Defendens Imperium Romanum: A Classical Problem in Military Strategy

Charles S. ReVelle; Kenneth E. Rosing

The American Mathematical Monthly, Vol. 107, No. 7. (Aug. - Sep., 2000), pp. 585-594.

Stable URL:
http://links.jstor.org/sici?sici=0002-9890%28200008%2F09%29107%3A7%3C585%3ADIRACP%3E2.0.CO%3B2-2

The American Mathematical Monthly is currently published by Mathematical Association of America.

Defendens Imperium Romanum*: A Classical Problem in Military Strategy

Charles S. ReVelle and Kenneth E. Rosing

INTRODUCTION AND PROBLEM DESCRIPTION. In the third century of the Common Era (CE), when Rome dominated not only Europe, but also North Africa and the Near East, it was able to deploy fifty legions throughout the empire. In this *forward defense* strategy even the furthermost areas of the empire were secured by the on-site presence of an adequate number of legions of the Roman army. However, the empire had lost much of its muscle by the fourth century CE and the forces of Rome had diminished to only about twenty-five legions. It had thus become impossible to station legions in sufficient strength at all of the forward positions of the empire without abandoning the core.

A new *defense in depth* strategy was devised by the Emperor Constantine (Constantine The Great, 274–337) to cope with the reduced power of the empire [**11**]. His defense in depth used local troops to disrupt invasion and deployed mobile Field Armies (FAs) to stop and throw back the intruding enemy, or to suppress insurrection. The earlier *forward defense* strategy had provided a wall around the empire denying any but the most modest of incursions—it even allowed Roman forces to sally into barbarian lands to disrupt invasions as they were being mounted. In place of Rome's forces, the defense in depth strategy substituted local part-time militias (who would be fighting for their own land and families) to slow and fragment any invading barbarian army until the heavier weight of an FA, dispatched from a distant area, could be brought to bear.

Each set of roughly six legions with ancillary cavalry, artillery, etc. forms an FA, a unit of forces whose numbers are sufficient to secure any one of the regions of the empire [**2**]. In the third century CE, Rome's fifty legions or about eight FAs could be allocated so that each of the eight provinces was secured by its own FA. However, by the fourth century CE, only four FAs were available for deployment. The regions of the empire are considered to be connected as shown in Figure 1, where each region is represented as a circle (node). Movement along a line (edge) between regions (nodes) represents a "step" and for a region to be *securable*, an FA must be able to reach it in just one step [**2**].

A region is considered to be *secured* if it has one or more FAs stationed in it already. On the other hand, the region may be *securable*—that is, an FA may be capable of deploying to protect that region in a single step, but only under a special condition. An FA can be deployed from one region to an adjacent region *only* if it moves from a region where there is at least one other FA to help launch it. This is analogous to the island-hopping strategy pursued by General MacArthur in World War II in the Pacific Theater—movement followed the chain of islands already secured by troops left behind.

