ON CONWAY'S RECURSIVE SEQUENCE

TAL KUBO AND RAVI VAKIL

October 28, 1993

1. Introduction

The recursion

$$a(n) = a(a(n-1)) + a(n - a(n-1))$$

for $n > 2$ with initial values $a(1) = a(2) = 1$, has been studied by Douglas Hofstadter, John H. Conway and others at various times since about 1975 [7, 2]. Conway in particular discovered many of its properties, and helped to publicize it by offering a cash prize for information about its asymptotic growth.

Colin Mallows received this prize [9]; properties of the sequence obtained from this recursion with initial values $a(1) = \ldots = a(k) = 1$, have been obtained by David Newman and Daniel Kleitman [8].

It is the purpose of this note to present the structure of this sequence (and its generalizations) in a new way, which makes its properties quite transparent.

In particular we correspond to each integer k a finite set of natural numbers $S(k)$ according to a simple rule, and find that the function $a(k)$ corresponds to a certain "compression" operation, C, on the set $S(k)$:

$$S(a(k)) = CS(k)$$

$S(k)$ is just the k^{th} term in a certain ordering of the finite sets of natural numbers (sort by largest element, then by decreasing order of size, then lexicographically). Compression here means (roughly) to drop the smallest element and subtract one from the rest. (See section 6 for a precise definition.) Then the compression of the k^{th} set is the $a(k)^{\text{th}}$ set.

The asymptotic properties of $a(n)$ are easy consequences of this characterization.

In the next section we give a brief summary of the properties of the sequence $a(n)$; it so happens that numbers between 2^k and $2^{k+1} - 1$ are mapped by a into those between 2^{k-1} and $2^k - 1$. One can represent a by an arborescence, whose level to level structure can be described by reading off the rows of a recursively defined tree.

We then define the tree representation, and prove its validity. We use it to prove several properties of the a's including a non-recursive formula for them, and Conway's doubling property.

The function $2n - a(n)$ vanishes when n is a power of two, and between each pair of consecutive powers of two, the $(k-1)^{\mathrm{st}}$ and k^{th}, is positive and symmetric and achieves its maximum $2^{([k/2]-1)}$ times. We prove this and give an algorithm for solving Conway's asymptotic question in general.

Finally we consider generalizations and present a discussion of open questions. Our characterization of $a(n)$ carries over to Newman's higher-order variants, and reveals a completely analogous picture. The structures obtained suggest that the $a(n)$ recursion is not so isolated and may be related to more traditional branches of mathematics such as the theory of games.

2. Properties of $a(n)$

The following chart gives the first few values of $a(n)$:

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
$a(n)$	1	1	2	2	3	4	4	4	5	6	7	7	8	8	8	8	9
n	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34
$a(n)$	10	11	12	12	13	14	14	15	15	15	16	16	16	16	16	17	18

A glance at it reveals an unexpected degree of structure. Some of the properties that are visible here:

- (1) $a(n) \leq n$ (so that $a(n)$ is well-defined by the recurrence).
- (2) $a(n) - a(n-1) = 0$ or 1, for all $n \geq 1$.
- (3) $a(2^n) = 2^{n-1}$ for all $n \geq 1$.
- (4) $a(n) = 2^k$ for exactly $k+1$ values of n $(k > 0)$.
- (5) $a(n) \geq \frac{n}{2}$, with equality iff n is a power of 2 and $n \neq 1$.
- (6) $\frac{a(n)}{n} \to \frac{1}{2}$ as $n \to \infty$
- (7) $A(n)$, defined as $2a(n)-n$, has zeros at the powers of 2, and has symmetric positive "humps"[1] between its zeros, *i.e.* $A(2^n + t) = A(2^{n+1} - t) > 0$ for $0 < t < 2^n$.
- (8) The maximum of $A(n)$ on the "hump" between 2^{n-1} and 2^n is achieved $2^{[n/2]-1}$ times.
- (9) $a(2n) \leq 2a(n)$ for all n.

The first two are easy to prove by induction [9, 8]. The next three can also be proved inductively. Conway proved the asymptotic property (6).

The symmetry property (7), so striking and unusual, is a strong clue that $a(n)$ has some combinatorial structure. The prominent appearance of powers of two further suggests some description involving finite sets, where the symmetry property might be induced by complementation. We will confirm this intuition below.

