Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

PROPERTIES OF THE NORTA METHOD IN HIGHER DIMENSIONS

Soumyadip Ghosh
Shane G. Henderson

School of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY 14853, U.S.A.

ABSTRACT

The NORTA method for multivariate generation is a fast general purpose method for generating samples of a random vector with given marginal distributions and given product-moment or rank correlation matrix. However, this method has been shown to fail to work for some *feasible* correlation matrices. (A matrix is feasible if there exists a random vector with the given marginal distributions and the matrix as the correlation matrix.) We investigate how this feasibility problem behaves as the dimension of the random vector is increased and find the problem to become acute rapidly. We also find that a modified NORTA procedure, augmented by a semidefinite program (SDP) that aims to generate a correlation matrix "close" to the desired one, performs well with increasing dimension.

1 INTRODUCTION

Cario and Nelson (1997) described the NORTA method for generating random vectors with prescribed correlation matrix. This method belongs to a family of methods available for multivariate generation that address the specific problem of generating samples of a finite dimensional random vector such that the generated samples match a given set of marginal distributions for the individual components, and some measure of dependence between them, typically chosen to be either the product-moment or the rank correlation matrix. (The product-moment correlation matrix for a random vector $X = (X_1, \ldots, X_d)$ is the matrix $\Sigma_X = (\Sigma_X(i, j) : 1 \le i, j \le d)$ where

$$\Sigma_X(i, j) = \frac{\text{cov}(X_i, X_j)}{(\text{var}X_i \ \text{var}X_j)^{1/2}}.$$

The rank correlation matrix is of the same form except that now

$$\Sigma_X(i, j) = \frac{\text{cov}(F_i(X_i), F_j(X_j))}{(\text{var}F_i(X_i) \ \text{var}F_j(X_j))^{1/2}},$$

where F_i and F_j are the distribution functions of X_i and X_j respectively.)

The philosophy of specifying marginals and correlations to model dependent random variates is clearly an approximate one, since the joint distribution is not completely specified. However, almost all the methods that have been suggested to model and generate from the full joint-distribution suffer from some serious drawbacks, for instance the enormous amount of information needed to specify (and fit) the joint distribution, and the specific nature of these methods that precludes an easy adaptation to cases where a joint distribution of a different nature is to be modelled. These drawbacks make their use impractical for a model of even moderate complexity. Hence, by aiming for the simpler goal of matching only the marginal distributions and the correlation matrix, one hopes to capture the essence of the dependence between the components while being able to work with easily implementable methods that work well in higher dimensions.

Another argument in support of modelling random vectors in this way involves the use of diffusion approximations to model queueing systems. In many cases the limiting diffusions depend only on the first two moments of the input distributions. Therefore, there is some insensitivity in performance measures computed from these models to the exact form of the input distributions. In general then, if a form of this insensitivity is present in a model, the approach discussed here for modelling random vectors is quite reasonable.

The NORTA method involves a componentwise transformation of a multivariate normal random vector, and capitalizes on the fact that multivariate normal random vectors are easily generated; see e.g., Law and Kelton (2000), p. 480. Cario and Nelson (1997) traced the roots of the method back to Mardia (1970) who looked at bivariate distributions, and to Li and Hammond (1975) who concentrated on the case where all of the marginals have densities (with respect to Lebesgue measure). Iman and Conover (1982) implemented the same transformation procedure to induce a given rank correlation in the output. Their method is only approximate, in that the output will have only very approximately the desired rank correlation.

The NORTA method is a very efficient, easy to implement general purpose generation method, and has seen adaptations to various contexts. Clemen and Reilly (1999) described how to use the NORTA procedure to induce a desired rank correlation in the context of decision and risk analysis. Lurie and Goldberg (1998) implemented a variant of the NORTA method for generating samples of a predetermined size. Henderson, Chiera, and Cooke (2000) adapt the NORTA method to generate samples of dependent quasi-random numbers.

