

On greedy algorithms, partially ordered sets, and submodular functions

B. L. Dietrich
A. J. Hoffman

Recent developments in the use of greedy algorithms in linear programming are reviewed and extended. We find a common generalization of some theorems of Queyranne–Spieksma–Tardella, Faigle–Kern, and Fujishige about greedy algorithms for linear programs in diverse contexts. Additionally, we extend a well-known theorem of Topkis about submodular functions on the product of chains to submodular functions on the product of lattices.

1. Introduction

A *greedy algorithm* for solving an optimization problem successively maximizes the variables in some chosen order. Greedy algorithms are widely used in many contexts as a heuristic, with a hope and a prayer (sometimes even a proof) that the rapid calculation will yield an acceptable answer. In this essay, however, we are interested specifically in linear programming problems and in cases where the greedy algorithms actually produce correct answers.

While such cases are rare, they have been influential in linear programming history. A sympathetic interpretation of the 1781 paper by Monge [1] allows us to interpret a passage in that remarkable pamphlet as the description of a linear programming problem (actually of the important subclass known as the *transportation problem* [2]) which is amenable to a greedy solution. In fact, this notion was eponymized in [3] to celebrate its pedigree. A survey of the impact of the Monge concept is given in [4]. We provide a technical description later.

A second influence has been the use of the idea of greedy algorithms for linear programming to give a general context to some famous combinatorial algorithms. A leading example is the introduction by Edmonds [5] of the concept of *polymatroid*, and the interplay between greedy algorithms on polymatroids with greedy algorithms on matroids (which, in turn, generalize some algorithms for constructing minimum spanning trees for a graph). We should also mention in this spirit the notion of $(0, 1)$ matrices in which the 1's in each row (or in each column) occur consecutively. Although introduced [6] in the first

place to provide a greedy (and correct) solution to the *caterer problem* [7], it became an essential ingredient in many arguments used to prove combinatorial theorems using linear programming duality.

There is, unfortunately, no survey of the subject of greedy algorithms in linear programming more recent than [8], which is more than 15 years old. Our aim in this paper is to review an interesting development in greedy theory in a series of related papers appearing in the last few years, to describe the common themes, to offer generalizations of two of these themes, and to comment on weaknesses or lacunae remaining in our treatment. Central to this work being summarized and extended are the concepts (to be defined below) of partially ordered set and submodular functions.

2. Definitions; products of chains

A partially ordered set P is a set of elements $\{a, b, c, \dots\}$, together with a binary relation “ $<$ ” on some pairs such that $a < b$, $b < c$ imply $a < c$, and $a < a$ is never true. Also, $a > b$ means $b < a$. We say that a and b are *comparable* if $a = b$, $a < b$, or $b < a$. A chain is a partially ordered set in which any two elements are comparable. P is said to be a *lattice* if, for every pair of elements a, b , there is a unique least upper bound $a \vee b$ in P satisfying

$$c \geq a, c \geq b \text{ imply } c \geq a \vee b, \quad (1a)$$

and there is a unique greatest lower bound $a \wedge b$ in P satisfying

$$c \leq a, c \leq b \text{ imply } c \leq a \wedge b. \quad (1b)$$

We later introduce the more general notion of pseudolattice, but we postpone this to Section 3. In the meantime, we frequently use L to denote the partially ordered set with a lattice structure.

A real function r defined on a lattice L is said to be modular if $r(a \vee b) + r(a \wedge b) = r(a) + r(b)$; (2a)
submodular if $r(a \vee b) + r(a \wedge b) \leq r(a) + r(b)$; (2b)
supermodular if $r(a \vee b) + r(a \wedge b) \geq r(a) + r(b)$. (2c)