*Defending the Roman Empire

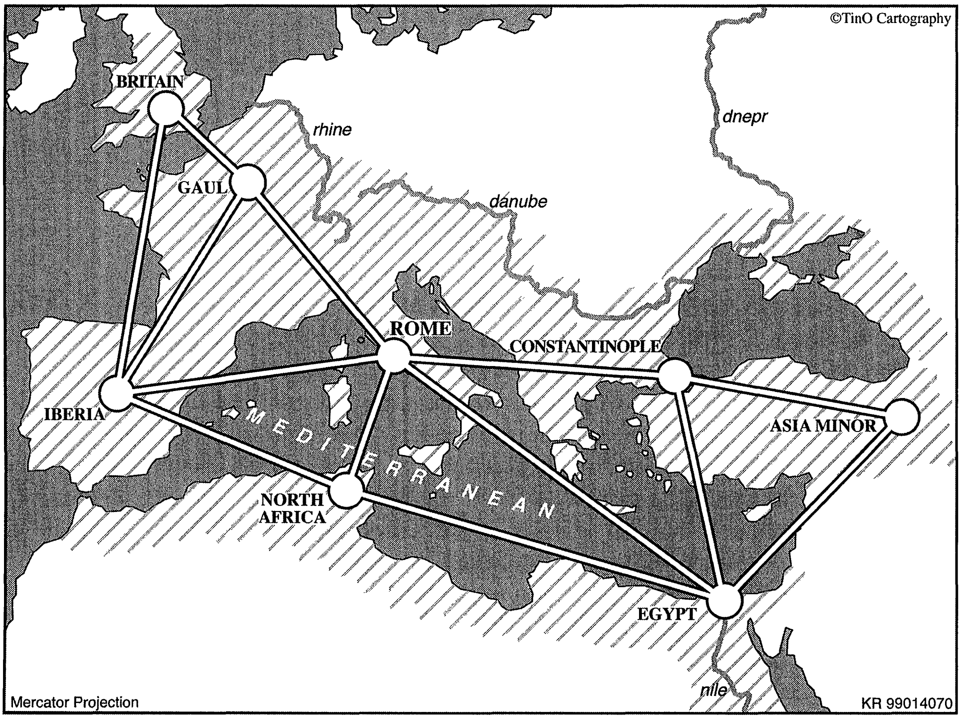

Figure 1. The Empire of Constantine and its Eight Provinces

The challenge for Constantine was to allocate just four FAs to positions in the eight regions of the empire. Constantine chose to place two at Rome, a symbolic as well as strategic choice, and two at his new capital, Constantinople. With this deployment, each region of the empire was either already secured or could be reached by an FA in just one step—with the exception of Britain. To reach Britian with an FA required an FA to move from Rome to Gaul, securing Gaul. Then, a second FA needed to be launched from Constantinople to Rome. Only then could an FA shift from Rome to Gaul, and finally an FA could move from Gaul to Britian, a total of four steps for the field armies. Each of the four steps began from a base that had two FAs present.

Here is another alternative, not necessarily better than Constantine's strategy, but illustrative of the moves that are legitimate. Suppose we place one FA in Gaul, two in Rome, and one in Constantinople. Britian can now be reached in two steps, which consist of moving an FA from Rome to Gaul and moving an FA from Gaul to Britian—better for Britian than before. However, Asia Minor is now *not* reachable in one step, but requires two steps: Rome to Constantinople and Constantinople to Asia Minor. All the rest of the empire is just one step away. It is not clear that this is better than Constantine's strategy. Although the number of steps to the worst-off region has been reduced to two, the number of regions that are more than one step away from any assistance has gone from one to two.

Can a modern analyst do better than Constantine's solution? "Better" may be measured with respect to several criteria. One criterion is the number of regions that cannot be reached in a single step; this number is to be reduced. For Constantine's solution, as we have just seen, that number is one. Another criterion is the number of steps it takes to reach the worst-off node. Again, for Constantine's choice, this number is four, the number of moves needed to reach Britian with assistance. To do better than Constantine, one would need to do better with respect to one of the criteria without degradation of the other. For example, relative to the first criterion, an improvement would be to reach *every* region in just one step. Alternatively, if one can reduce the number of steps to reach the worst-off node *without* increasing the number of regions that are more than one step away, the solution will be better. That is, if an analyst can keep the number of nodes that can't be reached in one step to just one, and can reduce the maximum number of steps to reach the worst-off node to *fewer than* four, then the solution will be better than Constantine's. And of course, if all regions were made either initially secure or reachable in one step, then the empire would be fully protected.

Still a third criterion concerns the consequences of a second war occurring somewhere in the empire. We might want to minimize the number of regions that can't be reached or secured in the event of a second war or, conversely, we might want to maximize the number of regions secured or covered in one step in the event of a second war.

We may envision the fundamental problem as having at least two phases. In the first phase, the number of FAs required may be an unknown, and we ask the question, "What is the least number of FAs to be placed, and where should they be sited, so that all regions of the empire are either secured or securable?" In the second phase, we ask "How should a limited number of FAs (the number available may be insufficient to secure or make securable all of the regions) be deployed to achieve optimally some security objective such as the maximum number of regions made secure or securable?" We present several models that address these criteria in the remainder of this paper.