Due to its attactive properties, the NORTA procedure, *when it works*, is often the method of choice for generating random vectors with arbitrary marginals and any given feasible correlation matrix. It is thus natural to ask whether the NORTA procedure can match *any* feasible correlation matrix for a given set of marginals.

For 2-dimensional random vectors, the NORTA method can match any feasible correlation matrix. This follows immediately from the characterizations in Whitt (1976). However, this does not hold for dimensions 3 and greater. Both Li and Hammond (1975) and Lurie and Goldberg (1998) postulate examples in 3 dimensions where the NORTA procedure might fail for feasible correlation matrices, but do not establish that the counterexamples exist. In Ghosh and Henderson (2001a), Ghosh and Henderson (2001b), we give a computational procedure based on chessboard distributions to determine whether a given correlation matrix is feasible for the marginal distributions or not. Using this procedure, we rigorously established that such counterexamples do exist. Let us call feasible correlation matrices that cannot be matched using the NORTA method *NORTA defective* matrices.

Based on the numerical results obtained in Ghosh and Henderson (2001b) we had also conjectured that this feasibility problem might get steadily worse as the dimension increases, in the sense that the NORTA method would be increasingly likely to fail for feasible correlation matrices as the dimension of the matrices grew. We investigate this aspect of the feasiblity problem further in this paper. We estimate, for each dimension, the probability that the NORTA procedure fails to work for a feasible rank correlation matrix chosen uniformly from the set of all feasible correlation matrices. Kurowicka and Cooke (2001) also looked at this problem, but they work with a probability distribution that is not uniform over the set of all feasible correlation matrices. Our results confirm their finding that the probability the NORTA procedure fails to work grows rapidly with dimension.

Suppose we are willing to trade off accuracy for the sake of the efficiency of the NORTA generation procedure, i.e., we wish to use NORTA to generate a random vector with the prescribed marginals, and a correlation matrix that is, at least approximately, the required correlation matrix. In Ghosh and Henderson (2001b) we describe a semidefinite programming approach that can assist in this regard.

The proposed augmented NORTA method works in exactly the same manner as the original method unless a NORTA defective matrix is encountered. For such a matrix, a semidefinite program is set up and solved, and the results are then used to modify the inputs given to the NORTA generation step in the hope that the generated random vector has a correlation matrix that is "close" to the desired one (it has the same marginal distributions). The numerical results in Ghosh and Henderson (2001b) indicate that this is typically true for the 3-dimensional case. In this paper, we examine higher dimensions, exploring how the augmented NORTA method performs as the dimension increases. The results indicate that NORTA can typically get very close to a target correlation matrix, even in very high dimensions. So in high dimensions, while NORTA is unlikely to be able to exactly match a desired correlation matrix, it may be able to match a correlation matrix that is very close to the desired one.

The next section reviews the NORTA procedure and indicates why some matrices may be NORTA defective. Section 3 studies how the NORTA feasibility problem affects its performance as the dimension of the random vector is increased. Section 4 briefly describes the SDP augmentation proposed in Ghosh and Henderson (2001b), and studies how this augmented method performs in higher dimensions. Finally, Section 5 summarizes the conclusions that we were able to draw from our studies.

2 THE NORTA PROCEDURE

Suppose that we wish to generate i.i.d. replicates of a random vector $X = (X_1, \ldots, X_d)$ with prescribed marginal distributions

$$F_i(\cdot) = P(X_i \leq \cdot), i = 1, \ldots, d,$$

and product-moment or rank correlation matrix

$$\Sigma_X = \Sigma_X(i, j),\ 1 \leq i, j \leq d.$$

If we assume Σ_X to be feasible for the marginals, then the NORTA method generates i.i.d. replicates of X by the following procedure.