Let k chains $C_1, C_2, \dots, C_k$ be given, with respective cardinalities $c_1, c_2, \dots, c_k$. There is a natural partial ordering on the Cartesian product $C_1 \times \dots \times C_k$ that yields a lattice. A submodular function $r(\cdot)$ on this lattice is called a Monge array in case $k = 2$, and we use the same term for general k . Now consider the linear programming problem arising as follows: Construct a $(0, 1)$ matrix A with $c_1 c_2 \dots c_k$ rows and $c_1 + c_2 + \dots + c_k$ columns defined as follows: The rows correspond to the Cartesian product of the chains, and the columns to the union of the chains, with a 1 in a particular row and column if the element corresponding to that column is one of the multiplicands in the Cartesian product corresponding to the row. Let d be a given nonnegative vector with $c_1 + \dots + c_k$ coordinates corresponding to the columns of A , with $r(\cdot)$ being a submodular function (Monge array) on the product lattice $C_1 \times \dots \times C_k$ corresponding to the rows of A . Our linear programming problem is the following: Choose a nonnegative y that satisfies $y'A = d'$, and of all such y , choose one that minimizes (y, r) . It is shown in [9] that the following greedy algorithm solves the problem: Let $\mathcal{A}_t$ be the row of A corresponding to the greatest element (say t) of the product lattice, $\bar{d}$ the least entry of d corresponding to a 1 in $\mathcal{A}_t$. Let $y_t = \bar{d}$, replace d with $d - \bar{d}\mathcal{A}_t$, delete from A the column containing $\bar{d}$, and continue. This greedy algorithm terminates with an optimum solution, or the news that the problem is not feasible. (This news would be using up the rows of A without attaining a feasible solution.)

Next, consider the following variation of the polymatroid. Let r be a submodular function defined on the lattice of subsets of a finite set U , with $r(\emptyset) = 0$. Let A be the $(0, 1)$ matrix whose rows are the incidence vectors of the subsets of U , d a nonnegative vector with coordinates corresponding to elements of U . Consider the linear programming problem $y'A = d, y \geq 0$, minimize (y, r) . Again, a greedy algorithm starting with the greatest element in the lattice— U itself—solves the optimization problem.

In [10] it was pointed out that the validity of the greedy algorithm for polymatroids can be shown to be a

consequence of the validity of the greedy chain-product algorithm described above. (Among other interesting items in [10] is a proof that the chain-product algorithm is valid over any sublattice of the chain-product, and an allusion to the theorem of Topkis about submodular functions on chain-products; we comment on both later.)

The reasoning in [10] is that, in the polymatroid problem, we can replace each element u of U with a two-element chain. The greater element u' corresponds to including the element u in a subset; the lesser element u'' corresponds to omitting the element. This is how the subsets of U are coded. Of course, we need to choose values of the d vector to correspond to the different u'' , but it is easy to see how this is accomplished. Thus, [10] succeeds in folding this class of polymatroids into the multi-dimensional Monge array.

Further work [11–13] on topics related to those in [10] is discussed later.

3. Modular, consecutive (0, 1) functions

In this section, we present an alternative approach to the greedy algorithms discussed in Section 2.

Let U be a finite set and L a lattice, with least element m and greatest element M . We are given functions $f: L \rightarrow 2^U$ and $r: L \rightarrow \mathbb{R}$. We assume that

$$\text{if } f(a) = \emptyset, \text{ then } a = m \text{ and } r(m) = 0. \quad (3)$$

For each $u \in U$, let $f(L, u)$ be defined by

$$f(a, u) = 1 \text{ if } u \in f(a), 0 \text{ if } u \notin f(a). \quad (4)$$

We assume that the function f satisfies the following:

for all $u \in U$, $f(L, u)$ is modular,

and

$$\text{for all } a \neq b, f(a) \neq f(b). \quad (5)$$

In addition to the given functions f and r , we are also given a nonnegative function $d: U \rightarrow \mathbb{R}$, and we consider the following linear programming problem:

$$\text{minimize } \sum_{a \in L} r(a) x(a), \quad (6)$$

where

$$x(a) \geq 0 \quad \text{for all } a \in L,$$

and

$$\sum_{a \in L} x(a) f(a, u) = d(u) \quad \text{for all } u \in U.$$

For this problem, we propose a greedy algorithm GL . In the first step, we set $x(M)$ to be the minimum of all $d(u)$ such that $u \in f(M)$, and we let u^* be one of the (possibly several) u s where the minimum is attained. If we delete from L all elements a such that $u^* \in f(a)$, subtract