BASIC FORMULATIONS. The formulations we present next belong to the class of $0, 1$ optimization problems that seek yes-or-no location decisions at discrete points. These discrete siting problems, which most often consider distances measured on a network, constitute one of two classes of location problems. The other class typically chooses sites from an infinite space of alternative locations and often makes use of an Euclidean or other distance measure, Together, these two classes of problems make up the family of problems referred to as Location Science or Topothesiology. Several current texts and collections survey these problems and the methods used to solve them; see **[6]**, **[7]**. or **[10]**.

The Set Covering Deployment Problem. The first formulation is called the Set Covering Deployment Problem (SCDP). This formulation is a novel derivative of a well-known problem, the Location Set Covering Problem (**[16]**, **[15]**, and **[13]**). In the Location Set Covering Problem with demand nodes and eligible facility sites scattered on a network or a plane, the problem is to find, and site, the least number of facilities so that all points of demand have at least one facility within some distance standard.

In the SCDP formulation, each region must either be *secured* by one or more FAs or *securable* by an FA that can reach the region in a single step from a two-FA region. We seek the least number of FAs to distribute among the regions so that all regions are either secured or securable.

We let:

I = the set of demand areas/deployment sites;

$x_i = 1, 0$; the variable is 1 if region i contains one or two FAs and 0 otherwise;

$y_1 = 1, 0$; the variable is 1 if region i contains two FAs, and 0 otherwise;
$N_i = \{j | \text{region } j \text{ is one step from region } i\}$

The objective is

$$\text{Minimize } z = \sum_{i \in I} x_i + \sum_{i \in I} y_i$$

Subject to:

$$x_i + \sum_{j \in N_i} y_j \geq 1 \qquad \text{for every } i \in I \tag{1}$$

$$y_i \leq x_i \qquad \text{for every } i \in I \tag{2}$$

The constraints (1) say that every region must be secured or securable. If $x_i = 1$, then region i is secured—without regard to the presence of any regions one step away that contain two FAs. If $x_i = 0$, then at least one of the regions one step away from region i must have two FAs—this to make region i securable in a single step. Of course, it is possible that for some i, $x_i = 1$ and also there is at least one region one step away from region i that has two FAs. Nonetheless, one of these events is required to occur for each $i \in I$. The constraints (2) say that there cannot be two FAs in a region unless there is at least one FA in the region.

The objective function of the SCDP is the number of regions with one or two FAs plus the number of regions with two FAs. This is exactly the number of FAs deployed. For example, suppose there is just one FA in region A and two FAs in region B. The value of the first term of the objective function is two since both regions A and B have either one or two FAs. The value of the second term is one since only region B has two FAs. The objective function is $2 + 1 = 3$. Minimizing the objective function minimizes the number of FAs deployed subject to the constraints.

If it were felt that an FA could not be isolated (i.e., must have at least one other FA within one step), then we could modify the problem by adding an additional constraint:

$$\sum_{j \in N_i} x_j \geq x_i \qquad \text{for every } i \in I \tag{3}$$

The Maximal Covering Deployment Problem. There is also a natural problem that is complementary to the SCDP. Its goal is to allocate a limited number of FAs to the regions in such a way as to maximize the number of regions that are securable in a single step or already secured by the presence of an FA. We call this problem the Maximal Covering Deployment Problem (MCDP). The MCDP is also a relative of a widely known problem, the Maximal Covering Location Problem [**3**]. The Maximal Covering Location Problem resembles the Location Set Covering Problem in that it assumes that demand nodes and eligible facility sites are dispersed on the plane or network. In contrast to the Location Set Covering Problem, it seeks to site a limited number of facilities in such a way that the greatest number of demand nodes has one or more facilities within a distance standard.

The logical constraints of the MCDP fix in advance the number of FAs, but say that an area is "covered" only if it has an FA present on site (it is secured) or if there is at least one position only one step away that has two FAs located there (it is securable). Coverage of every region is not required but is sought as a goal. One new type of variable needs to be introduced for the MCDP. It is:

$u_i = 1, 0$: the variable is 1 if region i is secured or securable in a single step. it is 0 otherwise.

