Erin Triples With Bobs

Don Knuth, Stanford Computer Science Department
(18 April 2026, revised 22 April 2026)

I learned a week ago about a fascinating unsolved problem about permutations, and I think the problem deserves to be more widely known. So I'm writing these notes in hopes that somebody will be inspired to solve it.

A wonderful tradition called "change ringing" was developed at the large churches of 17th century England, where independently ringable bells of $n+1$ different pitches could be found in the bell towers. The idea was to ring all of the bells in sequence — always ending with the lowest note, but with $n!$ possible permutations in which the other notes could appear. With patience and persistence, all of those permutations could actually be played, once each, until returning to the starting order and thereby achieving a "full peal." (I gave a brief introduction to this history on pages 322–324 of [6], because it has a close connection to computer programming. One of the early ways to obtain a full peal — published in 1668, and called "plain changes" — is still often the method of choice on modern hardware, because each permutation differs from its predecessor by simply interchanging the order of two adjacent bells. For example, a full peal of plain changes for $n = 3$ goes '123 — 132 — 312 — 321 — 231 — 213 — 123'. When $n = 4$ it's '1234 — 1243 — 1423 — 4123 — 4132 — 1432 — 1342 — 1324 — 3124 — 3142 — 3412 — 4312 — 4321 — 3421 — 3241 — 3214 — 2314 — 2341 — 2431 — 4231 — 4213 — 2413 — 2143 — 2134 — 1234'.)

Change ringing is also a great social pastime, as the $n+1$ people who control the bells do their best to play the correct melody without getting confused. When $n = 5$, the most popular peals try to interchange *two* adjacent pairs of bells, going for example from 12345 to either 13254 or 21354 or 21435. This is called Doubles. And when $n = 7$ we get Triples: The successor of 1234567 is either 1325476 or 2135476 or 2143576 or 2143657, always swapping *three* pairs of neighboring bells.

Of course $7! = 5040$ is a rather large number of permutations. Even so, a full peal can be performed in three or four hours. Such peals are especially nice because $7+1 = 8$ gives us a full octave, a complete "scale."

Let $p_1 = (23)(45)(67)$, $p_3 = (12)(45)(67)$, $p_5 = (12)(34)(67)$, and $p_7 = (12)(34)(56)$ be the four triple swaps, namely the four permutations that interchange three adjacent pairs, written in cycle notation. It turns out that the product of adjacent triple swaps is a permutation of order 3; for example, $p_1p_3 = (123)$ and $p_3p_5 = (345)$.

There is undoubtedly a vast number of total peals on $7+1$ bells for which every permutation is obtained from its predecessor by performing one of the triple swaps. Such peals correspond to Hamiltonian cycles on the graph of $7! = 5040$ permutations, where each vertex α has four neighbors: $p_1\alpha$, $p_3\alpha$, $p_5\alpha$, $p_7\alpha$. (Think of a permutation $\alpha = abcdefg$ as the mapping $1 \mapsto a$, $2 \mapsto b$, $\ldots$, $7 \mapsto g$. Then $p_3\alpha$ is the mapping $1 \mapsto b$, $2 \mapsto a$, $3 \mapsto c$, $4 \mapsto e$, $5 \mapsto d$, $6 \mapsto g$, $7 \mapsto f$, namely $bacedgf$.) This 4-regular graph is the so-called "Cayley graph" on the generators $\{p_1, p_3, p_5, p_7\}$. But most of those Hamiltonian cycles are unsuitable for an actual performance, because their haphazard rules of succession are too complicated for humans to deal with. Bellringers have therefore tried to construct peals for which the rules are reasonably straightforward.