[1] This term comes from Mallows [9]

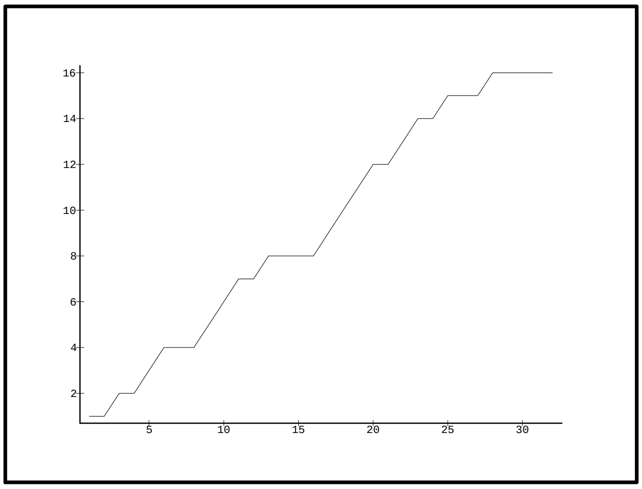

FIGURE 1. $a(n),\ 1 \le n \le 32$

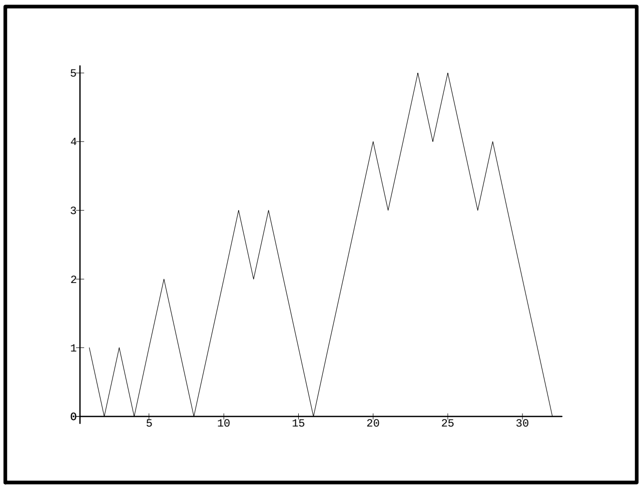

FIGURE 2. $A(n) = 2a(n) - n,\ 1 \le n \le 32$

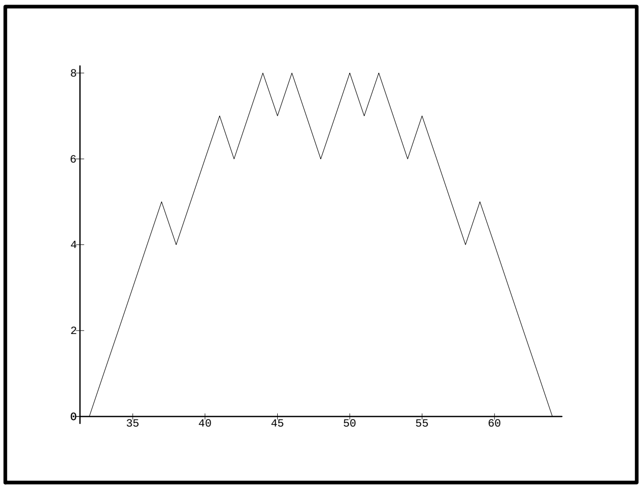

FIGURE 3. $A(n),\ 32 \le n \le 64$

3. Frequency counts and the Tree Representation of $a(n)$

Since $a(n)$ grows monotonically by at most one as n increases by one, it is determined by its frequency count, *i.e.* by the number of times it takes each value. The additional property that $a(n) < n$ for $n > 1$ allows us to represent the function $a(n)$ as a tree. We may display $a(n)$ using a directed graph, with vertices {1,2,3,...} and an arrow $i \to j$ whenever $i > j$ and $a(i) = j$. Then a frequency count is associated to each row, by counting the number of incoming arrows at the successive vertices of that row. We obtain a collection of finite sequences $(Q_i)_{i\ge 0}$, each giving the frequency count of one layer of this graph. (See Figure 4.)

Notice that the layers of the graph extend between consecutive powers of 2, which follows from $a(2^n) = 2^{n-1}$. We define the n^{th} *level* to be the set of integers on the $(n+2)^{\text{nd}}$ row of our graph. Thus for $n \ge 0$, level n consists of the numbers from 2^n+1 to 2^{n+1}, a maps level n to level $n-1$, and Q_n is the frequency count of level $n-1$.

The graph is governed by a peculiar growth pattern: If a vertex V, not at the end of a row, has indegree N, then the vertices pointing to V, have indegrees $1,2,...,N$, in that order. If V, at the end of a row, has indegree N, then the vertices pointing to it have indegrees $1,2,..,N-1,N+1$, in that order.

