SOME COMBINATORIAL LEMMAS

BY

DONALD E. KNUTH

STAN-CS-72-291
JUNE 1972

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY

SOME COMBINATORIAL LEMMAS

by Donald E. Knuth

This report consists of several short papers which are completely independent of each other:

1. "Wheels Within Wheels." Every finite strongly connected digraph is either a single point or a set of n smaller strongly connected digraphs joined by an oriented cycle of length n. This result is proved in somewhat stronger form, and two applications are given.

2. "An Experiment in Optimal Sorting." An unsuccessful attempt, to sort 13 or 14 elements in less comparisons than the Ford-Johnson algorithm, is described. (Coauthor: E. B. Kaehler.)

3. "Permutations With Nonnegative Partial Sums." A sequence of s positive and t negative real numbers, whose sum is zero, can be arranged in at least $(s+t-1)!$ and at most $(s+t)!/(\max(s,t)+1) < 2(s+t-1)!$ ways such that the partial sums $x_1 + \ldots + x_j$ are non-negative for $1 \leq j \leq s+t$.

This research was supported in part by the National Science Foundation under grant number GJ-992, and the Office of Naval Research under contract number N-00014-67-A-0112-0057 NR 044-402. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Title: Wheels Within Wheels

Author: Donald E. Knuth
Computer Science Department
Stanford, California 94305

Abstract: A simple means of representing the structure of all strongly-connected directed graphs is developed and applied to a certain problem involving the construction of "toll booths".

Keywords and Phrases: directed graph, strongly connected, toll booth problem, topological sorting, Petri nets.

AMS category: 05C20

WHEELS WITHIN WHEELS

by Donald E. Knuth

"Their appearance and their work was, as it were, a wheel within a wheel."

— Ezekiel 1:16

A <u>strongly-connected digraph</u> is a nonempty directed graph in which an oriented path exists from any vertex to any other. The following lemma shows that all finite strongly-connected digraphs can be constructed in a fairly simple way.

<u>Lemma 1</u>. Every strongly-connected digraph $\mathcal{D}$ is either a single vertex with no arcs, or it can be represented as follows for some $n \geq 1$:

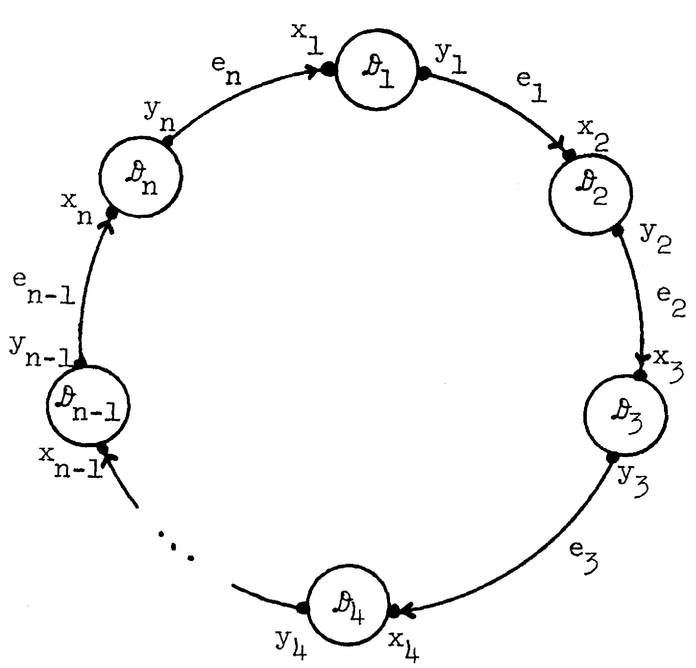

Here $\mathcal{D}_1,\ldots,\mathcal{D}_n$ are strongly-connected digraphs; x_i and y_i are (possibly equal) vertices of $\mathcal{D}_i$; and e_i is an arc from y_i to x_{i+1} . The original digraph $\mathcal{D}$ consists of the vertices and arcs of $\mathcal{D}_1,\ldots,\mathcal{D}_n$ plus the arcs $e_1,\ldots,e_n$.

In fact, if σ is any given oriented cycle of $\mathcal{D}$, there exists such a representation in which each of the e_i is contained in σ .

(Since each of the $\mathcal{D}_i$ can be further decomposed in the same way, every strongly-connected digraph essentially consists of "wheels within wheels".)