1. Generate an $\mathbf{R}^d$ valued joint normal random vector $Z = (Z_1, \ldots, Z_d)$ with mean vector 0 and covariance matrix $\Sigma_Z = (\Sigma_Z(i, j) : 1 \leq i, j \leq d)$, where $\Sigma_Z(i, i) = 1$ for $i = 1, \ldots, d$.

2. Compute the vector $X = (X_1, \ldots, X_d)$ via

$$X_i = F_i^{-1}(\Phi(Z_i)), \tag{1}$$

for $i = 1, \ldots, d$, where Φ is the distribution function of a standard normal random variable, and

$$F_i^{-1}(u) = \inf\{x : F_i(x) \geq u\}. \tag{2}$$

The vector X generated by this procedure will have the prescribed marginal distributions. To see this, note that each Z_i has a standard normal distribution, so that $\Phi(Z_i)$ is uniformly distributed on $(0, 1)$, and so $F_i^{-1}(\Phi(Z_i))$ will have the required marginal distribution.

The covariance matrix Σ_Z should be chosen, in a preprocessing phase, so that it induces the prescribed correlation matrix Σ_X on X. However, there is no general closed form expression that gives Σ_Z in terms of Σ_X. Indeed, determining the right Σ_Z is the most difficult step in implementing the NORTA method.

Each component of Σ_X has been shown to depend only on the corresponding component of Σ_Z. As in Cario and Nelson (1997), we can define $c_{ij}(z) = \Sigma_X(i, j)$ to represent the correlation between X_i and X_j as a function of the correlation z between Z_i and Z_j, when X_i and X_j are generated as in (1). Cario and Nelson (1997) show that under certain very mild conditions $c_{ij}(\cdot)$ is a non-decreasing, continuous function. This result helps us perform an efficient numerical search for a value $\Lambda_Z(i, j)$ that solves

$$c_{ij}(\Lambda_Z(i, j)) = \Sigma_X(i, j). \tag{3}$$

Hence a numerical estimate Λ_Z of Σ_Z can be determined by solving a number of one-dimensional root-finding problems. Unless stated otherwise, we assume that a solution exists for (3).

Henderson, Chiera, and Cooke (2000) also show that under stronger assumptions the value of z in (3) is uniquely determined by $\Sigma_X(i, j)$. They infer from this that if their assumptions hold and if NORTA *can* work, then it *will*.

The matrix Λ_Z is constructed in a way that does not necessarily ensure that it is positive semidefinite. It might indeed turn out to be indefinite, in which case it cannot be a valid covariance matrix for a joint normal distribution. Can this happen, i.e., can there exist a feasible correlation matrix that, under exact numerical estimation in (3), gives an indefinite Λ_Z?

Li and Hammond (1975) postulated the following counterexample. Suppose $X = (X_1, X_2, X_3)$ is a random vector with uniform $(0, 1]$ marginals, and correlation matrix

$$\Sigma_X = \begin{pmatrix} 1 & -0.4 & 0.2 \\ -0.4 & 1 & 0.8 \\ 0.2 & 0.8 & 1 \end{pmatrix}.$$

For this special case of uniform marginals, the equations (3) can be solved analytically as (Kruskal 1958)

$$\Lambda_Z(i, j) = 2 \sin\left(\frac{\pi}{6} \Sigma_X(i, j)\right). \tag{4}$$

The unique solution Λ_Z for the given Σ_X turns out to be indefinite.

This counterexample is of course valid only if such a uniform random vector X exists. Li and Hammond (1975) did not show this, and no general purpose method previously existed to determine the feasibility of a correlation matrix for a given set of marginals. We have since been able to develop a computational procedure in Ghosh and Henderson (2001b) that can determine, for *almost* any (in a Lebesgue measure sense) given correlation matrix, whether it is feasible for a given set of marginal distributions or not.

Applying this algorithm to the Li and Hammond example gives a construction of the random vector, so that it does, indeed, exist. In Ghosh and Henderson (2001b) we generate a number of such feasible matrices for three-dimensional uniform random vectors that are NORTA defective. The numerical results suggest a structure to the failure of NORTA. To explain this observation more carefully we need some notation.