Fabian Stedman (1640–1713), one of change ringing's great pioneers, devised a famous composition called "Stedman Doubles," which achieves half of the permutations for $n = 5$ by performing the sequence

$$(p_3, p_1, p_3, p_1, p_3, p_5, p_1, p_3, p_1, p_3, p_1, p_5)^5 \tag{1}$$

of 60 changes. (Here I'm converting Triples to Doubles, by implicitly removing '(67)' from p_1, p_3, and p_5.) Since every double swap is an even permutation, that's the best we can do with Doubles. To get a full peal of length $5! = 120$, we can include also the odd permutations, by replacing any one of the 60 changes in (1) with a single swap such as '(12)'; all permutations will then be heard when the resulting sequence is carried out twice. (Indeed, that's one of the nicest possible rules for a full peal when $n = 5$.)

Therefore campanologists hoped to discover a similarly memorable rule for peals of triples when $n = 7$. The "Holy Grail" for Stedman Triples, sought for a long time, was to be a rule of the form

$$(p_3, p_1, p_3, p_1, p_3, p_{5/7}, p_1, p_3, p_1, p_3, p_1, p_{5/7})^{420}, \tag{2}$$

where '$p_{5/7}$' means either p_5 or p_7. If such a rule could be found, five of every six steps would be predetermined, and there would be only two choices for the unspecified steps (thus 2^{840} possibilities, not 4^{5040}). The use of p_7 every time causes a repeat after 84 changes; the use of p_5 every time causes a repeat after 60 changes, as when $n = 5$; but much longer sequences are obtainable if p_5 and p_7 are judiciously mixed.

A convenient code was invented, for compositions whose structure conforms to (2): One of the ringers, called the "conductor," would call out "Bob!" at appropriate times, in order to signal that the current change should be p_5 instead of the default p_7.

There were partial successes, as many ingenious sequences of Stedman-like compositions with bobs were explored and found to yield quite a few different permutations. But for more than 200 years nobody was able to go all the way. It seemed that a full peal would require a little bit of fudging, namely a few special steps that *weren't* triple swaps according to the prescribed pattern of (2).

Finally, in 1994, the riddle of Stedman Triples With Bobs was solved at last. Colin John Edward Wyld, working by hand, composed a full peal [10] with 705 well-chosen bobs! Rumors of its existence spurred others to tackle the problem anew, notably Andrew Johnson and Philip Andrew Bruce Saddleton, whose 579-bob peal was actually the first to be rung successfully [5]. Johnson went on to discover dozens more full peals of the Stedman Triples With Bobs until his death in 2024, some with as few as 438 bobs. While doing this he developed highly sophisticated computational methods — for example, by devising new ways to exploit SAT technology for Hamiltonian cycle problems [2]. One of his noteworthy achievements [3] was a solution of the form A^3; in other words, he showed that the same sequence of $840/3 = 280$ bob calls could be repeated thrice.

Meanwhile, in 1908, a more uniform approach had been introduced by Gabriel Lindoff (1869–1941), another famous composer. He hoped to simplify Stedman's pattern by finding a full peal of the form

$$(p_3, p_1, p_3, p_1, p_3, p_{5/7})^{840} \tag{3}$$

instead. And he gave the name "Erin Triples" to such patterns, because he lived in Dublin.

The title of this note, "Erin Triples With Bobs," encapsulates the problem of discovering a full peal that matches the template (3). All we have to do is choose p_5 or p_7, wherever $p_{5/7}$ is specified. I believe that a solution almost certainly exists, somewhere among the 2^{840} possibilities. Indeed, there "must be" zillions of solutions. But none are yet known. Perhaps you, dear reader, are destined to find the first example.