Proof. Let e be an arc of $\mathcal{D}$, and let the relation

$$x \leftrightarrow y \quad [\text{without } e]$$

mean that oriented paths excluding e exist from x to y and from y to x . This is an equivalence relation which partitions the vertices of $\mathcal{D}$ into components, namely the so-called strong components of $\mathcal{D} - e$.

Suppose that $e' = (x',y')$ and $e'' = (x'',y'')$ are distinct arcs such that $x' \leftrightarrow x''$, $x' \not\leftrightarrow y'$, $x'' \not\leftrightarrow y''$ [without e] . Let $\mathcal{D}^0$, $\mathcal{D}'$, $\mathcal{D}''$ denote the respective components of x' , y' , and y'' ; possibly $\mathcal{D}' = \mathcal{D}''$. The shortest oriented path leading from a vertex of $\mathcal{D}''$ to a vertex of $\mathcal{D}^0$ involves no arcs leading from vertices in $\mathcal{D}^0$, hence $\mathcal{D} - e'$ contains an oriented path from $\mathcal{D}''$ to $\mathcal{D}^0$. Since e'' goes back from $\mathcal{D}^0$ to $\mathcal{D}''$, we have

$$x'' \leftrightarrow y'' \quad [\text{without } e'] \ .$$

Furthermore

$$x \leftrightarrow y \ [\text{without } e] \ \text{ implies } \ x \leftrightarrow y \ [\text{without } e'] \ ,$$

for all x and y, since $x \leftrightarrow y$ [without e] means that x and y belong to the same component, and e' does not lie within any component. Thus $\mathcal{D} - e'$ has fewer strong components than $\mathcal{D} - e$.

Two distinct arcs leading out of the same component, such as e' and e'' in the above discussion, may be called "mates". We shall now prove that there is an arc e in σ such that $\mathcal{D} - e$ contains no mates. This will prove the lemma, since the components $\mathcal{D}_1, \mathcal{D}_2, \ldots$ must then have the cyclic form shown, with $e = e_1$, say, and with all the other e_i belonging to σ.

If $\mathcal{D}$ is finite we simply let e be the arc in σ such that $\mathcal{D} - e$ has the fewest components. This will imply the nonexistence of mates, since another arc e' of σ which is not included in some component cannot have a mate, lest $\mathcal{D} - e'$ have fewer components.

If $\mathcal{D}$ is infinite the argument is slightly more tricky. We choose e so that the minimum possible number of arcs of σ fail to lie within components of $\mathcal{D} - e$; if two arcs e are equally good by this criterion, we choose one that minimizes the number of arcs of σ that have mates. Now $\mathcal{D} - e$ contains no mates; for if it did, there would be an arc e' of σ which has a mate e''. But this contradicts the choice of e, since e must not lie in a component of $\mathcal{D} - e'$, and $\mathcal{D} - e'$ contains at least one less mated arc of σ. $\square$

If σ contains an arc e such that $\mathcal{D} - e$ is strongly connected, then $n = 1$ and the lemma holds rather trivially. But if $\mathcal{D}$ contains no such "redundant" arcs, then $n \geq 2$ and the same will be true for the $\mathcal{D}_i$.

The representation is not unique, even if σ is specified to be a simple cycle and $\mathcal{D}$ contains no redundant arcs; for example,

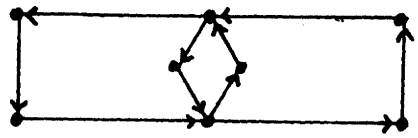

has two representations with $n = 3$ when σ is the outermost cycle.

Lemma 1 is sometimes useful when proving properties of strongly-connected digraphs by induction, or when finding counterexamples to conjectures. We shall consider one application here, namely the <u>free route to Las Vegas</u> problem:

<u>Theorem</u>. Given a finite, strongly-connected network $\mathcal{D}$ of one-way roads between cities, and a designated city called Las Vegas, it is possible to erect toll booths on these roads in such a way that the following three conditions are satisfied:

(i) There is no toll-free cycle. (It is impossible to drive indefinitely without paying a toll.)

(ii) Every road is part of a one-toll cycle. (It is possible to start on any road and return to your starting point, paying only one toll.)

(iii) There is a toll-free route from every city to Las Vegas.

