Defying God: the Stanley-Wilf Conjecture, Stanley-Wilf Limits, and a Two-Generation Explosion of Combinatorics

Eric S. Egge
Department of Mathematics and Statistics
Carleton College
Northfield, MN 55057
eegge@carleton.edu

In March of 2005, at the Third International Conference on Permutation Patterns in Gainesville, Florida, Doron Zeilberger declared that "Not even God knows $a_{1000}(1324)$." Zeilberger's claim raises thorny theological questions, which I am happy to ignore in this article, but it also raises mathematical questions. The quantity $a_{1000}(1324)$ is the one thousandth term in a certain sequence $a_n(1324)$. God may or may not be able to compute the thousandth term in this sequence, but how far can mortals get? If we can't get beyond the fortieth or fiftieth term, can we at least approximate the one thousandth term? How fast does $a_n(1324)$ grow, anyway? And what does $a_n(1324)$ even mean?

The answers to these questions involve fast computers, fascinating mathematics, and remarkable human ingenuity. But their stories, which are ongoing, also reflect important undercurrents and developments that have influenced all of mathematics, but especially combinatorics, over the past two generations and more.

Knuth's Railroad Problems

The story of $a_{1000}(1324)$ begins with a gap in the railroad literature, which Donald Knuth began to fill in 1968 in the first edition of the first volume of his masterpiece *The Art of Computer Programming*. In the second section of Chapter 2, Knuth included several exercises exploring a problem involving sequences of railcars one can obtain using a turnaround. One of Knuth's exercises is equivalent to the following problem.

> At dawn we have n railroad cars positioned on the right side of the track in Figure 1, numbered 1 through n from right to left. During the day we gradually move the cars to the left side of the track, by moving each car into and back out of the turnaround area. There can be any number of cars in the turnaround around area at a time, and at the end of the day the cars on the left side of the track can be in many different orders. Each possible order determines a permutation of the numbers $1, 2, \ldots, n$. Show that a permutation $\pi_1, \ldots, \pi_n$ (this time reading from left to right along the tracks) is attainable in this way if and only if there are no indices $i < j < k$ such that $\pi_i < \pi_k < \pi_j$.

The solution to this problem is a fun exercise in careful bookkeeping. If such a subsequence exists, then consider the situation when car π_i enters the turnaround. Since π_i is the smallest of our

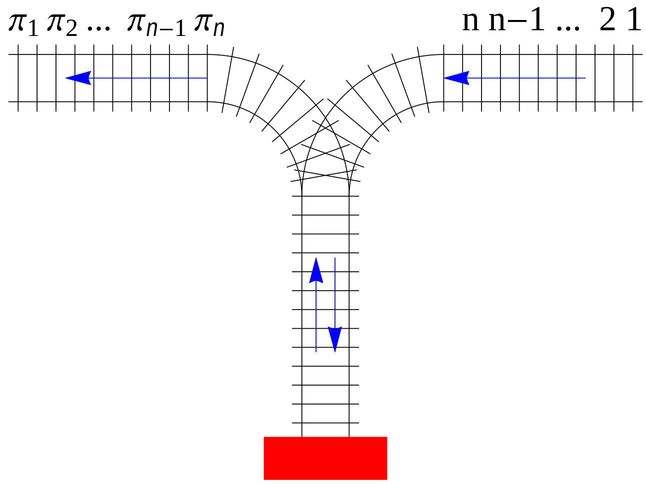

Figure 1: Knuth's railroad tracks.

three car numbers, cars π_j and π_k have already entered the turnaround, in that order. Furthermore, in order for car π_i to appear to the left of cars π_j and π_k, cars π_j and π_k must both still be in the turnaround when car π_i enters. But now cars π_j and π_k will leave the turnaround in the wrong order.

Conversely, suppose we have a target permutation $\pi_1, \ldots, \pi_n$ with no subsequence of the forbidden type. We can always move π_1 into position, and when car π_1 leaves the turnaround, the cars in the turnaround are, from bottom to top, $n, n-1, \ldots, \pi_1 + 1$. Now notice that π_2 cannot be larger than $\pi_1 + 1$, since this would mean π_1, π_2, and $\pi_1 + 1$ form a forbidden subsequence. So if π_2 is in the turnaround, then it is the top car there. Either way, we can move car π_2 into position. In general, if we have just moved car π_j into position, and b is the smallest entry greater than π_j which has not yet left the turnaround, then $\pi_{j+1} \leq b$, since otherwise π would have a forbidden subsequence π_j, π_{j+1}, b. Therefore, if π_{j+1} has entered the turnaround then it is the top car there, and we can move it into place.

Knuth was interested in this railcars problem because it models the data structure commonly called a stack, which arises in numerous programming problems, so he introduced no particular notation for the permutations he obtained. Indeed, no general notation for these permutations appeared in print until 1985, when Simion and Schmidt [31] published the first systematic study of permutations with forbidden subsequences of the type Knuth uses.

Today, if π and σ are permutations of lengths n and k respectively, then we say a subsequence of π of length k has *type* σ whenever its entries are in the same relative order as the entries of σ. For example, the subsequence 829 of the permutation 718324695 has type 213, since its smallest entry is in the middle, its largest entry is last, and its middle entry is first. In this context, we say π *avoids* σ, or π is σ-avoiding, whenever π has no subsequence of type σ, and we write $S_n(\sigma)$ to denote the set of all permutations of length n which avoid σ. We might also say that σ is a *forbidden subsequence* or a *forbidden pattern*. With this terminology, the permutations Knuth obtains with his railcars are exactly the 132-avoiding permutations, and the term $a_{1000}(1324)$ that Zeilberger's God finds so perplexing is none other than the size of $S_{1000}(1324)$. In Table 1 we have the first ten values of $|S_n(1324)|$.

n	0	1	2	3	4	5	6	7	8	9
$\lvert S_n(1324)\rvert$	1	1	2	6	23	103	513	2762	15793	94776

Table 1: The first ten values of $|S_n(1324)|$.

A Scattered History

Of course, combinatorial problems predate computers, the MAA, and even railroads. In fact, we can find evidence of people studying combinatorial problems nearly as far back in time as we can find evidence of people doing mathematics. For example, in the following quotation from Plutarch's *Table-Talk* [25, VIII.9, 732], we find Chrysippus (circa 200 BCE) and Hipparchus (circa 300 BCE) discussing how to form logical expressions.

> Chrysippus says that the number of compound propositions that can be made from only ten simple propositions exceeds a million. (Hipparchus, to be sure, refuted this by showing that on the affirmative side there are $103,049$ compound statements, and on the negative side 310,952.)

Plutarch doesn't say what a compound proposition is, but it seems reasonable to assume it's a combinatorial object of some kind. Indeed, in 1994 David Hough, who was then a graduate student at George Washington University, observed that 103,049 is the tenth small Schröder number [34, 17, 1]. This means 103,049 is, among other things, the number of ways to parenthesize a sequence of eleven letters. To give a sense of what this means, there are three ways to parenthesize a sequence of three xs, namely xxx, $(xx)x$, and $x(xx)$, and eleven ways to parenthesize a sequence of four xs, namely $xxxx$, $(xxx)x$, $((xx)x)x$, $(x(xx))x$, $x(xxx)$, $x((xx)x)$, $x(x(xx))$, $(xx)(xx)$, $(xx)xx$, $x(xx)x$, and $xx(xx)$.