The problem as structured can be stated as:

$$\text{Maximize } Z = \sum_{i \in I} u_i$$

Subject to:

$$u_i \leq x_i + \sum_{j \in N_i} y_i \qquad \text{for all } i \in I \tag{4}$$

$$y_i \leq x_i \qquad \text{for all } i \in I \tag{5}$$

$$\sum_{i \in I} x_i + \sum_{i \in I} y_i = p \text{ (} p \text{ Field Armies)} \tag{6}$$

Of course, a positive coefficient reflecting the relative importance of each region i could multiply each u_i term in the objective.

The maximal covering deployment model doesn't deal with how far away (the number of steps) any of the uncovered nodes are. That is, if a node is not covered in one step, there is no requirement that it can be covered in two steps, or three steps, etc. Furthermore, the model as stated does not deal with the issue of a second war.

Several methods are available to "solve" integer programs. Each method's utility varies with the characteristics of the program. We choose to solve the MCDP as a relaxed linear program and then impose the integer requirements if necessary. This means that the three sets of zero-one variables are relaxed in the linear program and are allowed to range continuously between zero and one. The replacement definitions are: $0 \leq x_i \leq 1$; $0 \leq y_i \leq 1$; and $0 \leq u_i \leq 1$. The resulting program can be solved easily by any of the numerous linear programming packages available. Some or all of the variables in the solution obtained via linear programming may lie strictly between zero and one. Such non-zero-one solutions are resolved by the technique of branch and bound, an add-on option generally packaged with linear programming solvers.

Branch and bound begins with the linear programming solution and fixes one non-zero-one variable first to 1 and then to 0. Both of the two resulting problems, the one fixing the variable to 0, the other fixing it to 1, are then solved using linear programming. This process of successively fixing non-zero-one variables to 0 or 1 is called *branching* and produces a bifurcated or tree-like structure of sequential solutions. The branching process continues until an integer feasible solution is found.

The objective value of such an integer feasible solution forms a *bound* for the problem. If other nodes are branched from and if solutions with objective values higher than the bound result, none of these solutions can be optimal. In a minimising problem, all solution nodes (solutions that have not yet been branched on) with more costly objectives whose variables are not all 0, 1 can be cut off. That is, they need not be searched further because they can never develop an integer feasible solution whose objective value is less than the one already found. Other integer-infeasible solutions on the tree may be branched from until either a feasible solution with a lower value of the objective is developed (giving a new bound) or until all developed solutions have objective values that exceed the objective value of the best integer solution. Branch and bound cannot be counted on to resolve all problems to 0, 1 solutions efficiently. It has shown itself to be effective on many 0, 1 location problems [**12**], but for certain types of constraint sets, the number of nodes that must be resolved is so huge that practical limits on computer memory or time preclude the use of this method.

We solved the MCDP problem for the empire as shown in Figure 1, with $p = 4$ FAs to be allocated among the eight regions. Six alternate optima were found, all of which protected the entire empire—all regions of the empire were either secured or securable in a single step. Of course, these alternatives could also have been obtained by complete enumeration (evaluating all possibilities) in this small problem or by solving the SCDP. Only one of the six alternatives placed FAs at Rome.

Dantzig cuts were used to generate six alternate optima. Dantzig discovered and developed the technique of cutting away (making invalid) unwanted or integer infeasible solutions in the context of the famous Knapsack Problem [**5**]. These cuts, which we added in the re-solution of the linear programming problem, simply require that the current solution be excluded in all subsequent analysis. In order to apply the cuts, the zero-one variables that are 1 in the current solution are summed, and their sum, in a new problem, is constrained to be 1 less than the number of such variables.

Of the alternate optimal solutions produced by solving the MCDP, the first all zero-one solution (0 cuts, see Table 1) placed two FAs in Iberia and two FAs in Egypt. We then solved the MCDP a second time, adding a constraint that says the sum of the x_i plus the sum of the y_i in Iberia and Egypt is less than or equal to three. This constraint excludes no possibilities other than the previous solution. The result is the second solution (cut 1) listed in Table 1—namely, two FAs in Iberia and two in Constantinople. To expose further alternate optimal solutions would require not only a constraint that excludes the first solution, but one that excludes this second solution as well. That is, we add a constraint that says that the sum of FAs in Iberia and Constantinople is less than or equal to three. The cuts concluded when the first sub-optimal solution was found. In total, six alternate optimal deployments of four FAs were found, each deployment securing or making securable all of the eight regions of the empire.

