding seems to be a fair compromise: Being by far less grainy than independent randomized rounding on the one hand, it is unlikely to produce unwanted structures on the other.

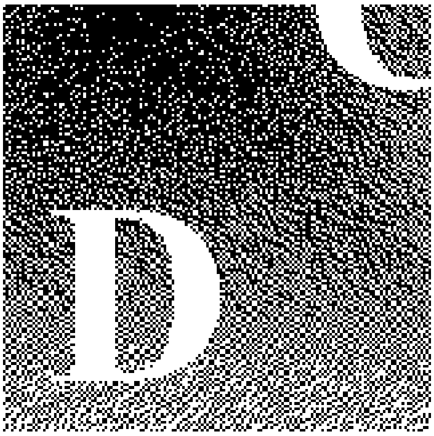

Fig. 6. Non-ind. Rand. Rounding.

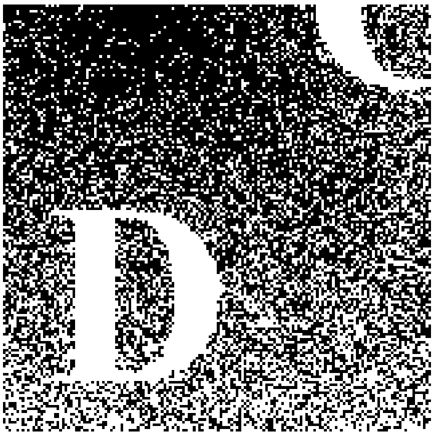

Fig. 7. Randomized Rounding.

5 Summary and Outlook

This paper describes a new approach in randomized rounding. By imposing suitable dependencies, we improve the expected rounding error significantly. For a particular problem arising in digital halftoning this improves previous algorithms both according to run-time and rounding error. In particular, we presented the first algorithm solving the rounding problem proposed by Asano et al. fast enough for practical application, namely in linear time.

On the methodological side, this paper shows that non-independent randomized rounding can be very effective if one succeeds in finding the right dependencies. We believe that this is a fruitful approach, in particular for problems where smallish structures do not allow to use large deviation estimates.

Acknowledgments

We would like to thank Tetsuo Asano, Naoki Katoh, Koji Obokata and Takeshi Tokuyama for providing us with a preprint of their SODA 2002 paper and some useful discussion on the topic.

References

1. N. Alon, B. Doerr, T. Łuczak, and T. Schoen. On the discrepancy of combinatorial rectangles. Tentatively accepted for publication in *Random Structures & Algorithms*, 2001.
2. T. Asano, N. Katoh, K. Obokata, and T. Tokuyama. Matrix rounding under the L_p-discrepancy measure and its application to digital halftoning. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2002.
3. B.E. Bayer. An optimum method for two-level rendition of continous-tone pictures. In *Conference Record, IEEE International Conference on Communications*, volume 1, pages (26–11)–(26–15). IEEE, 1973.
4. D. Bertsimas, C. Teo, and R. Vohra. On dependent randomized rounding algorithms. *Oper. Res. Lett.*, 24:105–114, 1999.
5. B. Doerr. Lattice approximation and linear discrepancy of totally unimodular matrices. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 119–125, 2001.
6. B. Doerr. Structured randomized rounding and coloring. In R. Freivalds, editor, *Fundamentals of Computation Theory*, volume 2138 of *Lecture Notes in Computer Science*, pages 461–471, Berlin–Heidelberg, 2001. Springer Verlag.
7. R.W. Floyd and L. Steinberg. An adaptive algorithm for spatial grey scale. In *SID 75 Digest*, pages 36–37. Society for Information Display, 1975.
8. D.E. Knuth. Digital halftones by dot diffusion. *ACM Trans. Graphics*, 6:245–273, 1987.
9. D.E. Knuth. Two-way rounding. *SIAM J. Discrete Math.*, 8:281–290, 1995.
10. B. Lippel and M. Kurland. The effect of dither on luminance quantization of pictures. In *IEEE Trans. Commun. Tech. COM-19*, pages 879–888, 1971.
11. P. Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *J. Comput. Syst. Sci.*, 37:130–143, 1988.
12. P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
13. A. Srinivasan. Improved approximation guarantees for packing and covering integer programs. *SIAM J. Comput.*, 29:648–670, 1999.