Plutarch's account of Chrysippus and Hipparchus's debate over the number of compound propositions is not the most ancient combinatorial reference, nor does it involve the most common combinatorial quantities. In *Sushruta Samhita*, a sixth century BCE Sanskrit text on surgery attributed to the Indian physician Sushruta, the author observes that we can make sixty-three combinations out of six different tastes, when we take them one at a time, two at a time, etc. This gives us all but one of the entries of the sixth row of Pascal's triangle, more than two millennia before Pascal. And Sushruta's discussion of what would come to be called the binomial coefficients is not an isolated occurrence in ancient Indian literature. The *Bhagabati Sutra*, a religious text of the Jains which appeared around 300 BCE, contains a more general rule for computing binomial coefficients, and the Jain mathematician Mahavira gives a completely general rule in his *Ganita Sara Sangraha*, which was written around 850 AD [14, p. 27].

Combinatorics is ancient, but for much of its history it has also been scattered, arising in diverse contexts but having few or no adherents of its own. Hipparchus solves a combinatorial problem to refute Chrysippus, but he does it to make a larger point. Sushruta solves a combinatorial problem in the midst of a landmark text on surgery, which barely contains any other mathematics at all. And this is how it goes for centuries: Euler invents graph theory to solve the Königsberg bridge problem, Pascal (re)discovers the binomial coefficients in his quest to resolve interrupted games of chance, Kempe "proves" the four color theorem to answer a question posed almost thirty years earlier by a student trying to color a map of England, Cayley studies partitions as a tool in invariant theory, Young introduces the tableaux that now bear his name to build on Cayley's work in invariant theory, and Kirchhoff proves his matrix-tree theorem for counting spanning trees in a graph to solve a problem in electrical engineering.

To be sure, some combinatorial problems were studied for their own sake a century or two ago, and there were even people we would recognize today as combinatorialists. For instance, in the early 1880s J. J. Sylvester and his students devoted themselves to the study of partitions, pioneering the use of Ferrers diagrams to develop a general theory [24]. But in the nineteenth century and the first half of the twentieth, Sylvester's work and life are the exceptions, not the rule, as most combinatorial problems were treated as isolated curiosities, not fundamental examples around which one might build a theory.

How Fast Does the Number of Knuth Railcar Permutations Grow?

Fortunately for us, the problem of estimating the rate at which $|S_n(1324)|$ grows is not an isolated curiosity, but is instead part of a general theory, which means we can gain some insight by looking at a related problem. For instance, since 132 is part of 1324, every permutation which avoids 132 also avoids 1324, so let's try to first estimate $|S_n(132)|$. To get a crude upper bound on this quantity, first recall that a *left-to-right minimum* in a permutation π is an entry which is smaller than every entry to its left. For example, the left-to-right minima in 694853127 are 6, 4, 3, and 1. Somewhat surprisingly, if π avoids 132, then it is determined by the values and positions of its left-to-right minima. For example, suppose $\pi \in S_9(132)$ has left-to-right minima 6, 3, and 1, which are in the first, third, and fifth positions. Since the left-to-right minima must be in decreasing order, we can start to construct π as in Figure 2. Now the second entry of π must be

$$6 \_ 3 \_ 1 \_ \_ \_ \_$$

Figure 2: A permutation $\pi \in S_9(132)$ with prescribed left-to-right minima.

greater than 6, since it's not a left-to-right minimum. But if we put 8 (resp. 9) there, then the 6, the 8 (resp. 9), and the 7, which must appear somewhere to the right, will form a subsequence of type 132. Therefore, the second entry must be 7. Similarly, the fourth entry must be the smallest available number greater than 3, and each successive entry is the smallest available number which is greater than the nearest left-to-right minimum on its left. Following this recipe, we find that in our example $\pi = 673412589$.

Returning to our bound on $|S_n(132)|$, the leftmost entry of a permutation is always a left-to-right minimum, as is 1, so there are at most 2^{n-1} sets of values for the left-to-right minima, and 2^{n-1} sets of positions for these values. Therefore, $|S_n(132)| \leq 4^{n-1}$ when $n \geq 1$. Since 4^{n-1} is much less than $n!$ when n is even reasonably large, the chances we can put our railcars into any particular order is almost zero for any train more than a few cars long.

We expect our 4^{n-1} bound on $|S_n(132)|$ to be terrible, because there are many ways a given choice of numbers and positions can fail to be the values and positions of the left-to-right minima for any permutation. For example, we made no effort to ensure that we have the same number of left-to-right minima as we have positions for them. And even when we have the same number of values as positions, the values and positions may not be compatible. For instance, if we insist on having left-to-right minima 1, 5, and 7 in the first, third, and eighth positions of a permutation of length 9, then we're going to have a bad time, since at least one of 2, 3, and 4 must appear between the 5 and the 1, and this entry will be another left-to-right minimum. Nevertheless, a miracle of sorts occurs: while it is possible to improve our bound on $|S_n(132)|$ to $\dfrac{4}{n^{3/2}\sqrt{\pi}}4^{n-1}$, the base 4 of the exponential factor cannot be replaced with a smaller number.

To see why no base less than four will work, we return to Knuth, who doesn't bother to estimate $|S_n(132)|$ at all in his solution to the railroad problem. Instead, he notes that we can encode each permutation in the set uniquely by writing down the sequence of moves that generate it. In particular, if we write N each time we move a railcar into the turnaround and E each time we move a railcar out of the turnaround on the other side, then we get a bijection between $S_n(132)$ and the set of sequences of n Ns and n Es in which every initial segment has at least as many Ns as it has Es. Such sequences are called *ballot sequences*, and it was already well known in 1969 that they are counted by the Catalan number $C_n = \frac{1}{n+1}\binom{2n}{n}$. In particular, we have $\frac{C_n}{C_{n-1}} = \frac{4n-2}{n+1}$, so $\lim_{n\to\infty} \frac{C_n}{C_{n-1}} = 4$ and $\lim_{n\to\infty} \sqrt[n]{|S_n(132)|} = 4$.

Introducing the notion of the type of a subsequence, and the associated idea of pattern avoidance, opens a panorama of fruitful generalizations and questions. For example, for any permutation π of length n, let π^r denote the *reverse* of π, which has $\pi^r(j) = \pi(n+1-j)$ for $1 \leq j \leq n$. Similarly, let π^{-1} denote the group-theoretic inverse of π. Then it's not hard to show that π avoids σ if and only if π^r avoids σ^r, which occurs if and only if π^{-1} avoids σ^{-1}. Combining these, we find $|S_n(132)| = |S_n(213)| = |S_n(231)| = |S_n(312)|$ and $|S_n(123)| = |S_n(321)|$ for all n. Furthermore, it's already implicit in work of Percy MacMahon [21] in 1915 and Craige Schensted [30] in 1961 that $|S_n(123)| = C_n$. As a result, we have $\lim_{n\to\infty} \sqrt[n]{|S_n(\sigma)|} = 4$ for all $\sigma \in S_3$. All of which leads to the natural question: what happens if σ is bigger?