(Condition (i) calls for comparatively many toll booths, while conditions (ii) and (iii) call for comparatively few. By (iii) and (i), every road leaving Las Vegas must contain a toll booth.)

<u>Proof</u>. We argue by induction on the number of roads (i.e., arcs) in $\mathcal{D}$, since the theorem is vacuously true when there are no roads. Using the

representation of $\mathcal{D}$ in Lemma 1, we can erect tollbooths in each $\mathcal{D}_i$ such that (i) and (ii) hold and such that there are toll-free routes from x_i to y_i, for $1 \leq i \leq n$. Establishing one further toll-booth on e_1 makes conditions (i) and (ii) hold in the entire digraph.

Now the proof is completed by applying another lemma.

Lemma 2. Using the terminology of the Theorem, if (i) and (ii) can be achieved in $\mathcal{D}$, it is possible to modify the placement of toll booths so that all three conditions are achieved.

Proof. The following move operation preserves both (i) and (ii), because it neither increases the number of toll-booths on any cycle nor decreases that number to zero:

"Let x be a city with toll booths on all roads leading into it. Destroy all these toll booths, and erect new ones on all roads leaving x, except on those roads which already have toll booths."

The proof will be complete if we can show that a sequence of move operations will produce condition (iii).

Given an arrangement of toll booths in $\mathcal{D}$, let $\mathcal{D}'$ be the subnetwork consisting of all cities x from which there exists a toll-free route to Las Vegas (including Las Vegas itself), and all toll-free roads between such cities. By condition (i), $\mathcal{D}'$ contains no oriented cycles, hence we can "topologically sort" the cities of $\mathcal{D}'$ into the order $x_1,\ldots,x_n$ such that no road of $\mathcal{D}'$ goes from x_i to x_j for $i \geq j$. (It follows that x_n is Las Vegas.)

Now consider the grand move operation, which consists of successively doing a move operation on vertices $x_1,x_2,\ldots,x_n$, in this order. We must show that such a grand move is well-defined, in the sense that all moves are legal. After having moved toll booths past $x_1, \ldots, x_{j-1}$, the roads leading into x_j are of two kinds:

(a) Those which had toll-booths before the grand move began. (These toll booths are still present.)

(b) Those which had no toll-booths before the grand move began. (These belong to $\mathcal{B}'$, so they lead from x_i for some $i < j$; they received toll booths when we moved past x_i.)

Therefore we can move past x_j.

After the grand move operation, all the roads in $\mathcal{B}'$ are still toll-free, since the road from x_i to x_j received a toll booth on move i and it was eliminated on move j. Furthermore if $\mathcal{B}'$ did not originally include all cities, there will be a road from some city x not in $\mathcal{B}'$ to some city x_i in $\mathcal{B}'$, since $\mathcal{B}$ is strongly connected. The grand move operation removes the toll from that road. Hence $\mathcal{B}'$ will grow until eventually every city belongs to it, i.e., until condition (iii) holds. □ □

It is plausible to guess that the theorem could be extended, replacing condition (ii) by a similar one:

(ii') There is a one-toll route from Las Vegas to every other city.

But the following counterexample (obtained by considering Lemma 1) shows that this cannot be done in general, since there is a unique way to install the toll booths meeting conditions (i) and (iii):