Let's look closer. We can assume, without loss of generality, that the first $p_{5/7}$ is the default, p_7, and that the second one is a bob, p_5. (Because p_7 must be followed by p_5 somewhere; and we can move our starting point to anywhere we want in the cycle.) The first thirteen permutations are therefore

$$\begin{array}{ll}
1234567 & \\
 & p_3 \\
2135476 & \\
 & p_1 \\
2314567 & \\
 & p_3 \\
3215476 & \\
 & p_1 \\
3124567 & \\
 & p_3 \\
1325476 & \\
 & p_7 \\
3152746 & \\
 & p_3 \\
1357264 & \\
 & p_1 \\
1532746 & \\
 & p_3 \\
5137264 & \\
 & p_1 \\
5312746 & \\
 & p_3 \\
3517264 & \\
 & p_5 \\
5371246 &
\end{array} .$$

These permutations are alternately even and odd, because each of $\{p_1, p_3, p_5, p_7\}$ is odd. And we can ignore the odd ones, because they differ by p_3 from their predecessors; we'll have all 5040 of the permutations if and only if we have all 2520 of the even permutations. Thus we're left with

$$\begin{array}{l}
1234567 \\
2314567 \\
3124567 \\
3152746 \\
1532746 \\
5312746 \\
5371246
\end{array} .$$

Aha! Notice that the first three all end with 4567. These are, in fact, all of the even permutations that end with 4567. And they're followed by the three even permutations that end with 2746. And so on.

Let's therefore partition the set of even permutations into "clumps," where two permutations belong to the same clump if and only if they have the same last four elements. Notice that the total number of clumps is $2520/3 = 840 = 7 \cdot 6 \cdot 5 \cdot 4$.

Consider the directed graph D whose vertices are the 2520 even permutation of $\{1,2,3,4,5,6,7\}$, and whose arcs are

$$abcdefg \longrightarrow caebgdf \qquad \text{and} \qquad abcdefg \longrightarrow caebdfg \tag{4}$$

for every vertex $abcdefg$, hence 5040 arcs altogether. (For example, two of the arcs in D are $1234567 \longrightarrow 3152746 \longrightarrow 5371246$, matching the 1st, 4th, and 7th even permutation in the scenario above.) It's easy to see that *every solution to Erin Triples With Bobs corresponds to an oriented cycle of length 840 in D, where every clump appears exactly once in that cycle.** This cycle consists of every 3rd even permutation, hence every 6th permutation, of the full peal.

Every even permutation also has exactly two predecessors. Indeed, from (4) we readily deduce that

$$bdafcge \longrightarrow abcdefg \qquad \text{and} \qquad bdaecfg \longrightarrow abcdefg \tag{5}$$

are the arcs of D that lead to vertex $abcdefg$.

The fact that D has 5040 arcs suggests a one-to-one correspondence between arcs and permutations. And a bit of thought reveals a nice way to number the permutations from #0 to #5039 so that they represent the arcs conveniently in a computer, by listing the permutations in *colex order* — that is, by ordering them lexicographically from right to left, as shown in Table 1 on the next page.

Here's how to understand Table 1: All permutations appear in the 'perm' column; and their left-right reflections appear in the 'mrep' column, which is in lexicographic order. The even permutations are the ones with parity 0. All permutations with the same last four elements appear consecutively, in groups of six. Thus all clumps are easily identified, from 4321 to 4567. Permutations $\#(2k)$ and $\#(2k+1)$ differ by interchanging their first two elements; therefore one of them, say α_k, is even and the other is odd. (In symbols, $\alpha_k = \#(2k \oplus \mathrm{parity}[2k])$, where '$\oplus$' is the exclusive-or operation on binary numbers.) The two arcs leading to α_k are $\#(2k)$ and $\#(2k+1)$, and they come from permutations $\mathrm{permto}[2k]$ and $\mathrm{permto}[2k+1]$. For example, when $k = 1573$ and $2k = 3146$, the kth even permutation α_k is $\#3147 = 4761235$; and its predecessors in D are 7143652 and 7142635, via arcs #3146 and #3147, in agreement with (5).

Similarly, the two arcs leading from α_k are $\mathrm{arcfrom}[2k]$ and $\mathrm{arcfrom}[2k+1]$. Continuing our example, the successors of $\alpha_{1573} = 4761235$ are 6427513 and 6427135, via arcs #1507 and #3138, in agreement with (4).