$d(u^*)$ from all $d(u)$ such that $u \in f(M)$, and delete u^* from U , it is observed from (3), (4), and (5) that the remaining elements form a lattice, and these same stipulations still hold for the remaining lattice and remaining U , on which d is still a nonnegative function. We continue this algorithm indefinitely, until we are compelled to stop. We may stop because the remaining d still has at least one positive coordinate, but the remaining part of L is empty or consists just of m with $f(m) = 0$, or we may stop because d has been reduced to 0.

Theorem 1

Assume that (6) is feasible. Assume also that (3), (4), and (5) hold. Then GL produces an optimum solution to (6) if the following additional conditions hold:

$$r \text{ is submodular,} \quad (7)$$

$$f(a) \supset f(b) \text{ implies } a > b, \quad (8)$$

$$a < b < c, f(a, u) = f(c, u) = 1 \text{ imply } f(b, u) = 1. \quad (9)$$

The reader should note that the *consecutivity condition* (9) from [6] is significant in lattice polyhedra generally (see [8, 14–16]). What is different here from general lattice polyhedra is (8), which specifies relations among the various $f(L, u)$ for different u , whereas we are usually concerned just with properties of $f(L, u)$ for each u by itself. But assumption (8) is unavoidable: In fact, if any one of (7), (8), or (9) is false, there is an instance of (6) in which the greedy algorithm GL does not produce an optimum solution.

Proof

To prove the theorem, all we need prove is the following:

There exists a feasible solution x^* to (6) for which

$$x^*(M) = d(u^*). \quad (10)$$

The reason is that induction will then justify every step of GL .

By the hypothesis of the theorem, (6) is feasible. Further, the feasible region is bounded, so an optimum solution to (6) exists. Let x be an optimum solution in which $x(M)$ is as large as possible, but $x(M)$ is less than $d(u^*)$. This means that, for each $u \in f(M)$, there exists at least one $a \in L$, $a < M$, such that

$$x(a) > 0 \text{ and } u \in f(a). \quad (11)$$

Let a be a maximal element of the partially ordered set L satisfying (11). By (8) and (5), there is at least one $u \in f(M)$ (say u'), that is not in $f(a)$. Thus, there exists at least one element $b \in L$, $b \neq a$, so that

$$u' \in f(b), u' \notin f(a). \quad (12)$$

Suppose that b and a are comparable. Then, by the definition of a , $b < a$, and we would have, with $M > a > b$, and (12), a violation of (9). Hence, a and b are incomparable. Let ϵ be a positive number not greater than the minimum of $x(b)$ and $x(a)$. Reduce $x(a)$ and $x(b)$ by ϵ , and raise $x(a \wedge b)$ and $x(a \vee b)$ by ϵ . By (5), the new x is feasible; by (7), the new x is optimal. If $M = a \vee b$, we have contradicted the definition of x . If $a \vee b \neq M$, we still have $a \vee b > a$, and have contradicted the definition of a . Hence, we have proved (10) and the theorem.

It remains to show that each of (7), (8), and (9) is needed. The necessity of (7) is obvious, and indeed was pointed out in [17]. For (9), let L be the chain $a < b < c$, $U = \{1, 2, 3\}$, $f(a) = \{1, 2\}$, $f(b) = \{1, 3\}$, and $f(c) = \{2, 3\}$. Let $d(1) = d(2) = d(3) = 1$, and let $r(a) = r(b) = r(c) = 1$. Then (7) and (8) are true, (9) is not, and GL does not produce an optimal solution; it does not even produce a feasible solution, though one exists.