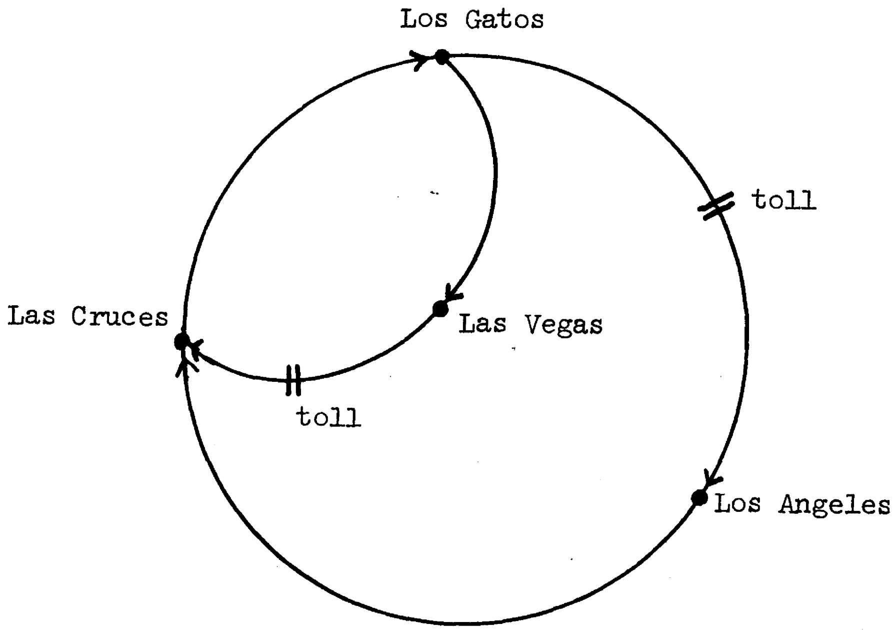

Acknowledgments: I wish to thank Anatol W. Holt for originally suggesting the toll-booth problem to me (without the Las Vegas constraint). This problem originated in his work on Petri nets.

AN EXPERIMENT IN OPTIMAL SORTING

by Donald E. Knuth and E. B. Kaehler

Since Ford and Johnson published their "merge-insertion" method of sorting in 1959 [1], nobody has been able to discover a sorting algorithm which uses fewer comparisons in its worst case. Their method has been proved optimal when 12 or less elements are being sorted, and it appears reasonable to conjecture that a better algorithm exists for 13 or 14 elements; but the possibilities are so enormous that an exhaustive computer search appears to be out of the question. The purpose of this note is to report the results of an unsuccessful attempt to improve on merge insertion when $n = 13$ or 14, in the hope that our experiments might suggest a new approach to the problem.

The maximum number of comparisons needed by merge insertion is known [2] to be

$$F(n) = \sum_{1 \le k \le n} \lceil \log_2(\tfrac{3}{4}n) \rceil \quad .$$

Since $F(n) - F(n-1) = \lceil \log_2 n \rceil$ when $12 \le n \le 16$, or when $24 \le n \le 32$, or $48 \le n \le 64$, etc., the Ford-Johnson procedure is no better in its worst case than simply sorting $n-1$ elements in $F(n-1)$ comparisons, then inserting the n-th element by binary insertion, for all such n. Surely there must be a better way than this!

An independent confirmation of Wells's proof [4] that there is no better way when $n = 12$, revealed a particularly efficient line of attack with respect to 9 elements. By combining this with efficient constructions for 4 or 5 elements, we hoped to come up with sorting procedures for 13 or 14 elements.

Suppose that the elements to be sorted are $K_1, \ldots, K_n$, and that they are distinct. Let the relation " $i < j$ " mean that K_i has been compared to K_j and that $K_i < K_j$. Let S be a set of such relations, and let $P(S)$ be the number of permutations $K_1 K_2 \ldots K_n$ of the set $\{1,2,\ldots,n\}$ such that the relations of S are valid. For example, if $n = 4$ and $S = \{1 < 2, 3 < 4, 2 < 4\}$ then $P(S) = 3$ since there are only three permutations $K_1K_2K_3K_4$ of $\{1,2,3,4\}$ such that $K_1 < K_2 < K_4$ and $K_3 < K_4$, namely 1234, 1324, 2314.

When $P(S) = 1$, the sorting has been completed. When $P(S) > 1$, we need to make another comparison, say between K_i and K_j, and then the two cases $S_1 = S \cup \{i < j\}$ and $S_2 = S \cup \{j < i\}$ must both be dealt with in the same way. When $P(S) \geq 2^k$, we must have either $P(S_1) \geq 2^{k-1}$ or $P(S_2) \geq 2^{k-1}$ (or both), hence at least k more comparisons must be made in some branch of the algorithm before the sorting is complete. Intuitively it seems best to choose i and j so that $P(S_1)$ and $P(S_2)$ are each approximately half of $P(S)$.

Suppose 8 elements are to be sorted, and that the first four comparisons are $K_1: K_2$, $K_3: K_4$, $K_5: K_6$, $K_7: K_8$. By renaming the elements if necessary, we may assume without loss of generality that the results are $1 < 2$, $3 < 4$, $5 < 6$, $7 < 8$. Then we may compare $K_1: K_3$ and $K_5: K_7$, and assume by symmetry that $1 < 3$ and $5 < 7$. And then we may compare $K_4: K_8$ and assume that $4 < 8$. All of these comparisons split the number of possibilities perfectly into two equal parts, hence if $S^{(7)}$ is the set of seven relations known so far we have $P(S^{(7)}) = 8!/2^7 = 315$.