If this pattern persists, then for all n, Q_{n+1} is generated from Q_n by replacing each occurence of 2 by 12, each 3 by 123, and so on, and increasing the last term by 1. For example:

Q_3 =1124 expands to 11121234 $\Rightarrow$ Q_4=11121235

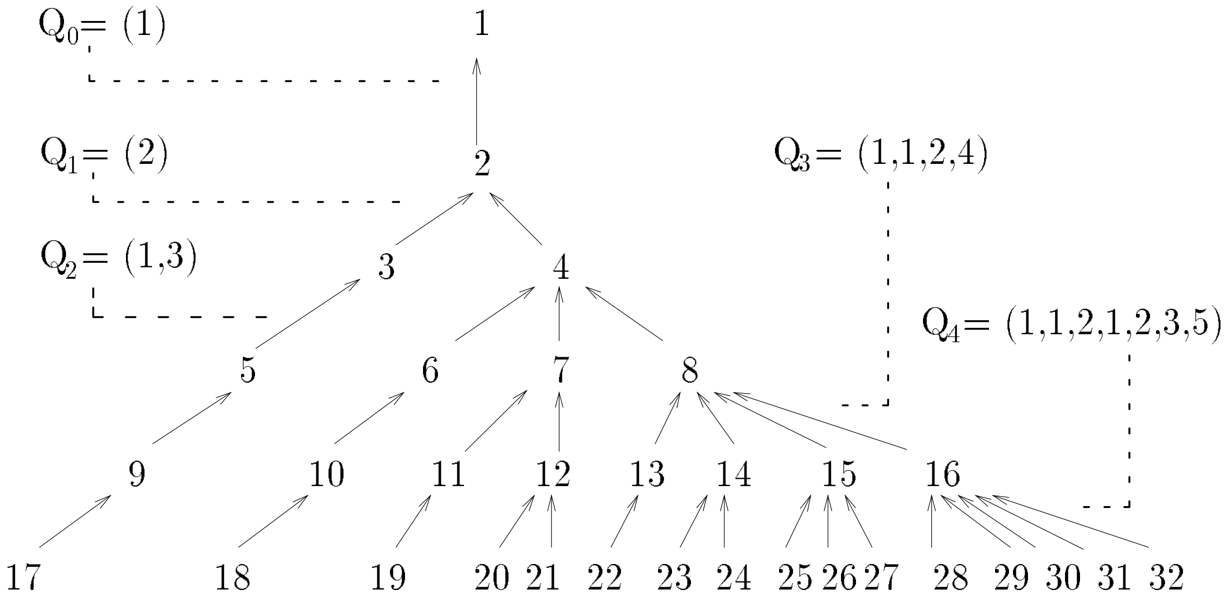

FIGURE 4. Tree Representation of $a(n)$

A proof that this rule does indeed give the frequency counts would provide a more structured method of producing $a(n)$, one layer at a time. A look at the definition of $a(n)$ shows that to prove anything non-trivial about $a(n)$, we must also understand its "sister function" $b(n) = n - a(n-1)$. (There is another function closely related to $a(n)$, $c(n) = n - a(n)$, which satisfies the "dual" relation $c(n) = c(a(n-1)) + c(n - a(n-1))$. This function, however, is just $b(n)$ in disguise: it is its conjugate with respect to the change of coordinates $n \to n-1$.)

We may display $b(n)$ by a graph and compute its row frequency counts ${Q_n}^*$, as we did with $a(n)$ (see Figure 5). Here we find a surprise: ${Q_n}^*$ is Q_n backwards, that is, the $b(n)$ tree is the mirror image of the $a(n)$ tree. This is a clue that we are about to uncover some concealed structure. Our approach will be to *define* Q_n by the recursive expansion rule, to study the sequences obtained (which are interesting in their own right), then to prove later that they are the frequency counts of $a(n)$.