For (8), let L , U , r , and d be as above. Let $f(a) = \{1, 2\}$, $f(b) = \{1, 2, 3\}$, $f(c) = \{3\}$. Let $r(a) = r(b) = r(c) = 1$. Then GL does not produce an optimal solution, even though (7) and (9) are true. $\square$

Observe that the theorem clearly covers the case of polymatroids. It also covers the case of chain-products, as we see in the next section. Besides its conclusions, our theorem has the feature of divorcing (partly, but not totally) the lattice L from the set U . In the other examples, and others to be cited later, U comes from the definition of L . However, our separation is not complete because of conditions (5) and (8).

4. Products of lattices

In this section we consider, more generally than the product of chains in [10], the product of lattices and the extension of the results in Section 3 to such products.

Theorem 2

Let $L_1, \dots, L_k$ be lattices; $U_1, \dots, U_k$ disjoint sets; $f_i: L_i \rightarrow U_i$, $i = 1, \dots, k$, each satisfying the conditions in Section 3. If we define the product lattice $L \equiv L_1 \times \dots \times L_k \equiv \{(a_1, \dots, a_k) : a_i \in L_i, i = 1, \dots, k\}$, with $a \leq b$ if and only if $a_i \leq b_i$ for all i ; and we define

$$f \equiv (f_1, \dots, f_k) \rightarrow U = U_1 \cup \dots \cup U_k$$

by

$$f(a_1, \dots, a_k) = f(a_1) \cup \dots \cup f(a_k),$$

we will have satisfied the conditions of Section 3 for f , L , and U .

If, in addition, we introduce a submodular $r: L \rightarrow R$, we have extended the material on chain-products in [10] to lattice products. It is interesting to observe, however, some facts about submodular r defined on lattice products.

Topkis [18] proved that for a product $C = C_1 \times \cdots \times C_k$ of chains, r is submodular on C if

$$r(a \vee b) + r(a \wedge b) \leq r(a) + r(b) \quad (13)$$

holds when $a = (a_1, \dots, a_k)$ and

$$b = (b_1, \dots, b_k) \text{ differ in at most two coordinates.} \quad (14)$$

We will prove Theorem 3.

Theorem 3

Let $L = L_1 \times \cdots \times L_k$, $r: L \rightarrow \mathbb{R}$. Let $a = (a_1, \dots, a_k)$ and $b = (b_1, \dots, b_k)$. Assume that (13) holds in the following cases:

For each i , a_i and b_i are comparable, and

$$a_i = b_i \text{ with two exceptions,} \quad (15)$$

and

$$a_i = b_i \text{ with one exception.} \quad (16)$$

Then $r: L \rightarrow \mathbb{R}$ is submodular.

Note that if L is the product of chains, (16) is immediate, and (15) is the same as (14). We think it interesting that proving r submodular on L needs only (15) when L is restricted to be the product of two element chains.

To prove that (15) and (16) imply (13), we first consider the case $k = 2$. Assume $a = (a_1, a_2)$, $b = (b_1, b_2)$, and define

$$c_1 = (a_1 \vee b_1), d_1 = (a_1 \wedge b_1);$$

$$c_2 = (a_2 \vee b_2), d_2 = (a_2 \wedge b_2).$$

By (16),

$$\begin{aligned} r(a_1, c_2) + r(a_1, d_2) &\leq r(a_1, a_2) + r(a_1, b_2), \\ r(b_1, c_2) + r(b_1, d_2) &\leq r(b_1, a_2) + r(b_1, b_2), \\ r(c_1, a_2) + r(d_1, a_2) &\leq r(a_1, a_2) + r(b_1, a_2), \\ r(c_1, b_2) + r(d_1, b_2) &\leq r(a_1, b_2) + r(b_1, b_2). \end{aligned} \quad (17)$$

By (15),

$$\begin{aligned} r(c_1, c_2) + r(a_1, b_2) &\leq r(c_1, b_2) + r(a_1, c_2), \\ r(c_1, c_2) + r(b_1, a_2) &\leq r(c_1, a_2) + r(b_1, c_2), \\ r(d_1, d_2) + r(a_1, b_2) &\leq r(d_1, b_2) + r(a_1, d_2), \\ r(d_1, d_2) + r(b_1, a_2) &\leq r(d_1, a_2) + r(b_1, d_2). \end{aligned} \quad (18)$$

Adding one half of (17) and (18) yields (13).