The Rise of the Machines in Combinatorics

Computer programming questions motivated Knuth's railcars problem, and they continued to drive the study of patterns in permutations for the next decade and a half. Rotem [29] was the next person to look at permutations which can be produced with one pass through a stack. He called these permutations *stack-sortable*, and in his main results he gave a bijection from these permutations to binary trees, which he used to analyze the average length of a monotonic subsequence in these permutations. Even after Simion and Schmidt kicked off the study of restricted permutations for their own sake, others continued to investigate restricted permutation problems arising from data structures. The next major step in this direction came in Julian West's thesis [39], where he studied permutations one can obtain using two passes through a stack, which he called *2-stack-sortable*. Relying on the language of pattern-avoidance, West showed that these permutations are the ones which avoid 2341 and whose only subsequences of type 3241 are those which are contained in a subsequence of type 35241. West also conjectured that the number of 2-stack-sortable permutations of length n is $\dfrac{2}{(n+1)(2n+1)}\dbinom{3n}{n}$, which Zeilberger proved [41] shortly after hearing West describe his conjecture at a talk in May of 1991. Since then a thriving industry has developed, in which people use the language of permutation patterns to study permutations generated by parallel queues, dequeues, networks of stacks, ordered stacks, token passing in graphs, and even forklifts, to name just a few. There are more relevant references than I can comfortably list here, but the interested reader would do well to start with the summary in Section 2.1 of Kitaev's encyclopedic monograph [19].

Questions from computer science lead to interesting problems throughout combinatorics, but for the past thirty years computers have played a more explosive role in the field, letting mathematicians generate reams of data that would be impossible to produce by hand. These data, in turn, suggest new conjectures, and eventually lead to new theorems. Simion and Schmidt's paper, for example, enumerates more than a dozen families of restricted permutations, with sequences like powers of two, the Fibonacci numbers, the central binomial coefficients $\binom{n}{\lfloor \frac{n}{2} \rfloor}$, and the triangular numbers plus one. With a computer, one can essentially peek in the back of the book, generating the first ten or so terms of each sequence, and then seeing what needs to be proved. West's conjecture on the number of 2-stack sortable permutations of length n would also be much harder to discover without computational assistance: one could certainly write down the 91 2-stack-sortable permutations of length five in half an hour or so, and perhaps get the 408 2-stack-sortable permutations of length six in another hour, but the 1938 2-stack-sortable permutations of length seven would present a challenge, and examining all 40320 permutations of length eight to find the 9614 which are 2-stack-sortable would take days, or even weeks, to do by hand. With a computer, obtaining all of these values and more is the work of a pleasant afternoon.

When evening comes, and our programming work has told us that our sequence begins with the terms $1, 2, 6, 22, 91, 408, 1938$, and 9614, we still need to formulate a conjecture about the general term before we can make more progress. Here, too, computers have become indispensable. Forty years ago we might have asked a dozen of our closest friends whether they had ever seen this sequence. Sometimes, this actually worked! For example, David Robbins first discovered that alternating-sign matrices are connected with the descending plane partitions first introduced by George Andrews by asking Richard Stanley whether he had ever seen the sequence $1, 2, 7, 42, 429, 7436, \ldots$ [28]. But most of the time, our friends' memories are no better than our own.

In the area of remembering integer sequences, computers have completely replaced human beings. Neil Sloane began this process in 1964, when he started collecting sequences he encountered on index cards, and in 1973 he published his collection in his book *A Handbook of Integer Sequences*, which included 2372 sequences. In 1995 Sloane and Simon Plouffe published a second edition of Sloane's book, with the new title *The Encyclopedia of Integer Sequences*, which included 5487 sequences. Books can't keep up with computers, though, and in 1996 Sloane established the On-Line Encyclopedia of Integer Sequences [23], which he singlehandedly maintained until 2002. The OEIS, as it is affectionately known, is now run as a wiki by a devoted group of users. It grows by more than 10000 entries each year, and would require more than 750 volumes of more than 500 pages each if it were published in book form today.

When we enter our terms $1, 2, 6, 22, 91, 408, 1938, 9614$ into the OEIS search box, we are rewarded with a description of the sequence that includes several more terms, a general formula, code for generating even more terms, a list of related references, comments from other users about the sequence, families of objects the sequence is known or conjectured to count, and more. In hindsight it's amusing that Sloane called his book "A Handbook," as though there might be competitors. There are none, and the OEIS is a required stop for anyone who encounters an integer sequence they don't recognize. It's no exaggeration to observe that in certain parts of combinatorics, the OEIS alone has increased the rate of new discoveries by an order of magnitude.

The Stanley-Wilf Conjecture

In keeping with the prevailing sense that combinatorics was just an offshoot of other subjects, few papers even mentioned patterns in permutations before 1985, and those that did generally followed Knuth in approaching the subject from a computer science point of view. Nevertheless, throughout the 1970s and into the 1980s, combinatorics was becoming a proper, and popular, subject of study in its own right, and the study of patterns in permutations was a current in this wave of acceptance. In particular, by 1980 a handful of people behind the scenes had begun to ask what sorts of sequences can appear as $|S_n(\sigma)|$, and in particular, how fast such a sequence can grow.

Just as we bounded $|S_n(132)|$ with an exponential function, we can also bound $|S_n(12\cdots k)|$ with an exponential function in an elementary way. Inspired by the classic pigeonhole principle proof of the Erdős-Szekeres theorem on increasing and decreasing subsequences in permutations, we first label each entry m of a given $\pi \in S_n(12\cdots k)$ with the length of the longest increasing subsequence of π whose last entry is m. In Figure 3 each entry of the permutation 381426975 has

1	2	1	2	2	3	4	4	3
3	8	1	4	2	6	9	7	5

Figure 3: Labeling the entries of the permutation 381426975.

its label above it in red. Since π has no increasing subsequence of length k, all of these labels will be among $1, 2, \ldots, k-1$. In addition, for each j the entries with label j will be in decreasing order: if an entry labelled j has a larger entry to its right, then that larger entry must be labelled at least $j+1$. Therefore π is a disjoint union of $k-1$ or fewer decreasing subsequences. There are $(k-1)^n$ ways to choose one of these subsequences for each of $1, 2, \ldots, n$ to be in, and there are no more than $(k-1)^n$ ways to choose which positions belong to which subsequences. Therefore, $|S_n(12\cdots k)| \leq (k-1)^{2n}$. In fact, in 1981 Regev used [27] analytical machinery to prove results which imply

$$\lim_{n\to\infty} \sqrt[n]{|S_n(12\cdots k)|} = (k-1)^2. \tag{1}$$

However, the relationship between Regev's work and subsequences of permutations is not visible to the naked eye, in part because his actual results have to do with certain sums arising in the study of the representations of the symmetric group.

Unaware at first of Regev's work and the bound on $|S_n(12\ldots k)|$, around 1980 Herb Wilf asked whether $|S_n(\sigma)| \leq (k+1)^n$ for all $\sigma \in S_k$. At nearly the same time, Richard Stanley independently asked whether $\lim_{n\to\infty} \sqrt[n]{|S_n(\sigma)|} = (k-1)^2$ for all $\sigma \in S_k$. Perhaps responding to Stanley's conjecture or Regev's results, Wilf soon asked whether there exists, for each permutation σ, a finite constant $c(\sigma)$ such that $\lim_{n\to\infty} \sqrt[n]{|S_n(\sigma)|} = c(\sigma)$. Eventually, two conjectures emerged: the upper bound conjecture and the limit conjecture, which together came to be known as the Stanley-Wilf conjecture.