Suppose that the marginal distributions $F_1, \ldots, F_d$ have densities with bounded support, and are fixed. We can, with an abuse of notation, view a $d \times d$ correlation matrix as an element of $d(d-1)/2$ dimensional vector space. This follows because there are $d(d-1)/2$ elements above the diagonal, the matrix is symmetric, and the diagonal elements are equal to 1. Let Ω denote the set of feasible correlation matrices. (Under the assumptions made on the marginals, the results hold identically for both rank and product-moment correlations, and hence no distinction will be made between them.) We view this set as a subset of $d(d-1)/2$ dimensional space. Ghosh and Henderson (2001b) prove that in this setting Ω is nonempty, convex, closed and full-dimensional.

Returning to the discussion above, we found that in 3 dimensions, NORTA defective matrices tended to occur near the boundary of Ω. Moreover, the indefinite correlation matrices Λ_Z determined for the joint normal distribution from (3) seemed to lie close to (but outside of) the set of symmetric positive semidefinite matrices. So NORTA defective matrices tended to occur near the boundary, and they were never too distant from a NORTA feasible matrix.

3 NORTA IN HIGHER DIMENSIONS

As mentioned above, NORTA appears to fail most often when the correlation matrix is close to the boundary of the set Ω. Now, in a sense that can be made precise, "most" points in certain sets in high dimensions lie close to the boundary. For example, consider the interior of the unit hypercube $[-\frac{1}{2}, \frac{1}{2}]^d$ in $\mathbb{R}^d$ represented by the hypercube $[-\frac{1-\epsilon}{2}, \frac{1-\epsilon}{2}]^d$. The ratio of the volumes of the interior and the whole set is then $(1-\epsilon)^d$, which decreases rapidly to 0 as d increases.

This suggests that feasible matrices within the set Ω may become increasingly likely to be NORTA defective as the dimension of the problem increases, so that the feasiblity problem that NORTA faces becomes increasingly acute as the dimension increases.

We investigate this dimensionality aspect of the NORTA feasiblity problem in the context of generating samples of a uniform random vector, i.e., a random vector with uniform $(0, 1]$ marginal distributions. This case has special significance to the NORTA method because, by construction, the method has to generate a uniform random vector as the first (intermediary) step. Furthermore, the rank correlation matrix of a NORTA generated vector with continuous marginal distributions coincides with the product moment correlation matrix for the intermediate uniform random vector.

This special case also has two advantages. First, the function c_{ij} is explicitly known; see (4). Hence any feasible correlation matrix for a uniform random vector can be easily tested for NORTA feasibility.

Second, it has recently been established (Kurowicka and Cooke 2001) that the set of all feasible correlation matrices for uniform marginals, say Ω, coincides with the set of all symmetric positive semidefinite matrices with ones on the diagonal. Thus the problem of estimating the probability of NORTA infeasibility reduces to the following algorithm.

1. Let $n \geq 1$ be given.
2. Let $\Sigma_X(1), \ldots, \Sigma_X(n)$ be an i.i.d. sample chosen uniformly from

$$\Omega = \{\Sigma : \Sigma = \Sigma^T, \Sigma \succeq 0, \ \Sigma_{ii} = 1 \quad \forall i\}.$$

3. For each $i = 1, \ldots, n$ let $\Lambda_Z(i)$ be obtained from $\Sigma_X(i)$ using the componentwise relation (4).
4. Estimate the probability of NORTA infeasibility by the proportion of matrices in $\{\Lambda_Z(i) : i = 1, \ldots, n\}$ that are not positive semidefinite.

(The matrix inequality $A \succeq 0$ signifies a constraint that the matrix A be positive semidefinite.)