4. Reduced frequency counts

More symmetries are revealed through the "reduced" frequency counts R_n, formed from Q_n by reducing its last element by 1 (for $n > 0$). The first few R_n are:

$$R_1 = 1, R_2 = 12, R_3 = 1123, R_4 = 11121234, R_5 = 1111211212312345$$

The recursive rule to generate the R_n directly is slightly simpler than for Q_n:

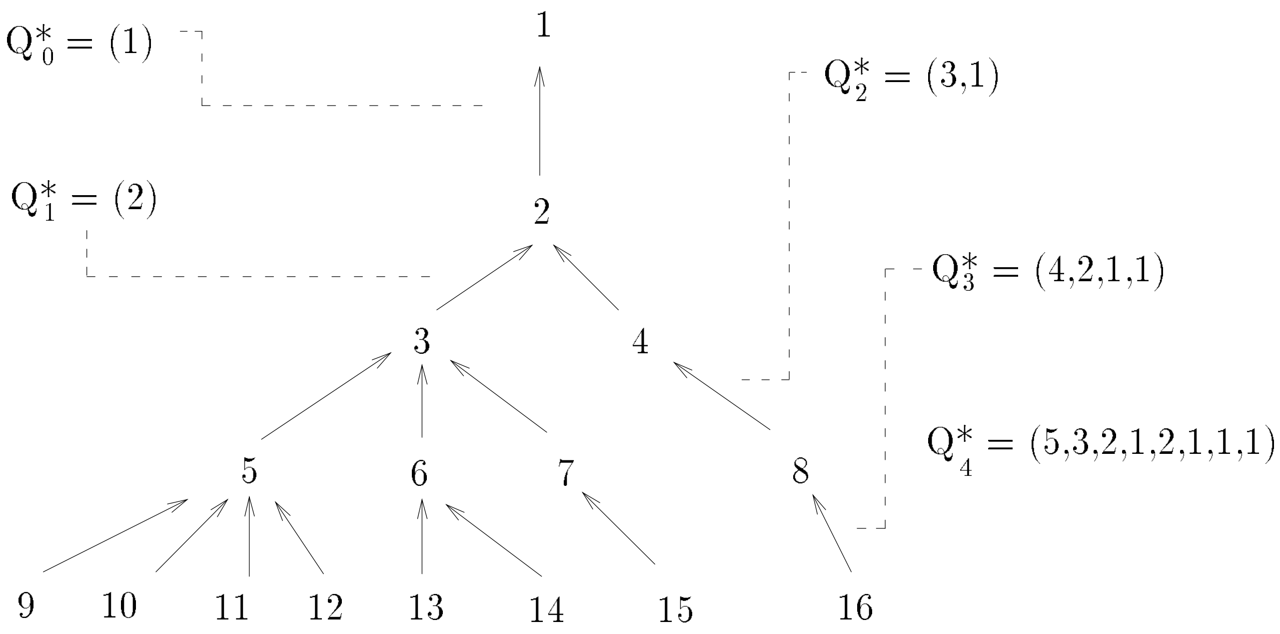

FIGURE 5. Tree Representation of $b(n)$

$R_1 = 1$, $\quad R_n = \mathbf{e}(R_{n-1})n$ (the concatenation of $\mathbf{e}(R_{n-1})$ and n), where $\mathbf{e}$ is the expansion operator which replaces each t by $12...t$, for $t = 1, 2, ...$

Some interesting properties of the R_n:

- (a) R_n has length 2^{n-1}
- (b) The sum of all terms of R_n is $2^n - 1$
- (c) Reducing each element of R_n by 1 and suppressing zeros, we get R_{n-1}
- (d) The lengths of the successive blocks of 1's in R_n are Q_{n-2} backwards. For example,
  $$\begin{array}{lcccccccccc} R_5 = & 1111 & 2 & 11 & 2 & 1 & 23 & 1 & 2345 & \\ \text{block lengths} & 4 & & 2 & & 1 & & 1 & & = Q_3, \text{ reversed} \end{array}$$
- (e) *The "Flip-Flop" Property* Writing R_n backwards underneath itself, there is exactly one 1 in each column (for $n > 1$).
  For example, $\begin{array}{rcl} R_5 & = & 1111211212312345 \\ R_5 \text{ reversed} & = & 5432132121121111 \end{array}$
- (f) Write R_n in "unary form", replacing each number by sums of ones and writing commas between terms. Exchange + signs and commas. Then the result is the unary form of R_n backwards! For example,

$$\begin{aligned} R_3 &= 1,1,2,3 \\ &= 1,1,1+1,1+1+1 \\ &\rightarrow 1+1+1,\ 1+1,1,1 \\ &= R_3 \text{ backwards} \end{aligned}$$

Properties (a) and (b) are both equivalent to the equation $a(2^n) = 2^{n-1}$. Properties (e) and (f) are reformulations of the symmetry of $2a(n) - n$ mentioned in section 2. To see equivalence with (f), consider $A(n) = 2a(n) - n$ for $2^m < n \leq 2^{m+1}$: its successive differences are $\pm$ 1, and can be generated by changing commas to + signs and + signs to – signs in the unary expansion of R_m. This follows directly from the assumption that the Q_i do give the frequency counts of $a(n)$.