If we now compare $K_4 : K_6$ we find 157 cases with $4 < 6$, 158 cases with $6 < 4$. Let's work on the latter possibility, since it is probably a little harder; if we compare $K_1 : K_5$ it turns out that $1 < 5$ occurs 77 times but $5 < 1$ occurs 81 times, and again we focus attention on the latter case. Introducing a new element K_9, it may be in any of 9 relative positions with respect to the original 8 elements, hence there are $9 \times 81 = 729$ cases to consider. In 372 of these, $7 < 9$, while $9 < 7$ in the remaining 357. The 372 cases can be broken into 192 with $4 < 9$ and 180 with $9 < 4$.

Let's look at these 192 cases in detail. The eleven relations $S^{(11)} = \{1<2, 3<4, 5<6, 7<8, 1<3, 5<7, 4<8, 6<4, 5<1, 7<9, 4<9\}$ can be diagrammed as follows:

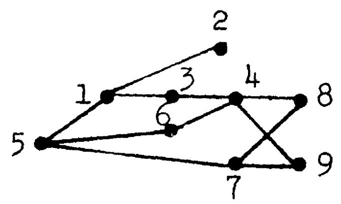

The symmetry between 8 and 9 implies that we can assume $8 < 9$ without loss of generality; this leads to $S^{(12)} = S^{(11)} \cup \{8 < 9\}$, with 96 possibilities. Curiously a perfect split occurs if we now compare K_2 with K_4. There are 48 possibilities with $2 < 4$, i.e., with

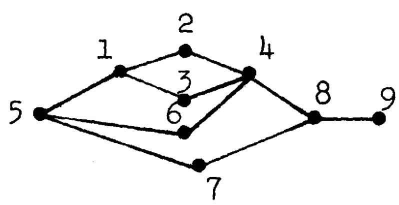

By symmetry we may assume $2 < 3$, and then 6 can be inserted into its proper place relative to $\{1,2,3\}$ in two more comparisons; and now we may compare K_7 with the middle element of $\{K_1, K_2, K_3, K_4, K_6\}$. The result is to produce one of the two configurations

or

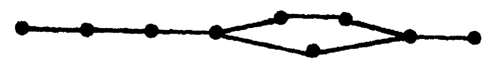

each of which has 3 remaining possibilities.

In the other branch, the 48 possibilities in $S^{(12)} \cup \{4 < 2\}$ can also be reduced to $48/16 = 3$ in four more comparisons. Starting with

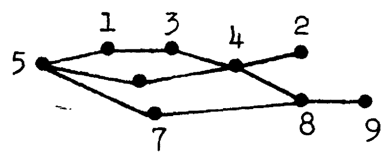

,

we compare K_3 with K_7. If $3 < 7$ we obtain the diagram

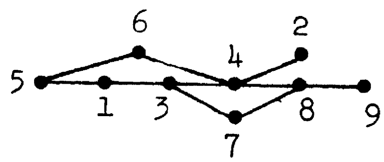

which is symmetrical with resepect to 4 and 7 (if we remove the least element, 5); after three more comparisons we obtain either

or

.

On the other hand of $7 < 3$, there is symmetry between 7 and 1, and again the reduction is straightforward.

We have now worked out the "hardest" lines of a partial sorting procedure; in each case considered, we discovered a way to reduce the 9! initial possibilities to only 3, after a grand total of 17 comparisons. This is remarkably efficient, since

$$\frac{9!}{2^{17}} = \frac{945}{1024} \cdot 3$$

is only slightly less than 3; in other words, each of the comparisons has split the possibilities very nearly in half.