The optimal solutions of Constantine's problem are shown in Table 1. The first column, headed "cuts", indicates the number of cut constraints required to get the particular solution shown on that particular line. The first line (0 cuts) shows the solution found with no additional constraints. The last column, headed "B & B", shows the number of branch and bound nodes that had to be evaluated in order to arrive at an optimal integer feasible solution. The remaining columns refer to the

TABLE 1. SOLUTIONS: ROMAN EMPIRE

cuts	BRI	IBE	GAU	NAF	ROM	EGY	CON	AMI	B & B
0		2				2			0
1		2					2		1
2		2						2	3
3	1				2			1	3
4	2					2			4
5			2			2			14

BRI = Britian, IBE = Iberia, GAU = Gaul, NAF = North Africa,
ROM = Rome, EGY = Egypt, CON = Constantinople, AMI = Asia Minor

eight regions of the Empire and each is headed by a three-letter abbreviation that is clarified in an underline to the table. The numbers in the body of the table indicate the number of FAs stationed in the various regions. Each of the following tables has the same form.

Discussion of the Alternate Optima of the MCDP. The single "Roman solution" deployed two FAs in Rome, one FA in Britian, and one in Asia Minor (Table 1, cut 3). Despite the advantage of placing legions in Rome (hence the term "Roman solution"), this deployment suffers from a reduced capability to respond to a second war occurring somewhere else in the empire. Looking at Figure 1 and envisioning this positioning of FAs, one can see that if a war occurred in any of the five unsecured regions, the response to that war would then leave four regions without protection in the event of a second war. That is, no FA could reach the remaining four regions in a single step. In fact, no FA could be launched anywhere in the empire as no two-FA regions would remain.

On the other hand, the equal optimal to the "Roman solution", which consists of two FAs in Iberia and two FAs in Egypt (Table 1, cut 0) performs better than the Roman solution when it comes to protection in the event of a second war. The worst case situation for this deployment would be a first war in Rome itself or in North Africa. It does not matter whether the response to this first war comes from Iberia or from Egypt. In either situation, two regions would be out of reach in a single step. This is a better outcome than occurs in the Roman solution where up to four regions could not be reached in the event of a second war. The as yet unanswered question is how to discover these robust solutions, robust in the sense that they do well even in the event of a second war.

FURTHER COMPUTATIONAL EXPERIENCE

Pax Britannica. Until well into the 19^{th} century Britian possessed sufficient resources to keep many capital ships in all key regions of interest. With (generally) six "Battle Fleets" (BFs), each composed of roughly twenty ships of the line, Britian pursued a forward defense policy for some 150 years. Around the end of the 19^{th} century, with declining British power, the change to steam propulsion, and the rise of Germany as a maritime power, this strategy had to be revised [**9**]. Even though the increasing mobility and power of the modern ship allowed the reduction of the BF to around eight capital ships each (with their support ships and auxiliaries), Britian had only four BFs in 1900 [**2**]. With six key regions this necessitated a shift to a defense in depth. First Sea Lord John Fisher brought three BFs to home (British) waters and stationed one BF in the Mediterranean. His principle reason was the heightening European tension and the increasing naval powers of Germany [**8**].

The rules of movement of BFs are the same as those of the FAs of the Roman Empire. *One BF must be present in a region to launch a second. Now however,* since we are all at sea, a step is counted as passing through one of the key regions. The six key regions are identified in Table 2 and their "one step contiguity" is shown in Figure 2; the identifiers of the key regions are across the top and down the left side. An "X" indicates that one region is directly reachable from the other, a one-step move for the BFs.

The strategy of Fisher results in only three of the six regions (Britian, The Mediterranean, and The West Indies) being secured or securable in one step. The worst off region (The Far East) is five steps away [**2**].