We've saved the most powerful symmetry for last! Consider a Pascal's triangle where the rule for forming a term from two adjacent ones in the previous row is concatenation, and the left "leg" of the triangle consists of all 1's as usual, but the right "leg" consists of the positive integers in order.

$$\begin{array}{ccccccccc} & & & & 1 & & & & \\ & & & 1 & & 2 & & & \\ & & 1 & & 12 & & 3 & & \\ & 1 & & 112 & & 123 & & 4 & \\ 1 & & 1112 & & 112123 & & 1234 & & 5 \end{array}$$

Notice that the $(n-1)^{\text{st}}$ row of this triangle (counting the top row as zero) is just R_n broken into pieces $R(n,k)$ of size $\binom{n-1}{k-1}$ $(1 \leq k \leq n)$. An easy inductive proof that $R(n+1,k) = \mathbf{e}(R(n,k))$ shows that this triangle does in fact yield R_n (base case $R(n+1,1) = 1 = \mathbf{e}(R(n,1))$ and $R(n+1,n) = 1,2,...,n = \mathbf{e}(n) = \mathbf{e}(R(n,n))$; inductive step (for $n > k > 0$) $\mathbf{e}(R(n,k)) = \mathbf{e}(R(n-1,k-1)R(n-1,k)) = R(n,k-1)R(n,k) = R(n+1,k)$).

We can use the Pascal's triangle representation of R_n to sketch quick proofs of properties (a) through (e).

(a) The n^{th} row of binomial coefficients has sum 2^n.

(b) Each term in R_n has 2 "daughters" in the next row of the Pascal's triangle, except for the ending $n+1$ which has daughters $n+1$ and $n+2$ in the next row. Thus (sum of the terms in R_{n+1}) = 2 (sum of the terms in R_n) + 1.

(c) By induction show that the reduction of $R(n,k)$ is $R(n-1,k-1)$.

(d) Ignore the left- and right-most elements of each row (*i.e.* $R(n,1)$ and $R(n,n)$). For $1 < k < n$, let $g(n,k)$ be the sequence of lengths of successive blocks of 1's in $R(n,k)$ (*e.g.* $g(5,3) = 21$ as there are two blocks of 1's in 112123 of length 2 and 1 respectively). Show that $g(n,k)$ also satisfies the concatenation property $g(n+1,k) = g(n,k-1)g(n,k)$ $(k > 2)$ and $g(n,2) = n$. Then $g(n,k)$ is $R(n-2,n-k)$ backwards, by induction (or by glancing at Figure 6 – the diagrams on the left and right are mirror images). Taking into account the additional 1 in $R(n,1)$, we get the

$R(n,k)$					$g(n,k)$				
		1					—		
	1		2			—		—	
1		12		3	—		1		—
1	112		123	4	—	2		1	—
1	1112	112123	1234	5	—	3	21	1	—

FIGURE 6

		+		
	+		–	
+		+ –		–
+	++ –		+ – –	–
+	+++ –	++ – + – –	+ – – –	–

FIGURE 7

desired result.

(e) Replace 1's by + signs and non-1's by – signs, as in Figure 7.

Ignore the top row, and reflect the triangle about its vertical axis. The + and – signs on the borders of the rows have switched signs, and as all rows of the triangle are formed inductively from these signs, the tree is invariant under flipping left to right and exchanging + and – signs. This is equivalent to the flip-flop property (e) of R_n. [2]

The corresponding symmetry of $2a(n) - n$ follows since the signs in the triangle reproduce its successive differences.

It is possible to show (by an unpleasant induction) that Q_n does indeed give the frequency counts of $a(n)$. This fact and the power of the Pascal's Triangle representation of R_n to explain mysterious properties suggests that the properties discussed before fully capture the underlying strucutre. Such is not the case - much more remains to be revealed. Perhaps the moral of this story is that one shouldn't give up too soon.

5. Trees

The analogy between our triangle and the ordinary Pascal's triangle enumerating subsets of finite sets, again suggests an interpretation of $a(n)$ by sets. We proceed to derive one.

We define the "ordered trees" T_n inductively. Place n at the top of the tree, and have T_1, ..., T_{n-1} as subtrees emanating (in order) from the top (see Figure 8).

Then R_n can be obtained by reading T_n row by row, starting from the bottom and going left to right! Specifically, the rows of T_n, read from the bottom up, are $R(n,1)$,

[2] An equivalent Pascal's triangle was also introduced in Mallows' letter [10].