Let us therefore assume that a 17-step procedure exists such that each of the sets $S^{(17)}$ that occurs has at most 3 corresponding permutations. (We haven't proved this, but we have grounds to suspect it is true since the lines not yet considered have fewer possibilities and comparatively more freedom.) It is plausible to suspect that this fact could be used to discover a sorting procedure for 14 items. If we add five more elements $K_{10},\ldots,K_{14}$ and sort them using seven comparisons, we obtain configurations $S^{(24)}$ which include

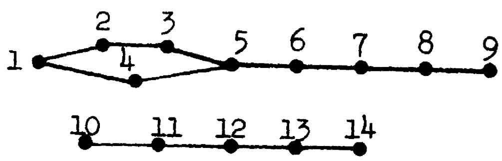

as one of the possibilities after 24 comparisons have been made. This configuration represents $3\binom{14}{5} = 6006$ possibilities; since 6006 is comfortably less than $2^{13} = 8192$, it appears likely that the above configuration can be sorted in 13 more comparisons, for a total of $24 + 13 = 37 < F(14)$.

Therefore an exhaustive search was programmed, based on the above configuration. The following matrix shows, for each i and j, the value of $P(S^{(24)} \cup \{i < j\})$:

0	6006	6006	6006	6006	6006	6006	6006	6006	3861	5346	5841	5976	6003
0	0	6006	4004	6006	6006	6006	6006	6006	2046	3996	5220	5788	5973
0	0	0	2002	6006	6006	6006	6006	6006	966	2604	4228	5348	5873
0	2002	4004	0	6006	6006	6006	6006	6006	1506	3300	4724	5568	5923
0	0	0	0	0	6006	6006	6006	6006	378	1428	3003	4578	5628
0	0	0	0	0	0	6006	6006	6006	168	798	2058	3738	5250
0	0	0	0	0	0	0	6006	6006	63	378	1218	2730	4620
0	0	0	0	0	0	0	0	6006	18	138	570	1650	3630
0	0	0	0	0	0	0	0	0	3	30	165	660	2145
2145	3960	5040	4500	5628	5838	5943	5988	3861	0	6006	6006	6006	6006
660	2010	3402	2706	4578	5208	5628	5868	5346	0	0	6006	6006	6006
165	786	1778	1282	3003	3948	4788	5436	5841	0	0	0	6006	6006
30	218	658	438	1428	2268	3276	4356	5976	0	0	0	0	6006
3	33	133	83	378	756	1386	2376	6003	0	0	0	0	0

The two median elements, K_5 and K_{12}, can be compared, giving a perfect 3003 to 3003 split. The next best choice is to compare K_7 to K_{13}, giving 3276 to 2730. Each of the 13 possible comparisons such that both numbers are ≤ 4096 was pursued further.

The program was designed so that if, for example, the "large" 3276 branch could be sorted, the "small" 2730 branch was tried too. The two branches were often found to be isomorphic, in which case only one was pursued. However, we discovered to our chagrin that it was impossible to complete any of the "large" branches further than 9 levels; thus, the isomorphism test (designed to simplify the solution we thought would be found) was never actually useful.

The program was written in IBM Assembler language and carefully tested on smaller cases. In the above 6006-case example, 14×14 matrices were generated a total of 1235 times, and the computation took about 23.6 minutes on a 360/67.

A similar experiment was tried on the 13-element configurations

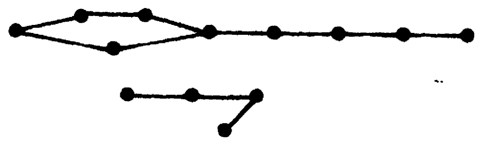

and

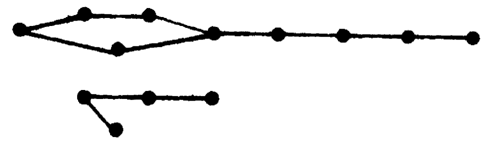

,

which each represent 6435 permutations. But neither of these could be sorted in 13 comparisons. Hence merge insertion still is the champion.

References:

[1] Lester Ford, Jr. and Selmer Johnson, "A tournament problem," American Math. Monthly 66 (1959), 387-389.

[2] A. Hadian, Ph.D. Thesis, University of Minnesota (1969), 38-42.

[3] D. E. Knuth, "Sorting and Searching," The Art of Computer Programming 3 (1972), Section 5.3.

[4] M. B. Wells, "Applications of a language for combinatorial computing," Proc. IFIP Congress 65, vol. 2, 497-498.