(There's actually something tricky about the 'arcfrom' column, explained in the appendix below. That's a side issue, however. The main point is that it's easy to compute the 'parity', 'perm', and 'permto' columns of Table 1; and the 'arcfrom' column is deducible from 'parity' and 'permto'.)

Our goal is to find a certain kind of 840-cycle in D. One way to approach it is to weaken the constraints, so that the "relaxed" problem is easier to describe locally: Let's try first to find a set of one *or more* cycles in D that hit every clump exactly once. We shall call that problem "Weak Erin Triples With Bobs."

Table 1 makes it easy to formulate this less-constrained problem as an instance of SAT (a Boolean satisfiability problem), namely as a conjunction of clauses (an AND of ORs of Boolean variables or their complements): There are 5040 Boolean variables, named #0 through #5039, representing the presence or absence of each arc in D. There are 840 "at-least-one" clauses, one for each clump C_k, representing the constraint that at least one arc enters C_k:

$$\big(\#(6k) \vee \#(6k+1) \vee \#(6k+2) \vee \#(6k+3) \vee \#(6k+4) \vee \#(6k+5)\big), \tag{6}$$

for $0 \le k < 840$. Another 840 at-least-one clauses force at least one arc to leave C_k, for $0 \le k < 840$:

$$\big(\mathrm{arcfrom}[6k] \vee \mathrm{arcfrom}[6k+1] \vee \mathrm{arcfrom}[6k+2] \vee \mathrm{arcfrom}[6k+3] \vee \mathrm{arcfrom}[6k+4] \vee \mathrm{arcfrom}[6k+5]\big). \tag{7}$$

And finally, there are "at-most-one" clauses, which state that certain pairs of arcs cannot both be present. Each clump C_k has 54 such clauses, which constrain the twelve variables that appear in (6) and (7).

---

* It might well be interesting to investigate the following generalization of the classic Hamiltonian Cycle Problem: "Given a digraph whose vertices are partitioned arbitrarily into N disjoint classes, find an oriented N-cycle with exactly one vertex in each class."

Table 1: REPRESENTING THE 5040 ARCS AS 5040 PERMUTATIONS

	parity	perm	mrep	permto		arcfrom	
#0	1	7654321	1234567	#1490	(7462513)	#978	(5637142)
#1	0	6754321	1234576	#50	(7463521)	#43	(5637421)
#2	0	7564321	1234657	#76	(5473621)	#31	(6735421)
#3	1	5764321	1234675	#1516	(5472613)	#966	(6735142)
#4	1	6574321	1234756	#1538	(6452713)	#973	(7536142)
#5	0	5674321	1234765	#98	(6453721)	#36	(7536421)
#6	0	7645321	1235467	#28	(6573421)	#61	(4736521)
#7	1	6745321	1235476	#1468	(6572413)	#1092	(4736152)
#8	1	7465321	1235647	#1512	(7542613)	#1099	(6437152)
#9	0	4765321	1235674	#72	(7543621)	#66	(6437521)
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
#3140	1	6247135	5317426	#415	(6723451)	#2131	(4216573)
#3141	0	2647135	5317462	#3169	(6721435)	#3234	(4216735)
#3142	0	4267135	5317624	#3195	(2741635)	#3222	(6412735)
#3143	1	2467135	5317642	#441	(2743651)	#2119	(6412573)
#3144	0	7641235	5321467	#3178	(6172435)	#3133	(4726135)
#3145	1	6741235	5321476	#1138	(6173452)	#1500	(4726513)
#3146	1	7461235	5321647	#1160	(7143652)	#1507	(6427513)
#3147	0	4761235	5321674	#3200	(7142635)	#3138	(6427135)
#3148	0	6471235	5321746	#3226	(4162735)	#3126	(7624135)
#3149	1	4671235	5321764	#1186	(4163752)	#1495	(7624513)
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
#5030	1	4123567	7653214	#3522	(4316275)	#3953	(2154736)
#5031	0	1423567	7653241	#4962	(4315267)	#4984	(2154367)
#5032	0	2143567	7653412	#5013	(1325467)	#4972	(4251367)
#5033	1	1243567	7653421	#3573	(1326475)	#3941	(4251736)
#5034	1	3214567	7654123	#3499	(3426175)	#4072	(1253746)
#5035	0	2314567	7654132	#4939	(3425167)	#5009	(1253467)
#5036	0	3124567	7654213	#4965	(1435267)	#4997	(2351467)
#5037	1	1324567	7654231	#3525	(1436275)	#4060	(2351746)
#5038	1	2134567	7654312	#3547	(2416375)	#4067	(3152746)
#5039	0	1234567	7654321	#4987	(2415367)	#5002	(3152467)