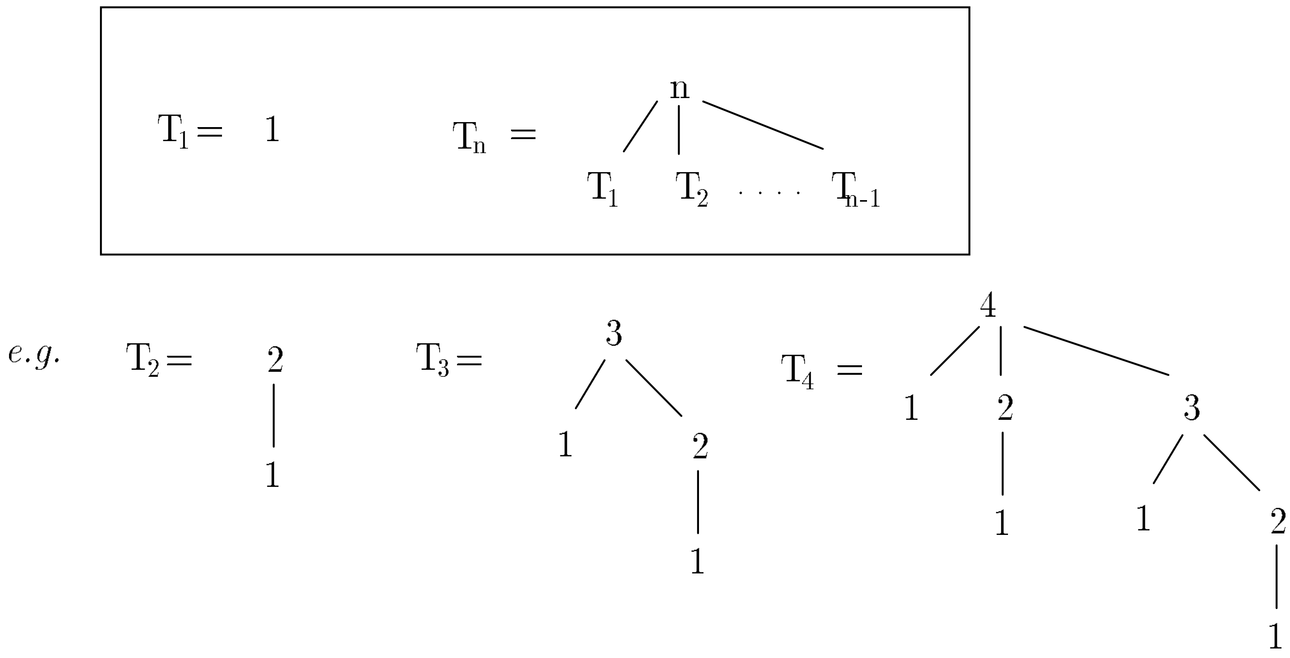

FIGURE 8. The trees T_n

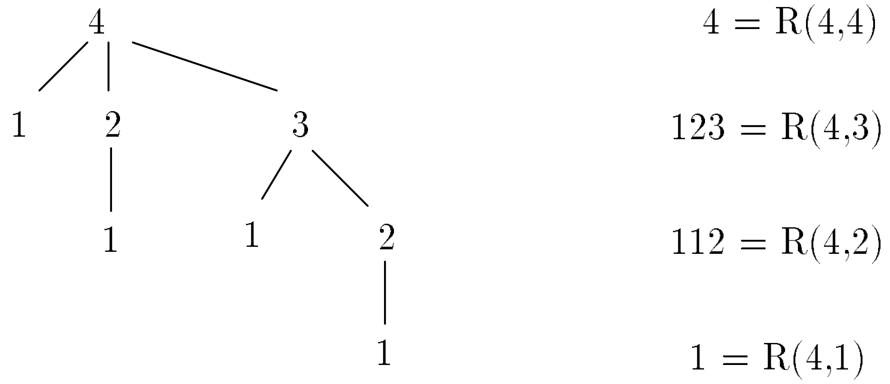

FIGURE 9. Reading $R(n,k)$ from T_n, when $n = 4$

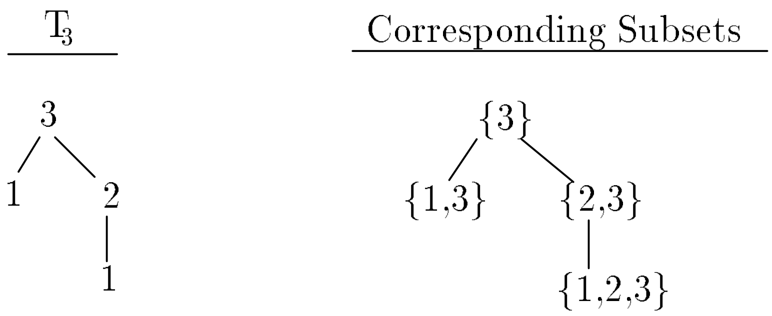

FIGURE 10

$R(n,2)$,... $R(n,n)$. More formally, define $T(n,k)$ to be the result of reading left to right the k^{th} row of T_n, where row 1 is the top and row n is the bottom. We claim that $R(n,k) = T(n, n+1-k)$ (see Figure 9).

Why does this rule work? From the definition of the triangle, $R(n,k) = R(n-1,k-1)R(n-1,k)$. On the other hand, the second row of T_n may be split into $\{1,2,..,n-2\}$ and $\{n-1\}$, which are respectively equal to $T(n-1,2)$ and $T(n-1,1)$. Therefore, apart from its top row, T_n is made from two copies of T_{n-1}, which are staggered by one row. It follows that for $j \neq 1, T(n,j) = T(n-1,j)T(n-1,j-1)$. Setting $j = n+1-k$, comparing the equations for R and T, and checking the anomalous case $k = n$, we get an inductive proof that $R(n,k) = T(n,n+1-k)$.

The graphs reveal that the vertices of T_n (and hence the terms of R_n) correspond naturally to the subsets of $1,2,...,n$ containing n. The elements of the set associated to a given vertex are the numbers along the (unique) path joining it to the root. The corresponding term of R_n can be recovered as the smallest element of this subset. (See Figure 10, for example.)

The bottom-up, left to right order in which we read T_n induces an ordering of the subsets of $1, .., n$ containing n:

Larger sets precede smaller ones. Sets of the same size are compared starting with the largest elements, then the next largest, and so on until there is a mismatch; this is the same as writing out the elements of the each set in descending order, then comparing the lists obtained, in lexicographic order.

To generate R_n, we can list these sets in order, then write under each set its smallest element. For example, when $n = 4$,

$$\begin{array}{lcccccccc} T_4 \rightarrow & \{1,2,3,4\} & \{1,2,4\} & \{1,3,4\} & \{2,3,4\} & \{1,4\} & \{2,4\} & \{3,4\} & \{4\} \\ & & & & \downarrow & & & & \\ R_4 \leftarrow & 1 & 1 & 1 & 2 & 1 & 2 & 3 & 4 \end{array}$$