We now consider the case of general $k \geq 3$, assuming (inductively) that the theorem has been proved for $k - 1$. Let us consider $L = L_1 \times \cdots \times L_k$ as $L = L_1^* \times L_k$, where $L_1^* = L_1 \times \cdots \times L_{k-1}$. This brings us back to the case of two lattices, but we must verify the hypothesis.

First, for fixed a_k , r is a submodular function on L_1^* by induction. For fixed $a \in L_1^*$, r is a submodular function on L_k by the original hypothesis. Next, (15) holds, from the theorem of Topkis. Thus, the induction step is complete.

5. Pseudolattices

In [11–13], further examples in the spirit of [10] were presented. To explain them, we first define the following: In a partially ordered set P , we say that b covers a if $a < b$ and there is no c such that $a < c < b$. Next, an ideal $\mathcal{I}$ of P is a subset such that $a \in \mathcal{I}$, $b < a$ imply $b \in \mathcal{I}$. Observe that the union (and also the intersection) of two ideals is an ideal. An antichain $\mathcal{A} \subset P$ is a set of mutually incomparable elements. If $\mathcal{A}$ is an antichain, then $[\mathcal{A}]$ denotes the ideal of all $b \in P$ such that $b \leq a$ for some $a \in \mathcal{A}$. If $\mathcal{I}$ is an ideal, $\mathcal{I}^+$ is the antichain consisting of the maximal elements of $\mathcal{I}$.

Let a partially ordered set P be given, and consider the lattice $L(P)$ of all ideals contained in P . $[L(P)]$ is itself a partially ordered set, ordered by inclusion, with union and intersection respectively the operations $\vee$ and $\wedge$. Let $\mathcal{U}$ be the set of elements of P , and define a mapping $f: L(P) \rightarrow 2^{\mathcal{U}}$ by

$$f(\mathcal{I}) = \mathcal{I}^+. \quad (19)$$

It is shown in [11] that, when P is a forest, a suitable generalization of the greedy algorithm described in [10] is valid. Note that, when P is a forest, the mapping (19) satisfies the hypothesis (and conclusion) of Theorem 1 above. If P is not a forest, then the greedy algorithm is not in general valid. But [12] points out that an additional hypothesis about the submodular function $r: L(P) \rightarrow \mathbb{R}$, if e is covered by at least two elements, then

$$r(\mathcal{I} \cup \{e\}) \geq r(\mathcal{I}) \quad (20)$$

is sufficient to justify the greedy algorithm even when P is not a forest. An interpretation of this result in terms of submodular flows is given in [13]. The key idea in both papers to establish that, if $\mathcal{I}_1$ and $\mathcal{I}_2$ are ideals in P , then the antichain

$$\mathcal{A} \equiv f(\mathcal{I}_1) + f(\mathcal{I}_2) - f(\mathcal{I}_1 \cup \mathcal{I}_2) \text{ is an antichain,} \quad (21)$$

and

$$[\mathcal{A}] \subset \mathcal{I}_1 \cap \mathcal{I}_2. \quad (22)$$

Further, every element in

$$(\mathcal{I}_1 \cap \mathcal{I}_2) - [\mathcal{A}] \text{ is covered by at least two elements.} \quad (23)$$

From (20), (22), and (23), we infer that $r([\mathcal{A}]) \leq r(\mathcal{I}_1 \cap \mathcal{I}_2)$, hence

$$r(\mathcal{I}_1 \cup \mathcal{I}_2) + r([\mathcal{A}]) \leq r(\mathcal{I}_1) + r(\mathcal{I}_2). \quad (24)$$

Instead of the arguments used in [11] and [12], we propose to include the result of [11] in the framework of our approach in Theorem 1 by slightly extending the definitions used therein. Note that Theorem 1, as it stands, cannot encompass [12] because (5) is false. On the other hand, if one wants to exploit (20)–(24), we have the difficulty that, in (20), $[\mathcal{A}]$ is not $\mathcal{F}_1 \cap \mathcal{F}_2$.