In detail, if we denote the variables of (6) by x_k, x'_k, y_k, y'_k, z_k, z'_k, and if we denote the variables of (7) by X_k, X'_k, Y_k, Y'_k, Z_k, Z'_k, the at-most-one clauses for clump C_k are

$$\begin{aligned}&(\bar x_k \vee \bar x'_k) \wedge (\bar x_k \vee \bar y_k) \wedge (\bar x_k \vee \bar y'_k) \wedge (\bar x_k \vee \bar z_k) \wedge (\bar x_k \vee \bar z'_k) \wedge \\ &\quad (\bar x'_k \vee \bar y_k) \wedge (\bar x'_k \vee \bar y'_k) \wedge (\bar x'_k \vee \bar z_k) \wedge (\bar x'_k \vee \bar z'_k) \wedge (\bar y_k \vee \bar y'_k) \wedge \\ &\quad (\bar y_k \vee \bar z_k) \wedge (\bar y_k \vee \bar z'_k) \wedge (\bar y'_k \vee \bar z_k) \wedge (\bar y'_k \vee \bar z'_k) \wedge (\bar z_k \vee \bar z'_k)\,,\end{aligned} \tag{8}$$

meaning that at most one arc enters C_k, together with

$$\begin{aligned}&(\overline{X}_k \vee \overline{X}'_k) \wedge (\overline{X}_k \vee \overline{Y}_k) \wedge (\overline{X}_k \vee \overline{Y}'_k) \wedge (\overline{X}_k \vee \overline{Z}_k) \wedge (\overline{X}_k \vee \overline{Z}'_k) \wedge \\ &\quad (\overline{X}'_k \vee \overline{Y}_k) \wedge (\overline{X}'_k \vee \overline{Y}'_k) \wedge (\overline{X}'_k \vee \overline{Z}_k) \wedge (\overline{X}'_k \vee \overline{Z}'_k) \wedge (\overline{Y}_k \vee \overline{Y}'_k) \wedge \\ &\quad (\overline{Y}_k \vee \overline{Z}_k) \wedge (\overline{Y}_k \vee \overline{Z}'_k) \wedge (\overline{Y}'_k \vee \overline{Z}_k) \wedge (\overline{Y}'_k \vee \overline{Z}'_k) \wedge (\overline{Z}_k \vee \overline{Z}'_k)\,,\end{aligned} \tag{9}$$