Conjecture 1 (The Stanley-Wilf Upper Bound Conjecture)**.** *For every permutation σ, there is a real number $c(\sigma)$ such that $|S_n(\sigma)| \leq c(\sigma)^n$.*

Conjecture 2 (The Stanley-Wilf Limit Conjecture)**.** *For every permutation σ, there is a real number $c(\sigma)$ such that $\lim_{n\to\infty} \sqrt[n]{|S_n(\sigma)|} = c(\sigma)$.*

It's routine to show that the limit conjecture implies the upper bound conjecture, and in 1999 Arratia showed [4] that these two conjectures are equivalent. More specifically, Arratia gave a simple combinatorial proof that for any permutation σ and all $m, n \geq 1$, we have $|S_{n+m}(\sigma)| \geq |S_n(\sigma)||S_m(\sigma)|$. This implies that $f(n) = -\ln(|S_n(\sigma)|)$ has the property that $f(n+m) \leq f(n) +$

$f(m)$ for all $m, n \geq 1$; we call such functions *subadditive* [40]. Now we can use a result, Fekete's subadditivity lemma, which says that if f is subadditive then $\lim_{n\to\infty} \frac{f(n)}{n}$ exists, though it might be $-\infty$. This means $\lim_{n\to\infty} -\frac{\ln(|S_n(\sigma)|)}{n}$ exists, so $\lim_{n\to\infty} \sqrt[n]{\frac{1}{|S_n(\sigma)|}}$ exists. The upper bound conjecture implies that $\sqrt[n]{\frac{1}{|S_n(\sigma)|}}$ is bounded below by $\frac{1}{c(\sigma)}$, and the limit conjecture follows.

For nearly a quarter of a century the Stanley-Wilf conjecture stood against all who tried to prove it. Indeed, by the dawn of the new millennium no proof was in sight, though many of the problem's sharp corners had been chipped away. Miklós Bóna sheared off one of the sharpest of these corners in 1999, when he showed that the Stanley-Wilf conjecture holds for layered permutations [7]. Bóna's result is an extension of the fact that the conjecture holds for monotone permutations, since a *layered permutation* is one obtained by first listing the smallest ℓ_1 numbers in decreasing order, then listing the next smallest ℓ_2 numbers in decreasing order, etc. For example, 321765498 is the layered permutation with $\ell_1 = 3$, $\ell_2 = 4$, and $\ell_3 = 2$. Noga Alon and Ehud Friedgut [3] chiseled away another sharp corner in 1999, when they showed that the Stanley-Wilf conjecture holds for all permutations consisting of an increasing sequence followed by a decreasing sequence, as well as for all permutations which avoid 123.

Fissures one might use to break the problem open were also noted throughout the 1990s. For instance, Alon and Friedgut showed [3] that the Stanley-Wilf conjecture almost holds, by showing that $|S_n(\sigma)| \leq c(\sigma)^{\gamma(n)n}$ for a certain slow-growing function γ, which they constructed from the inverse of the Ackermann function. In the same paper, Alon and Friedgut also proved that if there is a linear upper bound on the lengths of certain words over an ordered alphabet, then the Stanley-Wilf conjecture will follow.

Another fissure in the problem appeared in the early 1990s, but it was only visible from a certain angle. To describe this new point of view, we first recall that each permutation $\pi \in S_n$ has an associated *permutation matrix*, namely the $n \times n$ matrix of 0s and 1s whose ijth entry is 1 if $\pi(i) = j$ and 0 if $\pi(i) \neq j$. It is not difficult to translate the notions of pattern avoidance and containment we've been discussing into the language of permutation matrices, but our new point of view actually depends on a different, more permissive definition of containment. Specifically, we say an $m_A \times n_A$ matrix A of 0s and 1s *contains* an $m_B \times n_B$ matrix B of 0s and 1s whenever there is an $m_B \times n_B$ submatrix B_1 of A such that if the ijth entry of B is 1 then the ijth entry of B_1 is 1. For example, the matrix $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$ contains the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ exactly twice: once in red here $\begin{pmatrix} \color{red}{1} & 1 & \color{red}{1} \\ \color{red}{1} & 0 & \color{red}{1} \\ 1 & 1 & 0 \end{pmatrix}$ and once in red here $\begin{pmatrix} \color{red}{1} & \color{red}{1} & 1 \\ 1 & 0 & 1 \\ \color{red}{1} & \color{red}{1} & 0 \end{pmatrix}$.

One can ask the same question about matrices that we have asked about permutations, namely, how many $m \times n$ matrices do not contain a given matrix? However, the fact that our matrices need not contain any particular number of 1s inspired Zoltán Füredi and Péter Hajnal [16] to ask a different question: if we have a certain matrix C, all of whose entries are 0s and 1s, how many 1s can we put into an $n \times n$ matrix before it must contain C? Füredi and Hajnal answered this question for a variety of specific C, some of which were permutation matrices, and others of which were not. At the end of their paper, tucked in among several other open questions, they asked whether there exists, for any permutation matrix C, a constant $c(C)$ such that the number of 1s an $n \times n$ matrix can contain before it must contain C is bounded above by a linear function of n, namely $c(C) \cdot n$.

Füredi and Hajnal asked the question, but it was Martin Klazar who, in 2000, promoted Füredi and Hajnal's question to the status of a conjecture.

Conjecture 3 (The Füredi-Hajnal Conjecture)**.** *For every permutation matrix C, there is a real number $c(C)$ such that if an $n \times n$ matrix of 0s and 1s contains at least $c(C) \cdot n$ entries equal to 1, then it contains C.*

Klazar's promotion of this idea from a question to a conjecture may have been justified by his main result [20], which is that the Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. The problem had another hairline crack, but it wasn't clear whether anyone could get some explosives, or a crowbar, or even a chisel, into it.

More Machines in Combinatorics

As we mentioned earlier, over the past two generations computers have become essential tools that many mathematicians, and especially combinatorialists, use to generate mathematical data. These data lead to conjectures, and, if all goes well, to theorems. But computers can be more than just data-generating devices; more and more frequently, computers play a central role in generating conjectures, and even in proving theorems.

The first, and certainly most famous, proof in which computers played a substantial role is Appel and Haken's 1976 proof of the four color theorem. Although several people have found ways to streamline this proof over the past forty years, its overall structure remains the same: this is a proof by several hundred cases. More specifically, to prove the four color theorem we first show that if G is a planar graph which cannot be properly colored with four colors, and every planar graph with fewer vertices than G can be properly colored with four colors, then G must contain one of several hundred configurations of vertices. We then show that for each of these configurations, it's possible to reduce the number of vertices in G to obtain a smaller counterexample. The computer's role in this proof is twofold: it uses human-generated heuristics to help find a family of unavoidable configurations of vertices, and it aids in constructing proofs that each configuration can be reduced. (For more details on the four color theorem, its history, and its proof, see Robin Wilson's contribution to this volume.)