Our tactic is to extend Theorem 1 by using the notion of *pseudolattice*. A partially ordered set P is a pseudolattice if, for every pair of elements (a, b) there is a designated element $a \vee^* b$ satisfying $a \vee^* b \geq a, b$, and there is a designated element $a \wedge^* b$ satisfying $a \wedge^* b \leq a, b$.

Further, if a and b are comparable, $a \vee^* b = \max(a, b)$, $a \wedge^* b = \min(a, b)$.

Theorem 4

If, in Theorem 1, the term *pseudolattice* and the *pseudolattice operations* $\vee^*$ and $\wedge^*$ are substituted for the term *lattice* and the *lattice operations* $\vee$ and $\wedge$, then the theorem remains valid.

The proof is simply to revisit the proof of Theorem 1 and observe that it remains valid in the more general setting without changing a single word. Further, with this change, we accommodate [12] by observing that when we define

$$\mathcal{F}_1 \vee^* \mathcal{F}_2 \equiv \mathcal{F}_1 \cup \mathcal{F}_2$$

and

$$\mathcal{F} \wedge^* \mathcal{F}_2 \equiv [\mathcal{A}], \quad (25)$$

where $\mathcal{A}$ is defined by (21), then (5) holds along with all of the other hypotheses of Theorem 4, so [12] is included in the general setting.

6. Some questions

1. It is easy to see that in the partially ordered set $m < a, b, c < M$ with $\{a, b, c\}$ mutually incomparable, no function f satisfying our hypotheses is possible. It would be nice to know what partially ordered sets permit functions f .
2. A theory for combining greedy linear programming problems on polytopes of the form $\{x | \mathcal{A}x \leq b, x \geq 0\}$ is presented in [17]. We have no appropriate theory if the polytopes are of the form $\{x | \mathcal{A}x = b, x \geq 0\}$.

Acknowledgments

We thank Don Coppersmith, Milind Dawande, Jon Lee, and George Nemhauser for their contributions to this paper.

References

1. G. Monge, “Déblai et Remblai,” *Mem. de l’Academie des Sciences* (1781).
2. F. L. Hitchcock, “The Distribution of a Product from Several Sources to Numerous Localities,” *J. Math. Phys.* **20**, 224–230 (1941).
3. A. J. Hoffman, “On Simple Linear Programming Problems,” *Convexity: Proceedings of the Seventh Symposium in Pure Mathematics*, V. Klee, Ed., *Proceedings of Symposia in Pure Mathematics*, Vol. 7, American Mathematical Society, Providence, RI, 1963, pp. 317–327.
4. R. E. Burkard, B. Klinz, and R. Rodolf, “Perspectives of Monge Properties in Optimization,” *Discrete Appl. Math.* **70**, 95–161 (1996).
5. J. Edmonds, “Matroids and the Greedy Algorithm,” *Math. Program.* **1**, 127–136 (1971).
6. J. W. Gaddum, A. J. Hoffman, and D. Sokolowsky, “On the Solution to the Caterer Problem,” *Nav. Res. Logist. Quart.* **1**, 223–229 (1954).
7. W. Jacobs, “The Caterer Problem,” *Nav. Res. Logist. Quart.* **1**, 154–165 (1954).
8. A. J. Hoffman, “On Greedy Algorithms That Succeed,” *Surveys in Combinatorics*, I. Anderson, Ed., Cambridge University Press, Cambridge, England, 1985, pp. 97–112.
9. W. W. Bein, P. Brucker, J. K. Park, and P. K. Pathak, “A Monge Property for the d -Dimensional Transportation Problem,” *Discrete Appl. Math.* **58**, 97–109 (1995).
10. M. Queyranne, F. Spieksma, and F. Tardella, “A General Class of Greedily Solvable Linear Programs,” *Math. Oper. Res.* **23**, 892–908 (1998).
11. U. Faigle and W. Kern, “Submodular Linear Programs on Forests,” *Math. Program.* **72**, 195–206 (2000).
12. U. Faigle and W. Kern, “On the Core of Ordered Submodular Cost Games,” *Math. Program.* **87**, Ser. A, 483–489 (2000).
13. S. Fujishige, “A Note on Faigle and Kern’s Dual Greedy Polyhedra,” *Math. Program.* **88**, Ser. A, 217–220 (2000).
14. H. Gröflin and A. J. Hoffman, “Lattice Polyhedra II: Construction and Examples,” *Ann. Discrete Math.* **15**, 189–203 (1982).
15. A. J. Hoffman and D. E. Schwartz, “On Lattice Polyhedra,” *Proceedings of the 5th Hungarian Conference in Combinatorics*, A. Hajnal and V. T. Sos, Eds., 1976, North-Holland, Amsterdam, 1978, pp. 593–598.
16. A. J. Hoffman, “On Lattice Polyhedra III, Blockers and Anti-blockers of Lattice Clusters,” *Polyhedral Combinatorics, Math. Programming Stud.* **8**, 197–207 (1978).
17. W. W. Bein, P. Brucker, and A. J. Hoffman, “Series Parallel Composition of Greedy Linear Programming Problems,” *Math. Program.* **62**, Ser. B, 1–14 (1993).
18. D. M. Topkis, “Minimizing a Submodular Function on a Lattice,” *Oper. Res.* **26**, 305–321 (1978).