To illustrate all this, let's prove the following theorem once again using sets.

Theorem 1. *The flip-flop property (e) is true.*

PROOF: Introduce an operation on finite sets of positive integers, called **flip**:

Given a set X, its flip $F(X)$ is formed by keeping the largest element of X, and complementing the rest.

Explicitly, if $X = \{a, b, c, ..., n\}$ then $F(X) = \{1, 2, .., n\} - (X - \{n\})$.

It is not too difficult to see that F is bijective and order reversing on the sets with fixed largest element. Write down the subsets of with largest element N, in order, with the smallest element of each set listed under it; this just gives R_n. Applying the flip to the sets of this list, we get the same list in reverse order. The flip-flop property (e) reduces to the fact that among a set and its flip, exactly one has smallest element 1. (This proof fails when $n = 1$, which is precisely when the flip-flop property fails.) $\square$

6. THE SET MODEL

We began by studying layers of length 2^n of values of $a(n)$, then focussed on segments of length $\binom{n}{k}$, and finally zoomed in on the individual elements of the segment corresponding to vertices of the tree T_n. We next assemble the information implied by these facts into a complete picture of the function, by looking at different layers at once.

First, extend the T_n subset order to an order on all finite sets of positive integers:

1. **(1) The empty set is first.**
2. **(2) Sets are sorted by largest element.**
3. **(3) Sets with largest element n are sorted according to the T_n ordering (sort in decreasing order of size, then compare equal-size sets by writing them in descending order and comparing lexicographically).**

$$\underbrace{\{\}}_{T_0} < \underbrace{\{1\}}_{T_1} < \underbrace{\{1,2\} < \{2\}}_{T_2} < \underbrace{\{1,2,3\} < \{1,3\} < \{2,3\} < \{3\}}_{T_3} < \ldots$$

FIGURE 11. The $S(n)$ Ordering of Finite Sets of Positive Integers

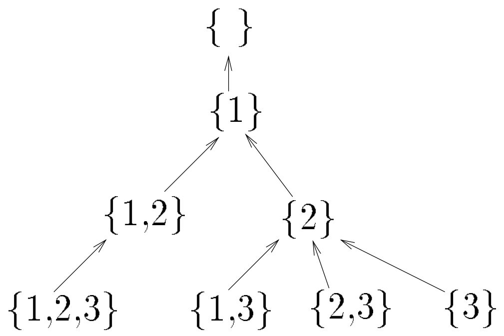

FIGURE 12. Graphic Representation of $a(n)$ in Terms of Subset Ordering (compare to Figure 4)

(See Figure 11.)