Since Appel and Haken's work, combinatorialists enumerating pattern-avoiding permutations have found ways to use computers that go beyond generating data (though this remains an important computer task) or checking numerous cases. For example, in the first few steps of his proof of West's formula for the number of 2-stack-sortable permutations of length n, Zeilberger used Maple to find a functional equation for the generating function for these permutations with respect to length and another statistic. But Zeilberger took the ability of computers to do mathematics to a new level in 1998, when he taught his computer Shalosh B. Ekhad to enumerate pattern-avoiding permutations. More specifically, Zeilberger wrote code that enables his computer to generate, for each set R of forbidden patterns an object called an *enumeration scheme*, which can be readily converted into a recursive formula for $|S_n(R)|$ [42]. Following Zeilberger's construction, Shalosh B. Ekhad was able to recover most of the enumerations then known, and even to discover (and in the process prove) some new ones. Since then Vince Vatter [37], Lara Pudwell [26] and others have refined Zeilberger's work on pattern-avoiding permutations, and extended it to more general objects and definitions of pattern containment.

The Proof of the Füredi-Hajnal Conjecture: Too Nice to be True

At the turn of the millennium the Stanley-Wilf conjecture was widely seen as one of the most important and difficult open problems in the study of permutation patterns, and several of the top researchers in the field were trying everything they could to crack the problem. In spite of their efforts, by late 2003 they had made no further progress. Meanwhile, unbeknownst to anyone studying permutation patterns, Adam Marcus and Gábor Tardos had become interested in a nice collection of questions involving 0-1 matrices that Füredi and Hajnal had posed a decade before. Although Marcus and Tardos didn't know it yet, these questions included the one Klazar had elevated to conjecture status. It didn't seem to Marcus and Tardos that these questions had attracted much attention, even though they were intriguing and approachable. Marcus and Tardos didn't care about permutations in particular, and they had never heard of the Stanley-Wilf conjecture, but they did find they could make progress on one of Füredi and Hajnal's questions involving permutation matrices. Soon they had settled one of Füredi and Hajnal's questions: they had shown that if C is a permutation matrix then there is a real number $c(C)$ such that if an $n \times n$ matrix of 0s and 1s contains at least $c(C)n$ entries equal to 1, then it contains C [22]. But they didn't hear the crowd cheering their result until weeks later, when Marcus found some of Klazar's other work related to Füredi and Hajnal's paper, and Klazar told him about the Stanley-Wilf conjecture.

Marcus and Tardos's proof of the Füredi-Hajnal conjecture was a major advance in the study of patterns in permutations, but it's not a long proof, and it doesn't involve any complicated technical machinery. It is a product of human ingenuity, not fast computers. Nevertheless, it took the permutation patterns community by storm. Bóna had just submitted the first edition of his book *Combinatorics of Permutations* [8] to the publisher when he learned of Marcus and Tardos's proof. Rather than publish a book that would be out of date before it went to press, Bóna insisted that the publisher wait for him to add a new section on Marcus and Tardos's work before moving ahead with publication. The proof is accessible enough that Zeilberger has given a one-page account in the personal journal of Ekhad and Zeilberger [43], and Tardos gave the entire proof in detail, on the blackboard, in a 45-minute talk at the Second International Permutation Patterns conference in Nanaimo, British Columbia, in early July of 2004. Marcus and Tardos's proof is so elegant that when they first discovered it, Marcus and Tardos themselves thought it was too simple to be correct. They gave themselves the weekend to find the error they were certain was hidden within.

We need not give all of the details of Marcus and Tardos's proof here, but we can outline it. Suppose C is a $k \times k$ permutation matrix, and for each positive integer n, let $f(n)$ denote the largest number of 1s an $n \times n$ matrix of 0s and 1s can contain without containing a copy of C. For simplicity, let n be a multiple of k^2. If M is an $n \times n$ matrix of 0s and 1s which does not contain C, then divide it into contiguous $k^2 \times k^2$ blocks. Call a block *wide* whenever it has at least k columns which contain a 1, and call a block *tall* whenever it has at least k rows which contain a 1. Essentially, Marcus and Tardos show that there are few wide blocks and few tall blocks, and that blocks which are neither wide nor tall contain few 1s.

Let's look at the blocks which are neither wide nor tall first. Any block with more than $(k-1)^2$ 1s must be tall or wide, so the number of 1s in a block which is not tall or wide is at most $(k-1)^2$. But we can also bound the number of these blocks which contain any 1s at all. In particular, if there are more than $f\left(\frac{n}{k^2}\right)$ blocks which contain a 1, then there is a pattern of type C of these blocks. By choosing appropriate rows and columns, we can select each of the 1s we need to form a copy of C in M, since C has exactly one 1 in every row and column. Therefore, there are at most $f\left(\frac{n}{k^2}\right)$

blocks which are neither tall nor wide, but which still contain a 1. Combining these observations, we find that the blocks which are neither wide nor tall contain no more than $(k-1)^2 f\left(\frac{n}{k^2}\right)$ 1s in total.

Now let's consider the wide blocks and the tall blocks. Each block is $k^2 \times k^2$, so each wide block and each tall block contains at most k^4 1s. Marcus and Tardos use the pigeonhole principle to show that there are no more than $\frac{n}{k}\binom{k^2}{k}$ wide blocks and no more than $\frac{n}{k}\binom{k^2}{k}$ tall blocks, so these blocks contain at most $2k^2\binom{k^2}{k}n$ 1s in total. Combining these estimates with our earlier work, we find

$$f(n) \leq (k-1)^2 f\left(\frac{n}{k^2}\right) + 2k^2\binom{k^2}{k}n.$$

Now it's not hard to show by induction that

$$f(n) \leq 2k^4\binom{k^2}{k}n. \tag{2}$$

The coefficient of n on the right side of (2) is large as a function of k, but we have what Füredi and Hajnal (through Klazar) promised: a linear bound on $f(n)$.

Undergraduates in Combinatorics

Marcus and Tardos's proof of the Füredi-Hajnal conjecture is remarkable in many ways, one of which is that Marcus had barely finished his undergraduate work when they found it. In particular, Marcus spent his junior year in 2001-02 as an undergraduate in the Budapest Semesters in Mathematics program, and after graduating from Washington University, he returned to Budapest on a Fulbright fellowship in 2003. It was during Marcus's Fulbright year that he and Tardos proved the Füredi-Hajnal conjecture, which places their work in the middle of another crucial development in the explosion of combinatorics that has taken place over the past two generations: student research, and especially undergraduate research.

The tradition of involving students in combinatorics research actually goes back more than a century: Sylvester did much of his foundational work on the theory of partitions [36] in collaboration with the nine graduate students who were taking a class with him on the subject at Johns Hopkins University in the spring of 1882. But the practice was rare in the United States when the National Science Foundation began funding URP (Undergraduate Research Program) in 1958, and undergraduate research in mathematics was still not widespread when the NSF began funding mathematics REUs in 1987. Since then, however, undergraduate research in mathematics has blossomed: there were eight NSF-funded REU sites in mathematics in 1987, and by 2003 there were more than fifty. Combinatorics has played a central role in this flowering. To see how extensively combinatorics is intertwined with undergraduate research, one need look no further than the NSF's own website: of the forty-seven REU sites listed there as being funded for the summer of 2014, twenty-three include projects on combinatorics and/or graph theory. (For more details on the history of undergraduate research in mathematics in the United States, see Joseph Gallian's contribution to this volume.) Combinatorics and graph theory are also well-represented among undergraduate research award winners: in three of the last four years a Morgan Prize winner or runner-up has been recognized for her or his substantial work in combinatorics.