Permutations with Nonnegative Partial Sums

by Donald E. Knuth

If $x_1, x_2, \ldots, x_n$ are real numbers whose sum is zero, it is well known that we can find some permutation $p(1)\ p(2) \ldots p(n)$ such that each of the partial sums

$$x_{p(1)} + \cdots + x_{p(j)} \tag{1}$$

is nonnegative, for $1 \leq j \leq n$. In fact, there are at least $(n-1)!$ such permutations; for if $p(1)\ p(2) \ldots p(n)$ is any permutation, it is not hard to see that some cyclic shift $p(k+1) \ldots p(n)\ p(1) \ldots p(k)$ will have this property.

Daniel Kleitman [2] has conjectured that the number of such permutations is always at most $2n!/(n+2)$, if the x's are nonzero. The object of this note is to prove the following sharpened form of his conjecture:

Theorem. Let $x_1, x_2, \ldots, x_n$ be real numbers, where $x_1 + x_2 + \cdots + x_n = 0$; and assume that s elements are greater than zero, t elements are less than zero, and $n-s-t$ elements are equal to zero. Let $P(x_1, x_2, \ldots, x_n)$ denote the number of permutations $p(1)\ p(2) \ldots p(n)$ such that each of the partial sums (1) is nonnegative. Then

$$P(x_1, x_2, \ldots, x_n) \leq \frac{n!}{\max(s,t)+1} . \tag{2}$$

Furthermore this bound is best possible, in the sense that equality is achieved for some x's whenever s, t, n are fixed values with $s + t \leq n$.

Proof. We may obviously assume that s and t are positive. Let ε be the smallest positive value such that the sum over a subset of the x_i is equal to ε. For example, if $n = 4$ and $(x_1,x_2,x_3,x_4) = (\pi, 2-\pi, 1, -3)$, then the 2^4-2 sums over proper subsets of the x's are

$$\pi, 2-\pi, 2, 1, 1+\pi, 3-\pi, 3, -3, -3+\pi, -1-\pi, -1, -2, -2+\pi, -\pi .$$

The smallest positive value among these is $\varepsilon = -3+\pi$. By symmetry, the largest negative value will always be $-\varepsilon$, since $x_1 + x_2 + \cdots + x_n = 0$.

Now let $x_0 = -\varepsilon$, and consider the $n+1$ values $-\varepsilon, x_1, x_2, \ldots, x_n$ whose sum is $-\varepsilon$. A permutation $q(0)\,q(1)\ldots q(n)$ of the indices $\{0,1,\ldots,n\}$ will be called <u>special</u> if each of the partial sums

$$x_{q(0)} + x_{q(1)} + \cdots + x_{q(j)}$$

is nonnegative, for $0 \le j < n$. (When $j = n$, of course, this sum will be $-\varepsilon$.)

The plan of the proof is to show first that there are exactly $n!$ special permutations. Then we shall map each of the $P(x_1,x_2,\ldots,x_n)$ permutations into $t+1$ distinct special permutations. This will prove that

$$(t+1)P(x_1,x_2,\ldots,x_n) \le n! \quad , \tag{3}$$

and by symmetry the same will be true with s replacing t, hence (2) will follow.

In order to count the special permutations, we shall use a cyclic-equivalence argument to show that exactly $1/(n+1)$ of the permutations are special. Let $q(0)\,q(1)\ldots q(n)$ be a permutation of the indices $\{0,1,\ldots,n\}$, and consider the quantity

$$f(j) = x_{q(0)} + x_{q(1)} + \cdots + x_{q(j)} + \frac{j+1}{n+1}\varepsilon \quad . \qquad (4)$$

This function takes on $n+1$ distinct values for $0 \le j \le n$, since $f(j) = f(j')$ and $j < j'$ imply that $x_{q(j+1)} + \cdots + x_{q(j')} = (j-j')\varepsilon/(n+1)$, contrary to our choice of ε. The permutation $q(k+1) \ldots q(n)\ q(0) \ldots q(k)$ is special if and only if

$$x_{q(k+1)} + \cdots + x_{q(j)} \ge 0 \qquad \text{for } k < j \le n \, ; \qquad (5)$$

$$x_{q(k+1)} + \cdots + x_{q(n)} + x_{q(0)} + \cdots + x_{q(j)} \ge 0 \quad \text{for } 0 \le j < k \, . \ (6)$$