Note that in estimating the probability of NORTA infeasibility we have had to choose a probability distribution on Ω. The uniform distribution (with respect to Lebesgue measure) is a natural choice, and is the one we prefer to work with. Kurowicka and Cooke (2001) also give estimates for the probability of NORTA feasibility but they use a different distribution on Ω.

3.1 Sampling Uniformly from Ω

Our first attempt at estimating the probability of NORTA infeasibility was to combine two well-known methods in simulation estimation: importance sampling and ratio estimation. We used importance sampling on the hypercube $[-1, 1]^{\frac{d(d-1)}{2}}$ (Ω is a strict subset of this hypercube) to choose correlation vectors from Ω. We then used ratio estimation (see, e.g., Henderson (2001)) to estimate the probability of NORTA infeasibility. The estimator of the probability of NORTA infeasibility was therefore of the form

$$\frac{\sum_{i=1}^{n}[I(\Sigma_X(i) \succeq 0, \Lambda_Z(i) \not\succeq 0)\frac{2^{-d}}{\phi(\Sigma_X(i))}]}{\sum_{i=1}^{n}[I(\Sigma_X(i) \succeq 0)\frac{2^{-d}}{\phi(\Sigma_X(i))}]},$$

where the matrices $\Sigma_X(i)$ were chosen independently with density ϕ from the hypercube $[-1, 1]^d$. We chose the density ϕ in a heuristic fashion.

This method of estimation worked well in lower dimensions but we found that it became excessively slow as the dimension increased. Indeed, it took more than two days to generate on the order of a thousand samples of positive definite matrices even for a dimension as low as $d = 12$. Clearly, a better sampling technique was needed.

Investing some further thought into the problem led us to construct a method that samples *exactly* from the uniform (in a Lebesgue measure sense) distribution on the set Ω. This method starts with the one-dimensional matrix [1] and then "grows out" the matrix to the dimension desired by successively adding an extra row (and the corresponding mirrored column) chosen from an appropriate distribution. To be more precise the method is as follows.

1. Let Σ be the 1×1 matrix 1.
2. For $i = 2, \ldots, d$
   (a) Let U be a column vector in $\mathbb{R}^{i-1}$ chosen, independently of all else, from distribution φ_i say.
   (b) Set
   $$\Sigma = \begin{bmatrix} \Sigma & U \\ U^T & 1 \end{bmatrix}.$$
   (c) Next i.

The distributions φ_i are conditional distributions that depend on the partial matrix Σ constructed thus far. We do not specify them further here.

This method has two key advantages over the first one. First, sampling from φ_i can be reduced to the problem of sampling from a univariate beta distribution, a very well-studied problem for which efficient algorithms are available (see Law and Kelton 2000, p. 453-458). Consequently this method scales very well with dimension. In our study we were able to generate samples consisting of many thousands of matrices up to dimension $d = 25$ in a matter of hours.

Second, this method does not involve a ratio-estimation step, which means that the estimation is more straightforward to implement. For a given sample size, we also found the results to be more accurate.

We used the exact sampling approach to estimate the probability of NORTA infeasibility for various dimensions. Our results are given in Figure 1, where the probability is plotted against dimension. The plot establishes that the feasiblity problem rapidly becomes acute as the dimension increases. As seen from the results, the probability of a matrix being NORTA defective is almost 1 even in as low a dimension as seventeen.