Not only is undergraduate research ubiquitous in combinatorics and vice versa, but the quality of the contributions undergraduates have made to the subject has been remarkably high. In 1989 Bill

Doran contributed [13] a key piece of the enumeration of the totally symmetric self-complementary plane partitions which fit in a $2n \times 2n \times 2n$ box, by explaining how to reformulate these objects as collections of nonintersecting lattice paths, which could then be counted using determinants or permanents of matrices of binomial coefficients. In 2002 Joshua Greene won the Morgan Prize for his simplified proof of the Kneser conjecture, which involves the chromatic numbers of certain Kneser graphs. Undergraduates have also contributed substantially to the study of patterns in permutations, even beyond Marcus and Tardos's work. For instance, in 2004 Reid Barton won the Morgan Prize for his work on the number of copies of a given pattern which can be packed into a permutation of a given length, and in the early 1990s Zvezdelina Stankova gave combinatorial proofs [32, 33] that $|S_n(4132)| = |S_n(3142)|$ and $|S_n(1234)| = |S_n(4123)|$ for all $n \geq 0$. These last two results will save us a substantial amount of work when we start to look beyond the results of Marcus and Tardos, by investigating the actual values of $\lim_{n\to\infty} \sqrt[n]{|S_n(\pi)|}$ for various permutations π.

Stanley-Wilf Limits

Marcus and Tardos's proof of the Füredi-Hajnal conjecture made it official: for every permutation π there is a constant $L(\pi)$ such that $\lim_{n\to\infty} \sqrt[n]{|S_n(\pi)|} = L(\pi)$. Today we refer to the number $L(\pi)$ as a *Stanley-Wilf limit*, and we'd like to know the value of $L(\pi)$ for as many permutations π as possible. We have seen that $L(123) = L(132) = 4$, and we've also seen how to use some easy symmetries to show that $L(\pi) = 4$ for every $\pi \in S_3$. In addition, we have seen evidence that $L(1234\cdots k) = (k-1)^2$, a fact which follows from the work of Regev [27].

The values of $L(\pi)$ given above are the only values that were known by the mid 1990s, and they all support the idea that $L(\pi)$ depends only on the length of π. In other words, it appears that all permutations of a given length are equally difficult to avoid, at least asymptotically. In fact, these data are also consistent with Stanley's much older conjecture that $L(\pi) = (k-1)^2$ for any permutation π of length k. It turns out, though, that making true conjectures about Stanley-Wilf limits is much harder than making false ones. Indeed, in his thesis [6] Bóna used a connection between 1342-avoiding permutations and plane trees to find an exact formula for $|S_n(1342)|$ as a sum of terms involving binomial coefficients, and from this formula he extracted the fact that $L(1342) = 8$. This is shocking! Not only does this result disprove two longstanding conjectures at once, it also says that at a fundamental level, 1342 is harder to avoid than 1234.

With $L(1234)$ and $L(1342)$ in hand, we might think that we have 22 more permutations of length four to consider. However, thanks to the symmetries we mentioned earlier and work of West and Stankova, we know that if π is permutation of length four, then $|S_n(\pi)|$ is equal to one of $|S_n(1234)|$, $|S_n(1423)|$, or $|S_n(1324)|$ for all n. As a result, we have just one more Stanley-Wilf limit to compute for a forbidden pattern of length four, namely $L(1324)$. All of which means that we have returned to our original question: how fast does $a_n(1324)$ grow, anyway?

We can start to get a feel for the growth rate of $a_n(1324)$ by looking at the data in Table 2, where we have the values of $\sqrt[n]{|S_n(1324)|}$ when n is a multiple of four. Our table ends at

n	4	8	12	16	20	24	28	32	36
$\sqrt[n]{\lvert S_n(1324)\rvert}$	2.19	3.348	4.141	4.728	5.186	5.559	5.869	6.134	6.363

Table 2: Values of $\sqrt[n]{|S_n(1324)|}$ for small n.

$n = 36$, because this is the largest value for which $|S_n(1324)|$ is currently known. All of these terms are available in the On-Line Encyclopedia of Integer Sequences, the terms with $n \leq 31$ are due to recent work of Fredrik Johansson and Brian Nakamura [18], and the rest are due to even more recent work of Andrew Conway and Anthony Guttmann [12]. It is known that $|S_n(1234)| \sim 9^n n^{-4}$ and $|S_n(1342)| \sim 8^n n^{-5/2}$, so we might ask for a constants μ and θ for which $|S_n(1324)|$ is asymptotic to a function of the form $n^\theta \mu^n$. In fact, Johansson and Nakamura, with a computational assist from Shalosh B. Ekhad, have found that if $|S_n(1324)|$ is asymptotic to a function of this form, then $\mu \approx 10.45$ and $\theta \approx -8.64$ are the most consistent with the values of $|S_n(1324)|$ for $n \leq 31$. If Johansson and Nakamura's values of μ and θ are in the right neighborhoods, then the fact that the value of $|\theta|$ for $|S_n(1324)|$ is so much greater than the corresponding values for the other two sequences explains the slow convergence we see in the data for $\sqrt[n]{|S_n(1324)|}$.

Unfortunately, Johansson and Nakamura's values might not be in the right neighborhoods. Even more troubling, as Johansson and Nakamura themselves point out, it's unlikely that there is an asymptotic formula for $|S_n(1324)|$ of the form $n^\theta \mu^n$ at all. Indeed, in May of 2014 Conway and Guttmann gave compelling evidence that $|S_n(1324)|$ actually behaves like a function of the form $B\mu^n \mu_1^{\sqrt{n}} n^g$ for constants B, μ, μ_1, and g. In addition, they estimated that $\mu = 11.60 \pm 0.01$, $\mu_1 = 0.0398 \pm 0.0010$, $g = -1.1 \pm 0.2$, and $B = 9.5 \pm 1.0$. Conway and Guttmann's methods strongly suggest they have the right asymptotic form for $|S_n(1324)|$, but these methods do not seem powerful enough to prove a claim like this. Which means that if we want to understand how fast $a_n(1324)$ grows, then we should probably try to bound it.

There is a natural lower bound on $L(1324)$: since $S_n(132) \subset S_n(1324)$, we must have $L(1324) \geq 4$. Beyond this, little progress was made on the problem of bounding $L(1324)$ below until 2006, when Michael Albert, Murray Elder, Andrew Rechnitzer, Paul Westcott, and Mike Zabrocki [2] used a finite automaton which accepts only sequences constructing certain elements of $S_n(1324)$ to show that $L(1324) \geq 9.47$. This result was the first nontrivial lower bound on $L(1324)$, but it also showed conclusively that 1234 is neither the most restrictive nor the least restrictive forbidden pattern of length four, a fact which is at odds with our intuition about which permutations are at the extremes of the set of all permutations. We continue to see incremental progress on this problem: in early June of 2014, David Bevan submitted a paper to arXiv [5] in which he uses interleaved trees and Łukasiewicz paths to construct a large class of permutations in $S_n(1324)$, thus showing that $L(1324) \geq 9.81$.