TABLE 2. SOLUTIONS: BRITISH EMPIRE

cuts	WIN	BRI	MED	CGH	SAS	FEA	B & B
0			2	2			0
1		1	1	2			0
2		2		2			1
3		2			2		3
4			2	1		1	2
5			2			2	0
6	2					2	8
7		2				2	7
8	2					2	4
9	1	1			2		5
10	2				1	1	6
11			2		2		6
12	2				2		14

WIN = West Indies, BRI-Britian, MED = Mediterranean,
CGH = Cape of Good Hope, SAS = South Asia, FEA = Far East

	WIN	BRI	MED	CGH	SAS	FEA
WIN		X	X	X		
BRI	X		X			
MED	X	X			X	
CGH	X				X	X
SAS			X	X		X
FEA				X	X	

Figure 2. Contiguity Matrix, British Empire. For clarification of abbreviations see Table 2.

Thirteen alternate optimal solutions to the MCDP are displayed in Table 2. All key regions of The British Empire are either secure with the presence of a BF or they can be reached by a BF, launched from a two-BF region, in one step. The greater flexibility (larger number of optimal configurations) is a function of the reduction in the number of regions that have to be covered from eight (Roman Empire) to six (British Empire).

Pax Americana. An application of the MCDP to the strategic network of regions of import to the USA demonstrates that the method is not merely of historical interest. While the USA is not being forced into a defense in depth strategy by declining economic and political power, in this post-colonial, post-right-by-conquest period, certain foreign stationings are problematic or impossible.

The definition of the Unit of Force (UF, a unit analogous to the FA or BF) for the Post-Cold War United States with its plethora of military arms is more problematic. Arquilla and Fredricksen [**2**] base their calculations on a study by Aspin [**1**] (a former secretary of defense). They arrive at four UFs, each consisting of roughly three divisions of infantry, three carrier battle groups, and five air wings with their required auxiliaries and support units.

The 15 equal optimal solutions are given in Table 3 and the one-step contiguity pattern on which the solutions are based is shown in Figure 3. This figure indicates the one-step moves possible for UFs. In all of the solutions in Table 3 the "empire" is either secured or securable with one step from a region hosting two UFs. The large number of equal optimals is hardly surprising when we are dealing with four UFs and a total of only five regions. The regions are named in the underline to the associated Table (Table 3).

TABLE 3. SOLUTIONS: AMERICAN EMPIRE

cuts	USA	EUR	NEA	SAS	EAS	B & B
0	2				2	0
1	2			1	1	1
2	2	1		1		1
3	1	1		2		2
4	2			2		2
5			2	1	1	3
6			2	2		5
7	2		1	1		6
8	1		2		1	1
9	2		2			5
10			2		2	9
11		1	2		1	12
12		1	1		2	12
13		2		1	1	13
14		2			2	12
15		2		2		15

USA = United States, EUR = Europe, NEA = Near East,
SAS = South Asia, EAS = East Asia

	USA	EUR	NEA	SAS	EAS
USA		X	X		X
EUR	X		X		
NEA	X	X		X	
SAS			X		X
EAS	X			X	

Figure 3. Contiguity Matrix, American Empire. For clarification of abbreviations see Table 2.

CONCLUSIONS. We have structured a classical problem on the deployment of military forces as a pair of 0, 1 programming problems. As far as we are aware, neither the Set Covering Deployment Problem nor the Maximum Covering Deployment Problem has previously been defined. Once a new problem is formally stated, other applications frequently occur. For example, several applications have appeared for the Maximum Covering Location Problem, a model developed for locating emergency service in order to maximize the population that can be served within a distance standard; see **[4]** and **[14]**.

We also demonstrate the utility of Dantzig cuts to reveal alternate optimal solutions by making infeasible each solution as it is found. We are not aware of similar uses of Dantzig cuts to expose a sequence of alternate optimal solutions (and stopping when the first sub-optimal is found). This technique could, however, also be useful to explore the sub-optimal region in the immediate neighborhood of the optimal solution. In such a case, as each cut is added, the next-best sub-optimal solution would be produced. The technique can be used to back off from the optimal solution or explore alternate optima in other 0, 1 problems as well.