meaning that at most one arc leaves C_k, together with

$$\begin{aligned}&(\bar x_k \vee \overline{Y}_k) \wedge (\bar x_k \vee \overline{Y}'_k) \wedge (\bar x_k \vee \overline{Z}_k) \wedge (\bar x_k \vee \overline{Z}'_k) \wedge (\bar x'_k \vee \overline{Y}_k) \wedge (\bar x'_k \vee \overline{Y}'_k) \wedge (\bar x'_k \vee \overline{Z}_k) \wedge (\bar x'_k \vee \overline{Z}'_k) \wedge \\ &\quad (\bar y_k \vee \overline{X}_k) \wedge (\bar y_k \vee \overline{X}'_k) \wedge (\bar y_k \vee \overline{Z}_k) \wedge (\bar y_k \vee \overline{Z}'_k) \wedge (\bar y'_k \vee \overline{X}_k) \wedge (\bar y'_k \vee \overline{X}'_k) \wedge (\bar y'_k \vee \overline{Z}_k) \wedge (\bar y'_k \vee \overline{Z}'_k) \wedge \\ &\quad (\bar z_k \vee \overline{X}_k) \wedge (\bar z_k \vee \overline{X}'_k) \wedge (\bar z_k \vee \overline{Y}_k) \wedge (\bar z_k \vee \overline{Y}'_k) \wedge (\bar z'_k \vee \overline{X}_k) \wedge (\bar z'_k \vee \overline{X}'_k) \wedge (\bar z'_k \vee \overline{Y}_k) \wedge (\bar z'_k \vee \overline{Y}'_k)\,,\end{aligned} \tag{10}$$

meaning that we must leave C_k from the permutation by which we entered.

A new technique for finding Hamiltonian cycles was introduced at last year's SAT conference [8], and it has outperformed all other approaches on a wide variety of difficult graphs. It's called a "CEGAR method enhanced by cut-set constraints," where CEGAR stands for "counterexample-guided abstraction refinement." The idea is to use SAT technology to solve a weaker problem, namely to cover the vertices with one *or more* cycles instead of insisting on a single cycle. If the solution turns out to have several cycles, we can add additional clauses that will prevent those particular ones, based on appropriate cut-sets of the graph; then we repeat the process. Eventually we'll either find a Hamilton cycle or we'll have a proof of impossibility — namely, a set of necessary clauses that cannot be satisfied.

A very similar method can be applied to the problem of Weak Erin Triples With Bobs: Suppose our SAT solver finds a solution, with cycles $\{C_1, \ldots, C_t\}$. If $t = 1$, we're done. Otherwise we can append t further necessary clauses, for $1 \le j \le t$, stating that at least one arc must proceed from a clump that's hit by C_j to a clump that isn't hit by C_j.

I got interested in Erin Triples With Bobs at the exact moment when I realized that this newfangled CEGAR method might well be the key to a solution. In fact I stopped working on everything else, and wrote a program [7] to construct clauses (6)–(11) and then to search for a solution.

Unfortunately that program has been running for a few days and it shows no sign of being close to success. And even if it does happen to come up a solution tomorrow, I'll probably have to add some cut-set-based clauses and wait another few days before knowing whether or not the second round succeeds. And so on.

In other words, I'm stuck.

I either need a better SAT solver, or a better idea, or both.

When somebody *does* solve Erin Triples With Bobs, I'm quite sure that the first performance of its full peal will be a splendid occasion indeed.

Appendix. Every entry in the 'permto' column of Table 1 is the number of an even permutation. In fact, that permutation is shown in parentheses. For example, the 'perm' column of row #1490 is 7462513, and parity[1490] = 0. Every even permutation appears exactly twice in the 'permto' column.

On the other hand, every entry in the 'arcfrom' column of Table 1 is the number of an *arc*. Every arc number appears exactly once in that column.

The permutation shown in parentheses after the 'arcfrom' column of row k is not really perm[arcfrom[k]]; it's perm[arcfrom[k]⊕parity[k]]. For example, perm[978] is 6537142 — an odd permutation — while perm[979] is 5637142. Table 1 shows '(5637142)', not '(6537142)'. (Think about it and you'll understand why.)

Incidentally, only the 'arcfrom' column of Table 1 is used in the construction of the SAT clauses that model cycles in the digraph D; not the 'parity' or 'perm' or 'mrep' columns. But the other columns are useful for interpreting the meaning of the 5040 Boolean variables, in printouts.