Received February 20, 2002; accepted for publication August 29, 2002

Brenda L. Dietrich IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (dietric@us.ibm.com). Dr. Dietrich joined the IBM Research Division in 1984, after completing studies for her Ph.D. in operations research and industrial engineering at Cornell University. She is currently Department Group Manager of Mathematical Sciences. Her research includes work in manufacturing modeling and scheduling, inventory management, transportation logistics, mathematical programming, and combinatorial optimization. She became the Manager of the Manufacturing Logistics group in 1991 and served on the Research Division technical planning staff in 1995. She founded the Logistics Application group in 1995 and became Manager of the Optimization Center in 1997. Dr. Dietrich is a member of the Advisory Board of the Department of Industrial Engineering and Management Sciences (IE/MS) at Northwestern University, a member of the Industrial Advisory Board for both the IMA (University of Minnesota) and DIMACS (Rutgers), and a member of the IBM Academy of Technology; and she has served on the Board of INFORMS. Dr. Dietrich holds eight patents; she is the author or coauthor of numerous publications, and she is coeditor of the book *Mathematics of the Internet, E-Auction and Markets*.

Alan J. Hoffman IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (ajh@us.ibm.com). Dr. Hoffman received A.B. and Ph.D. degrees from Columbia University and worked at the Institute for Advanced Study (Princeton), National Bureau of Standards, Office of Naval Research (London), and General Electric Company prior to joining IBM in 1961. He retired this year, and is now IBM Fellow emeritus. He is a member of the National Academy of Sciences, and a Fellow of the American Academy of Arts and Sciences and the New York Academy of Sciences. He has been an adjunct or visiting professor at the Israel Institute of Technology (Technion), City University of New York, Yale University, Stanford University, Rutgers University, and Georgia Institute of Technology; and he has fourteen Ph.D. students. He was the founding editor of *Linear Algebra and Its Applications*, and has served on the editorial boards of *Discrete Mathematics*, *Discrete Applied Mathematics*, *Naval Research Logistics Quarterly*, *Journal of Combinatorial Theory*, *Combinatorica*, *SIAM Journal of Discrete Mathematics*, *SIAM Journal of Applied Mathematics*, and *Mathematics of Computation*. Dr. Hoffman was Phi Beta Kappa lecturer in 1989–1990, and he was honored by a 65th birthday special issue of *Linear Algebra and Its Applications*. He was a co-winner of the von Neumann Prize of INFORMS in 1992, and the Founder's Award of the Mathematical Programming Society in 2000. He also holds an honorary D.Sc. degree from the Israel Institute of Technology (Technion). The book *Selected Papers of Alan J. Hoffman*, edited by Charles Micchelli, will be published in 2003 by World Scientific Publishing.