In the computational experience we have shown that the relaxed linear programming version of the problem solves in 0, 1 variables with either no branch and bound or with only modest amounts of it. In general, we have observed that the greater the number of cuts appended, the greater the amount of branch and bound required to resolve the integer infeasibilities (see Table 1–3).

Perhaps most interesting from the standpoint of combinatorial optimization, however, is the fact that a mathematical program has effectively solved a puzzle of some fame. The reduction of the puzzle to a mathematical programming form suggests the possibility that other puzzles—some of which may have significant applications—have related types of formulations.

REFERENCES

1. L. Aspin, *Report on the BOTTOM-UP REVIEW*, Government Printing Office, Washington, D.C., 1993.
2. J. Arquilla and H. Fredricksen, 'Graphing' an Optimal Grand Strategy, *Mil. Oper. Res. J.* 1 (1995) 3–17.
3. R. Church and C. S. ReVelle, The Maximal Covering Location Problem, *Papers Reg. Sci. Assoc.* 32 (1973) 101–118.
4. C. Chung, Recent Application of the Maximal Covering Location Planning Model, *J. Oper. Res. Soc.* 37 (1986) 735–746.
5. G. Dantzig, *Linear Programming and Extensions*, Princeton University Press, Princeton, NJ, 1963.
6. M. Daskin, *Network and Discrete Location: Models, Algorithms, Applications*, Wiley-Interscience, New York, 1995.
7. Z. Drezner, ed., *Facility Location: A Survey of Applications and Methods*, Springer, New York, NY, 1995.
8. J. A. Fisher, *Memories*, Hodder and Stoughton, London, 1919.
9. P. Kennedy, *The Rise and Fall of British Naval Mastery*, Ashfield Press, London, 1976.
10. R. Love, J. Morris, and G. Wesolowsky, *Facilities Location: Models and Methods*, North-Holland, Amsterdam, 1988.
11. E. N. Luttwak, *The Grand Strategy of the Roman Empire*, Johns Hopkins University Press, Baltimore, MD, 1976.
12. C. S. ReVelle, Facility Siting and Integer-Friendly Programming, *Eur. J. Oper. Res.* 65 (1993) 329–342.
13. C. S. ReVelle, C. Toregas, and L. Falkson, Applications of the Location Set Covering Problem, *Geog. Analy.* 8 (1976) 65–77.
14. D. Schilling, J. Vaidyanathan, and R. Barkhi, A Review of Covering Problems in Facility Location, *Loc. Sci.* 1 (1993) 25–56.
15. C. Toregas, and C. S. ReVelle, Binary Logic Solutions to a Class of Location Problems, *Geog. Analy.* 5 (1973) 145–156.
16. C. Toregas, R. Swain, C. S. ReVelle, and L. Bergmann, The Location of Emergency Services, *Oper. Res.* 19 (1971) 1363–1373.

CHARLES S. REVELLE is a professor at The Johns Hopkins University. He took his bachelor's and doctoral degrees at Cornell University. He finds endless entertainment in formulating 0-1 programming problems that describe real situations. He has created 0-1 problems in the field of water management, location science, biological conservation strategies, and forestry management science. He has also written with his wife, Penelope, five undergraduate textbooks on environmental science as well as a book on water resources management and a textbook on mathematical programming. "A window and a light will do for thinking."
Department of Geography and Environmental Engineering
313 Ames Hall, The Johns Hopkins University, Baltimore, MD 21218
revelle@jhu.edu

KENNETH E. ROSING is a recently retired former Reader in Economic Geography, Erasmus University Rotterdam, who received his Ph.D. (jointly) from King's College and University College. He has worked with computers, particularly for the visual representation of data (maps) since the mid 60's, when he began writing algorithms for the first computer generated atlas (published 1970). He believes that research exists for itself and that the serendipital spin-offs of research more than justify the investment. His avocation is developing inventive new tricks to force computers to solve the integer models developed by Professor ReVelle. "A window and a light will do for thinking" (but a Cray in the back room sure helps to check the answers).
Van Polanenpark 23, 4791 BP Klundert, The Netherlands
rosing@few.eur.nl