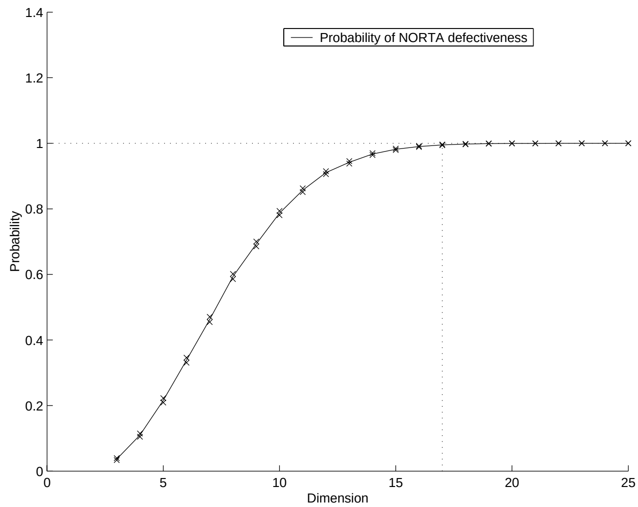

Figure 1: Probability of NORTA infeasibility, based on sampling 15,000 matrices uniformly from Ω in each dimension. Also shown are 95% confidence intervals.

4 FIXING NORTA

To recoup, we previously noted that NORTA defective matrices appear to lie close to the boundary of the set Ω. The reason matrices are NORTA defective is that the correlation matrix Λ_Z determined for the joint normal distribution turns out to be indefinite, and hence infeasible. Moreover, the results from the previous section confirm our intuition that since most points in a set lie near the boundary in higher dimensions, the NORTA infeasibility problem grows with dimension.

However, we also observed that the indefinite matrices Λ_Z lie very close to the set of feasible correlation matrices for joint normal random vectors (i.e., the set of positive semidefinite matrices with ones on the diagonal). This led to the suggestion in Ghosh and Henderson (2001b) that the setup stage of NORTA be augmented with an SDP that is used, if Λ_Z turns out indefinite, to find a matrix Σ_Z that is "close" to Λ_Z and is positive semidefinite. The matrix Σ_Z is then used within the NORTA method.

Why is this approach reasonable? In Theorem 2 of Cario and Nelson (1997) it is shown that under a certain moment condition, the output correlation matrix is a continuous function of the input covariance matrix Σ_Z used in the NORTA procedure. So if Σ_Z is "close" to Λ_Z, then we can expect the correlation matrix of the NORTA generated random vectors to be close to the desired matrix Σ_X. The moment condition always holds when we are attempting to match rank correlations, and we can expect it to hold almost invariably when matching product-moment correlations. Therefore, it is eminently reasonable to try and minimize some measure of distance $r(\Lambda_Z, \Sigma_Z)$ say, between Λ_Z and Σ_Z.

The SDP falls under the broad class of matrix completion problems; see Alfakih and Wolkowicz (2000), or Johnson (1990). For this case, given Λ_Z as data, we wish to choose a symmetric matrix Σ_Z to

$$\begin{array}{ll} \text{minimize} & r(\Sigma_Z, \Lambda_Z) \\ \text{subject to} & \Sigma_Z \succeq 0, \\ & \Sigma_Z(i,i) = 1. \end{array} \tag{5}$$

The metric $r(\cdot, \cdot)$ can be chosen as desired. In particular, choosing either the L_1 metric

$$r(A, B) = \sum_{i>j} |A_{ij} - B_{ij}|$$

or the L_∞ metric

$$r(A, B) = \max_{i>j} |A_{ij} - B_{ij}|$$

makes the minimization problem an SDP-constrained problem with a linear objective function. Efficient algorithms, and public domain codes implementing them, are available for solving semidefinite problems of this type; see Wolkowicz, Saigal, and Vandenberghe (2000).

The SDP framework allows us to include preferences on how the search for Σ_Z is performed. For example, we can require that for some (i, j), $\Sigma_Z(i, j) \geq \Lambda_Z(i, j)$, or that the value $\Lambda_Z(i, j)$ change by at most $\delta > 0$.

Numerical studies conducted in Ghosh and Henderson (2001b) indicate that in 3 dimensions this SDP augmentation yields NORTA generated random vectors with correlation matrices that are close to the desired ones. One might then ask whether this remains the case as the dimension increases.

We use a setting identical to that used in Section 3 for this study, and our measure of performance is the expected

L_1 distance that we have to move from the desired correlation matrix to reach a NORTA feasible one. This means that the minimization problem (5) is solved with $r(\cdot, \cdot)$ as the L_1 metric and no additional constraints are added.