Related work. As far as I know, the best attempts to solve Erin Triples With Bobs so far have been by Michael Haythorpe and Andrew Johnson [1] and by Andrew Johnson [2].

Paper [1] constructs an *undirected* graph with 13440 vertices and 21840 edges, with the property that its Hamiltonian cycles correspond one-to-one with solutions of Erin Triples With Bobs.

Paper [2] constructs a SAT instance for those solutions. It has 10911 variables and 119976 clauses.

"Near solutions" have been known to be possible for many years: In the late 1980s, Philip A. B. Saddleton composed a full peal that deviates from the required template (3) in just two of the 5040 changes. (This pattern [9] is quite complex, yet it was performed successfully in 1999.)

Then Andrew Johnson, while finishing papers [2] and [1], used his SAT techniques to create a *half peal* of Erin Bobs With Triples [4]. This is a directed cycle of length 420 in D that hits every "superclump" exactly once, where a superclump is a set of six vertices that can be obtained by combining clump *abcd* with clump *abcd*(12)(56). (Thus the superclumps are $1234 \cup 2134$, $1235 \cup 2136$, ..., $1256 \cup 2165$, ..., $4521 \cup 4612$, ..., $7564 \cup 7654$.) Johnson, who called this problem 'Erin2', obtained his first solution by "seeding" the SAT solver with a subset of Saddleton's composition [9].

Any solution to Erin2 yields a full peal, if we simply change one occurrence of the triple swap $p_7 = (12)(34)(56)$ to the single swap (34), and repeat the pattern twice.

Acknowledgment. I'm grateful to Michael Haythorpe and Philip A. B. Saddleton for answering several questions. Saddleton also pointed me to references [4] and [9].

References.

[1] Michael Haythorpe and Andrew Johnson, "Change ringing and Hamiltonian cycles: The search for Erin and Stedman Triples," *Electronic Journal of Graph Theory and Applications* **7** (2019), 61–75.

[2] Andrew Johnson, "Stedman and Erin Triples encoded as a SAT problem." Easy Chair Preprint 673 (9 December 2018), 16 pages. `https://easychair.org/publications/preprint/pp38`

[3] Andrew Johnson, "Three-part bobs-only peals of Stedman Triples." Published online as `https://ajohnson1.github.io/3part.html`.

[4] Andrew Johnson, Composition Library entry 65019 `https://complib.org/composition/65019`.

[5] A. Johnson and P. A. B. Saddleton, Peals of Stedman Triples with bobs only." Published online as `https://ajohnson1.github.io/ajart.html`. (See also *The Ringing World* #4375 (3 March 1995), 227, as well as letters published that year on pages 188, 216, 244, 440.)

[6] Donald E. Knuth, *The Art of Computer Programming*, Volume 4A: *Combinatorial Algorithms, Part 1* (Upper Saddle River, N. J.: Addison–Wesley, 2011), xvi + 883 pp.

[7] Donald E. Knuth, `https://cs.stanford.edu/~knuth/programs/hamsat-erin.ch` (18 April 2026), a change file for the CWEB program `https://cs.stanford.edu/~knuth/programs/hamsat.w`.

[8] Ryoga Ohashi, Takehide Soh, Daniel Le Berre, Hidetomo Nabeshima, Mutsunori Banbara, Katsumi Inoue, and Naoyuki Tamura, "SAT-based CEGAR method for the Hamiltonian Cycle Problem enhanced by cut-set constraints." In the proceedings of the 28th international conference on theory and applications of satisfiability testing (SAT 2025), edited by Jeremias Berg and Jakob Nordström, Leibniz International Proceedings in Informatics, *LIPIcs* **341**, 24:1–24:10.

[9] Philip A. B. Saddleton, Composition Library entry 29016, `https://complib.org/composition/29016`.

[10] Colin J. E. Wyld, "A 250-year-old problem solved." *The Ringing World* #4374 (24 February 1995), 197–198.