One of the nice features of Marcus and Tardos's proof of the Stanley-Wilf conjecture, in addition to its simplicity, is that it gives an explicit upper bound on $L(\pi)$ in terms of the length of π: if π has length k, then the proof tells us that $L(\pi) \leq 15^a$, where $a = 2k^4\binom{k^2}{k}$. Even better, in 2009 Josef Cibulka improved [10] this general bound significantly, to $L(\pi) \leq 2.88 \cdot 4k^8\binom{k^2}{k}^2$. Since $\binom{k^2}{k} = \frac{k^2(k^2-1)(k^2-2)\cdots(k^2-k+1)}{k(k-1)(k-2)\cdots 2\cdot 1}$ is at least a polynomial of degree k in k, the factor $\binom{k^2}{k}^2$ in Cibulka's bound is at least $2^{k \ln k}$. Ever since Stanley and Wilf's first conjectures about $L(\pi)$, people had hoped this bound could be improved to some polynomial in k. But Jacob Fox dashed these hopes in the fall of 2013, by showing [15] that $L(\pi)$ is exponential in k for almost all permutations π of length k. Nevertheless, Cibulka's work gives us a starting point in our quest to bound $L(1324)$ in particular: it says $L(1324) \leq 2.50078 \times 10^{12}$. However, because this bound is general, it's not at all sharp, so we turn our attention to finding tighter upper bounds on $L(1324)$ in particular.

To start to improve our upper bound on $L(1324)$, recall that we saw earlier how to reconstruct a permutation which avoids 132 from the positions and values of its left-to-right minima. This observation effectively divides the set of all permutations of a given length into classes, each of which contains no more than one element of $S_n(132)$. Since there are at most 4^{n-1} classes, there are at most 4^{n-1} elements of $S_n(132)$. In his thesis Bóna expanded on this idea, by proving that each of these classes contains at most 8^n elements of $S_n(1324)$. Thus, $L(1324) \leq 32$.

The numerical data regarding $\sqrt[n]{|S_n(1324)|}$ in Table 2 suggest that our upper bound on $L(1324)$ is worse than our lower bound, but Bóna's thesis work represented the state of the art for nearly a decade and a half. A new idea was required, and in 2012 Anders Claesson, Vít Jelínek, and Einar Steingrímsson provided one [11]. To each $\pi \in S_n(1324)$, we associate a sequence of colors, red or blue, one for each entry of π. To do this, we first color the leftmost entry π_1 red. Proceeding from left to right, if we have colored the entries $\pi_1, \ldots, \pi_{j-1}$, then we color the entry π_j according to the following rules.

1. If coloring π_j red would create a red subsequence of type 132, then we color it blue.
2. If one of $\pi_1, \ldots, \pi_{j-1}$ is blue, and is less than π_j, then we color π_j blue.
3. If neither of the first two rules applies, then we color π_j red.

For example, if $\pi = 749538612$ then we would color the entries as in 749538612, obtaining the color sequence $RRRBRBBBRR$. Note that we color the 5 blue because otherwise the 4, the 9, and the 5 would create a red subsequence of type 132. On the other hand, we color both the 8 and the 6 blue for two reasons: otherwise they would be part of a red subsequence of type 132, and they each have a smaller blue entry (namely, the 5) to their left.

We can show (go ahead! it's not hard) that if π avoids 1324, then its sequence of red entries avoids 132 and its sequence of blue entries avoids 213. Claesson, Jelínek, and Steingrímsson use some standard computations with binomial coefficients to show that this implies $L(1324) \leq 16$, but we can also follow Bóna [9] to really see where this bound comes from. Having colored each entry of $\pi \in S_n(1324)$ red or blue, we now separate the entries of each color into two classes. To do this, first replace each red entry which is a left-to-right minimum in π with the letter A, and then replace each remaining red entry with the letter B. Similarly, inspired by the fact that 1324 is its own reverse-complement and 132 and 213 are reverse-complements of each other, replace each blue entry which is a right-to-left maximum with the letter C, and replace each of the other blue entries with the letter D. For example, we've seen that if $\pi = 749538612$ then our color sequence is $RRRBRBBBRR$. Following our new prescription, we find that the associated string of As, Bs, Cs, and Ds is $AABDACCAB$.

We have now seen that for each $\pi \in S_n(1324)$, we obtain a string $w(\pi)$ of As, Bs, Cs, and Ds by listing which of these letters is assigned to the first entry, which to the second, etc. Similarly, we obtain a second such string $z(\pi)$ by listing which of these letters is assigned to the entry 1, which to the entry 2, etc. For instance, if $\pi = 749538612$ then we've seen that $w(\pi) = AABDACCAB$, and we also have $z(p) = ABAADCACB$. Note that our constructions of $w(\pi)$ and $z(\pi)$ amount to an elaboration on our earlier use of the positions and values of the left-to-right minima in a permutation $\pi \in S_n(132)$. Bóna's main result is that while some pairs of strings correspond to no permutation, no pair of strings corresponds to two or more permutations. Since there are 16^n pairs of strings of length n, we have $L(1324) \leq 16$. In fact, Bóna also shows that none of these strings contains a B which is immediately followed by a C, so some standard computations with linear recurrence relations gives us $L(1324) \leq 7 + 4\sqrt{3} \approx 13.928$, a result which Bóna has recently improved to $L(1324) \leq 13.73718$.

New Frontiers: Getting Closer to God

For more information on the state of the permutation pattern art, I would refer the reader to Bóna's book [8], Kitaev's monograph [19], and Steingrímsson's survey of open problems [35]. In the meantime, Table 3 summarizes the Stanley-Wilf limits we've discussed in this paper, and the reader

π	$L(\pi)$
132	4
1342	8
1324	[9.81, 13.73718]
$123\cdots k$	$(k-1)^2$

Table 3: The Stanley-Wilf limits in this paper.

can find more information about these limits in Section 6.1.4 of Kitaev's monograph. We still have much to discover about these quantities. In addition, there are several closely related topics that fall outside the scope of this article. For example, while it follows from the Marcus-Tardos theorem (as the Stanley-Wilf conjecture is now known) that for every set R of forbidden patterns there is a constant $c(R)$ such that $|S_n(R)| \le c(R)^n$, it is still not known whether $\lim_{n\to\infty} \sqrt[n]{|S_n(R)|}$ exists for every set R of forbidden patterns. Nevertheless, this limit does exist for many sets R, and Vatter has found some amazing structure in the set of real numbers which can appear as one of these limits [38].

Finally, Steingrímsson has recently stepped up to challenge Zeilberger's original claim about the difficulty of finding $|S_{1000}(1324)|$, saying, "I'm not sure how good Zeilberger's God is at math, but I believe that some humans will find this number in the not so distant future." In fact, in 2013 Steingrímsson and Zeilberger made a bet about whether someone will find $|S_{1000}(1324)|$ by 2030. If someone finds this value in year n for $n \le 2030$, then Zeilberger will pay Steingrímsson €10$(2030 - n)$. Otherwise, Steingrímsson will pay Zeilberger €170. Until Zeilberger and Steingrímsson settle their bet, we'll have to let Conway and Guttmann have the last word on the subject. They write, "While making no Messianic claims, our asymptotics permit the approximate answer 4.6×10^{1017}."