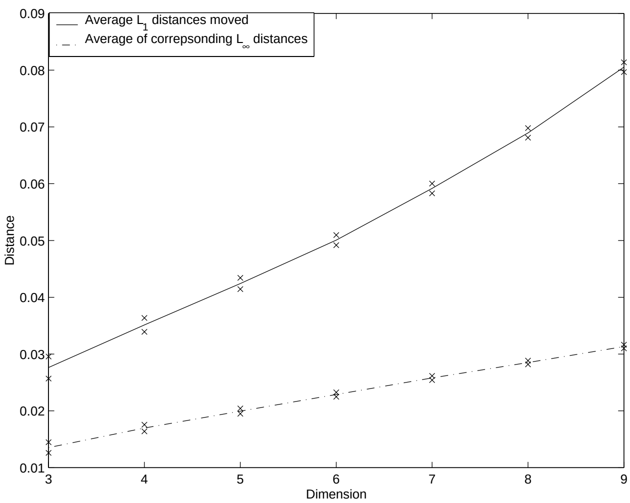

Figure 2: Performance of the SDP-augmented NORTA in higher dimensions. 15,000 matrices were generated uniformly from Ω and the semidefinite program, with r taken as the L_1 distance, solved for the NORTA defective cases. The solid line gives the expected L_1 distance with 95% confidence intervals as marked, with the average taken only over NORTA defective matrices. The dotted line gives the corresponding expected distance as measured in the L_∞ metric.

Figure 2 plots the results. We see that the expected L_1 distance increases as the dimension d increases at what might be perceived as a linear rate, although one could reasonably argue for a superlinear rate. If the rate of increase is indeed linear then, since there are $d(d-1)/2$ matrix entries above the diagonal, the *average* change per entry is (eventually) decreasing with dimension. Of course, it is possible that a small number of entries change by a large amount. The L_∞ distance is also shown, and we see that indeed, at least one entry is changed by an increasing amount as the dimension increases.

It might be preferable from a modelling standpoint to instead minimize the L_∞ distance, so that one tries to minimize the maximum deviation from the target correlations. The results in this case are shown in Table 1.

We see that the expected L_∞ distance appears to remain constant at around 0.005 or even decrease with dimension.

One might also attempt a hybrid of the L_1 and L_∞ approaches, perhaps by minimizing the L_1 distance subject to an upper bound on the L_∞ distance.

Table 1: The SDP-augmented NORTA in higher dimensions. For each dimension, 15,000 matrices were generated uniformly from Ω and the semidefinite program, with r taken as the L_∞ distance, solved for NORTA defective matrices. The second column gives the number of NORTA defective matrices encountered. The third column gives the point estimate, taken as an average only over the NORTA defective matrices, and not over all 15,000 matrices. The final column gives the halfwidth of 95% confidence intervals.

d	ND	L_∞	CI
3	524	0.0057	0.0004
4	1640	0.0053	0.0002
5	3271	0.0049	0.0001
6	4961	0.0045	0.0001
7	6988	0.0043	0.0001
8	8826	0.00414	0.00005
9	10428	0.00404	0.00004

Thus, the SDP-augmented NORTA problem performs well on average even in higher dimensions. It generates random vectors with correlation matrices which are close to the desired ones, while keeping changes to the individual correlations within reasonable limits.

5 CONCLUSIONS

We have empirically reached the following conclusions:

- The feasibility problem that the NORTA procedure faces becomes steadily worse with dimension. NORTA fails in the vast majority of cases even in as low a dimension as seventeen.
- The NORTA procedure, when augmented with the SDP optimization of Section 4, can generate samples with the required marginal distributions, and a correlation matrix that is a close approximation to the one desired, and the approximation remains accurate as the dimension increases.