If $j > k$, we have $f(j) > f(k)$ if and only if $x_{q(k+1)} + \cdots + x_{q(j)} > (k-j)\varepsilon/(n+1)$ if and only if $x_{q(k+1)} + \cdots + x_{q(j)} \ge 0$. And if $j < k$, we have $f(j) > f(k)$ if and only if $x_{q(j+1)} + \cdots + x_{q(k)} < (j-k)\varepsilon/(n+1)$ if and only if $x_{q(k+1)} + \cdots + x_{q(n)} + x_{q(0)} + \cdots + x_{q(j)} > (k-j-n-1)\varepsilon/(n+1)$ if and only if $x_{q(k+1)} + \cdots + x_{q(n)} + x_{q(0)} + \cdots + x_{q(j)} \ge 0$. In other words, $q(k+1) \ldots q(n)\ q(0) \ldots q(k)$ is special if and only if

$$f(k) = \min_{0 \le j \le k} f(j) \ ,$$

and this uniquely characterizes the value of k. It follows that exactly $n!$ of the $(n+1)!$ possible permutations $q(0)\ q(1) \ldots q(n)$ are special.

Now let $p(1)\ p(2) \ldots p(n)$ be a permutation of $\{1,2,\ldots,n\}$ such that all of the partial sums (1) are nonnegative, and let i be one of the t indices such that $x_{p(i)} < 0$. Then

$$p(1) \ldots p(i-1)\, 0\, p(i+1) \ldots p(n)\ p(i) \tag{7}$$

is a special permutation of $\{0,1,\ldots,n\}$, since $x_{p(i)} \leq x_0 = -\varepsilon$. Furthermore

$$p(1)\ p(2) \ldots p(n)\ 0 \tag{8}$$

is obviously a special permutation. In this way we can construct $(t+1)P(x_1,x_2,\ldots,x_n)$ special permutations, which clearly are all distinct. Therefore (3) holds, and (2) must be true.

To complete the proof of the theorem, we must show that (2) is best possible. This is equivalent to finding examples in which the permutations constructed in (7) and (8) exhaust all the special permutations. Such examples obviously arise whenever $x_i = -\varepsilon$ for some $i \geq 1$. Therefore equality holds in (2) when $1 \leq s \leq t$ and

$$x_1 = t-s+1$$
$$x_2 = \ldots = x_s = 1$$
$$x_{s+1} = \ldots = x_{s+t} = -1$$
$$x_{s+t+1} = \ldots = x_n = 0 \quad .$$

(For these x's , $\varepsilon = 1$. The fact that $P(x_1,\ldots,x_n) = n!/(t+1)$ in this case is well-known, since it is equivalent to other combinatorial problems; see, for example, Erdélyi and Etherington [1].) ∎

If none of the 2^n-2 sums over proper subsets of the x's is zero, it is easy to see by considering cyclic permutations that $P(x_1,x_2,\ldots,x_n) = (n-1)!$. Conversely, if $P(x_1,x_2,\ldots,x_n) = (n-1)!$, those partial sums must all be nonzero.

In the general case the possible values of $P(x_1,x_2,\ldots,x_n)$ seem to be spread out rather evenly between $(n-1)!$ and $n!/(\max(s,t)+1)$.

For example, let a_1, a_2, b_1, b_2, b_3 be positive numbers with $a_1 + a_2 = b_1 + b_2 + b_3$; the theorem tells us that $24 \le P(a_1, a_2, -b_1, -b_2, -b_3) \le 30$. In fact it is not difficult to verify in this case that $P(a_1, a_2, -b_1, -b_2, -b_3)$ equals 24 plus twice the number of pairs (i,j) such that $a_i = b_j$. Thus,

$$P(5,1,-2,-2,-2) = 24$$
$$P(4,1,-1,-2,-2) = 26$$
$$P(3,1,-1,-1,-2) = 28$$
$$P(2,1,-1,-1,-1) = 30 \quad .$$

Kleitman has used the above theorem to determine the asymptotic number of different score sequences possible in an n-person round-robin tournament, to within a factor of 2 .

References

[1] A. Erdélyi and I. M. H. Etherington, "Some problems of non-associative combinations. II.", Edinburgh Math. Notes 32 (1940), 7-12.

[2] D. Kleitman, "The number of tournament score sequences," Proc. Calgary International Conf. on Combinatorial Structures and their Applications, June, 1969 (New York: Gordon and Breach, 1970), 209-213.