Acknowledgements

Several people gave me helpful and extensive comments, suggestions and corrections on various drafts of this paper, and brought me up to date on recent developments related to Stanley-Wilf limits. In particular, I would like to thank Miklós Bóna, Stephen Kennedy, Darla Kremer, Adam Marcus, Lara Pudwell, Kailee Rubin, Einar Steingrímsson, Gábor Tardos, and Doron Zeilberger for their thorough and thoughtful feedback on various drafts of this article. Their comments led to a much-improved final version. Of course, blame for any remaining errors or shortcomings is mine alone.

References

[1] F. Acerbi, On the shoulders of Hipparchus: a reappraisal of ancient Greek combinatorics, *Arch. Hist. Exact Sci.* **57** (2003) 465–502.

[2] M. H. Albert, M. Elder, A. Rechnitzer, P. Westcott,, M. Zabrocki, On the Wilf-Stanley limit of 4231-avoiding permutations and a conjecture of Arratia, *Adv. in Appl. Math.* **36** no. 2 (2006) 95–105.

[3] N. Alon, E. Friedgut, On the number of permutations avoiding a given pattern, *J. Combin. Theory, Series A* **89** (2000) 133–144.

[4] R. Arratia, On the Stanley-Wilf conjecture for the number of permutations avoiding a given pattern, *Electron. J. Combin.* **6** (1999).

[5] D. Bevan, Permutations avoiding 1324 and patterns in Łukasiewicz paths (2014), available at http://arxiv.org/abs/1406.2890.

[6] M. Bóna, *Exact and Asymptotic Enumeration of Permutations with Subsequence Conditions.* Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1997.

[7] M. Bóna, The solution of a conjecture of Stanley and Wilf for all layered patterns, *J. Combin. Theory Ser. A* **85** (1999) 96–104.

[8] M. Bóna, *Combinatorics of Permutations*, second edition. Chapman & Hall, 2012.

[9] M. Bóna, A new upper bound for 1324-avoiding permutations (2012), available at http://arxiv.org/abs/1207.2379.

[10] J. Cibulka, On constants in the Füredi-Hajnal and the Stanley-Wilf conjecture, *J. Combin. Theory Ser. A* **116** (2009) 290–302.

[11] A. Claesson, V. Jelínek,, E. Steingrímsson, Upper bounds for the Stanley-Wilf limit of 1324 and other layered patterns, *J. Combin. Theory Ser. A* **119** no. 8 (2012) 1680–1691.

[12] A. R. Conway, A. J. Guttmann, On the growth rate of 1324-avoiding permutations (2014), available at http://arxiv.org/abs/1405.6802.

[13] W. F. Doran IV, A connection between alternating sign matrices and totally symmetric self-complementary plane partitions, *J. Combin. Theory Ser. A* **64** (1993) 289–310.

[14] A. W. F. Edwards, *Pascal's Arithmetical Triangle.* Oxford University Press, Oxford, 1987.

[15] J. Fox, Stanley-Wilf limits are typically exponential (2013), available at http://arxiv.org/abs/1310.8378.

[16] Z. Füredi, P. Hajnal, Davenport-Schinzel theory of matrices, *Discrete Math.* **103** (1992) 233–251.

[17] L. Habsieger, M. Kazarian, S. Lando, On the second number of Plutarch, *Amer. Math. Monthly* **105** no. 5 (1998) 446.

[18] F. Johansson, B. Nakamura, Using functional equations to enumerate 1324-avoiding permutations (2013), available at http://arxiv.org/abs/1309.7117.

[19] S. Kitaev, *Patterns in Permutations and Words.* Springer-Verlag, Berlin, 2011.

[20] M. Klazar, The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture, in *Formal Power Series and Algebraic Combinatorics: 12th International Conference, FPSAC'00, Moscow, Russia, June 2000, Proceedings*, Edited by D. Krob, A. Mikhalev, Springer-Verlag, Berlin, 2000. 250–255.

[21] P. MacMahon, *Combinatory Analysis*, Cambridge University Press, 1915.

[22] A. Marcus, G. Tardos, Excluded permutation matrices and the Stanley-Wilf conjecture, *J. Combin. Theory Ser. A* **107** (2004) 153–160.

[23] The On-Line Encyclopedia of Integer Sequences, published electronically at oeis.org.

[24] K. H. Parshall, *James Joseph Sylvester: Jewish Mathematician in a Victorian World.* The Johns Hopkins University Press, Baltimore, 2006.

[25] Plutarch, *Moralia*, vol. IX (introduction by E. L. Minar, Jr.), Loeb Classical Library, Harvard University Press, Cambridge, MA, 1961.

[26] L. Pudwell, Enumeration schemes for words avoiding permutations, in *Permutation Patterns*, London Math. Soc. Lecture Note Ser., Vol. 376, Cambridge University Press, 2010. 193–211.

[27] A. Regev, Asymptotic values for degrees associated with strips of Young diagrams, *Adv. Math.* **41** (1981) 115–136.

[28] D. P. Robbins, The story of $1, 2, 7, 42, 429, 7436, \cdots$, *Math. Intelligencer* **13** no. 2 (1991) 12–19.

[29] D. Rotem, Stack sortable permutations, *Discrete Math.* **33** (1981) 185–196.

[30] C. Schensted, Longest increasing and decreasing subsequences, *Canad. J. Math.* **13** (1961) 179–191.

[31] R. Simion, F. W. Schmidt, Restricted permutations, *Europ. J. Combin.* **6** (1985) 383–406.

[32] Z. Stankova, Forbidden subsequences, *Discrete Math.* **132** no. 1–3 (1994) 291–316.

[33] Z. Stankova, Classification of forbidden subsequences of length 4, *European J. Combin.* **17** no. 5 (1996) 501–517.

[34] R. P. Stanley, Hipparchus, Plutarch, Schröder and Hough, *Amer. Math. Monthly* **104** no. 4 (1997) 344–350.

[35] E. Steingrímsson, Some open problems on permutation patterns, in *Surveys in Combinatorics*, Edited by S. R. Blackburn, S. Gerke, M. Wildon, London Math. Soc. Lecture Note Ser., Vol. 409, Cambridge University Press, London, 2013.

[36] J. J. Sylvester, F. Franklin, A constructive theory of partitions, arranged in three acts, an interact and an exodion, *American Journal of Mathematics* **5** no. 1 (1882) 251–330.

[37] V. Vatter, Enumeration schemes for restricted permutations, *Combin. Probab. Comput.* **17** no. 1 (2008) 137–159.

[38] V. Vatter, Permutation classes of every growth rate above 2.48188, *Mathematika* **56** no. 1 (2010) 182–192.

[39] J. West, *Permutations with Forbidden Subsequences and Stack-sortable Permutations*, Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1990.

[40] Wikipedia contributors, Subadditivity, *Wikipedia, The Free Encyclopedia*, available at http://wikipedia.org/wiki/Subadditivity.

[41] D. Zeilberger, A proof of Julian West's conjecture that the number of two-stack-sortable permutations of length n is $2(3n)!/((n+1)!(2n+1)!)$, *Discrete Math.* **102** no. 1 (1992) 85–93.

[42] D. Zeilberger, Enumeration schemes and, more importantly, their automatic generation, *Ann. Comb.* **2** no. 2 (1998) 185–195.

[43] D. Zeilberger, A loving rendition of the Marcus-Tardos amazing proof of the Füredi-Hajnal conjecture, available at http://www.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/martar.html.