An added bonus is the exact sampling procedure of Section 3.1 which can be generalized to sample uniformly from the set of all positive semidefinite matrices with diagonals fixed at specific values. We are presently working on refining this procedure and plan to publish the results elsewhere.

ACKNOWLEDGMENTS

This research was partially supported by National Science Foundation Grant Number DMI 9984717.

REFERENCES

Alfakih, A., and H. Wolkowicz. 2000. Matrix completion problems. In *Handbook of Semidefinite Programming: Theory, Algorithms and Applications*, ed. H. Wolkowicz, R. Saigal, and L. Vandenberghe, 533–545. Boston: Kluwer.

Cario, M. C., and B. L. Nelson. 1997. Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Technical report, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois.

Clemen, R. T., and T. Reilly. 1999. Correlations and copulas for decision and risk analysis. *Management Science* 45:208–224.

Ghosh, S., and S. G. Henderson. 2001a. Chessboard distributions. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 385–393. Piscataway NJ: IEEE.

Ghosh, S., and S. G. Henderson. 2001b. Chessboard distributions and random vectors with specified marginals and covariance matrix. *Operations Research*. To appear.

Henderson, S. G. 2001. Mathematics for simulation. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 83–94. Piscataway NJ: IEEE.

Henderson, S. G., B. A. Chiera, and R. M. Cooke. 2000. Generating “dependent" quasi-random numbers. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 527–536. Piscataway NJ: IEEE.

Iman, R., and W. Conover. 1982. A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics: Simulation and Computation* 11:311–334.

Johnson, C. R. 1990. Matrix completion problems: a survey. *Proceedings of Symposia in Applied Mathematics* 40:171–198.

Kruskal, W. 1958. Ordinal measures of associaton. *Journal of the American Statistical Association* 53:814–861.

Kurowicka, D., and R. M. Cooke. 2001. Conditional, partial and rank correlation for the elliptical copula; dependence modelling in uncertainty analysis. Technical report, Delft University of Technology, Mekelweg 4, 2628CD Delft, Netherlands.

Law, A. M., and W. D. Kelton. 2000. *Simulation Modeling and Analysis*. 3rd ed. New York: McGraw-Hill.

Li, S. T., and J. L. Hammond. 1975. Generation of pseudorandom numbers with specified univariate distributions and correlation coefficients. *IEEE Transactions on Systems, Man, and Cybernetics* 5:557–561.

Lurie, P. M., and M. S. Goldberg. 1998. An approximate method for sampling correlated random variables from partially-specified distributions. *Management Science* 44:203–218.

Mardia, K. V. 1970. A translation family of bivariate distributions and Fréchet’s bounds. *Sankhya* A32:119–122.

Whitt, W. 1976. Bivariate distributions with given marginals. *The Annals of Statistics* 4:1280–1289.

Wolkowicz, H., R. Saigal, and L. Vandenberghe. (Eds.) 2000. *Handbook of Semidefinite Programming: Theory, Algorithms and Applications*. Boston: Kluwer.

AUTHOR BIOGRAPHIES

SOUMYADIP GHOSH is a doctoral student in the School of Operations Research and Industrial Engineering at Cornell University. He has a Masters degree in Operations Research from the University of Michigan, Ann Arbor. He can be contacted at `<sdghosh@orie.cornell.edu>`.

SHANE G. HENDERSON is an assistant professor in the School of Operations Research and Industrial Engineering at Cornell University. He has previously held positions in the Department of Industrial and Operations Engineering at the University of Michigan and the Department of Engineering Science at the University of Auckland. He is an associate editor for the ACM Transactions on Modeling and Computer Simulation, Mathematics of Operations Research, and Operations Research Letters, and the newsletter editor for the INFORMS College on Simulation. His research interests include discrete-event simulation, queueing theory and scheduling problems. His e-mail address is `<henderson@orie.cornell.edu>`, and his web page is `<www.orie.cornell.edu/~shane>`.