Let $S(n)$ denote the n^{th} set in the enumeration. We again represent $a(n)$ by a graph, but with the label of each vertex n replaced by $S(n)$. The vertices on level n will then be labelled by the subsets associated to T_{n+1} (the ones with largest element $n+1$). (See Figure 12, and compare it to Figure 4.)

Define the operation C, **compression**, as follows:

$$\begin{aligned} C(\varnothing) = C(\{1\}) &= \varnothing \\ C(\{n\}) &= \{n-1\} \text{ for } n > 1 \\ C(\{a_1, \ldots, a_k, n\}) &= \{a_2 - 1, .., a_k - 1, n - 1\} \end{aligned}$$

In other words, given a set $X \neq \emptyset$, $\{1\}$, form $C(X)$ by removing the smallest element of X if it has 2 or more elements, then subtracting 1 from the remaining elements.

Theorem 2. *$S(a(n))$ is the compression of $S(n)$.*

In words: if the n^{th} set is compressed, the result is the $a(n)^{\text{th}}$ set.

We will prove the theorem in the next section. Before doing so, let us check that it agrees with our earlier observations.

It is simple to prove that compression is order-preserving and surjective; this corresponds to the fact that the successive differences of $a(n)$ are all 0 or 1. That the 2^n'th set in our ordering is $\{n\}$, and compresses to $\{n-1\}$, is equivalent to $a(2^n) = 2^{n-1}$.

Next, let us see how to extract frequency counts from the set model. If S has more than one element, and $min(S) = a$, then the sets which compress to S are, in order, $\{1\} \cup (S+1), \{2\} \cup (S+1), ..., \{a\} \cup (S+1)$, where $S+1$ denotes S with all elements increased by 1. If S has one element, say $S = \{n\}$, then the sets which compress to S are $\{1, n+1\}, \{2, n+1\}, ...\{n, n+1\}, \{n+1\}$. Therefore, for a set of size at least 2, the frequency count is just its smallest element; for a singleton, the frequency count is one more than its unique element.

The set model shows that sequences Q_{N+1} are the frequency counts of the sets of level N. Here is a proof, by induction on $N \geq -1$. This is clearly the case for the top of the graph (level -1). To obtain the level N count from level $N-1$, read the sets for level $N-1$, but at each set S read, instead of the frequency count of S, the counts for the sets which compress to S. By the above reasoning, this means replacing each t in the sequence by $1, 2, 3, ..., t$, except at the end of the level where we replace $N+1$ by $1, 2, .., N, N+2$. Since this is exactly the rule that describes the formation of Q_{N+1} from Q_N, the induction is complete.

The set model therefore reproduces the earlier descriptions of $a(n)$, clarifying the relationship between Figure 4 and Figure 12. We will soon see that the symmetry of $A(n) = 2a(n) - n$ has an explanation in terms of the flip map, and that Conway's original problem on the growth of $a(n)$ can also be solved in this framework. But first we must show that the set model works, that it gives the same sequence $a(n)$ defined by the recurrence.

7. Proof that the set model claims are correct

Let $x(n)$ be the (well-defined) function $\mathbb{Z}^+ \to \mathbb{Z}^+$ defined by $S(x(n)) = C(S(n))$. We will prove that $x(n)$ coincides with $a(n)$, *i.e.* that the set model and the recursion give the same sequence.

To carry out the proof, we also need to understand $b(n) = n - a(n-1)$ in terms of finite sets. Recall that the graph for $b(n)$ was observed to give the same frequency counts as $a(n)$, but backwards. If true in general, this implies that $b(n) = FaF^{-1} = FaF$, *i.e.* $b(n)$ is the conjugate of $a(n)$ with respect to the flip. (Here we use F interchangeably as the "flip" function on sets defined in section 5, and the corresponding function on numbers exchanging $2^n + t$ and $2^{n+1} + 1 - t$ $(1 \leq t \leq 2^n)$.) The equivalent operation on sets is $T = FCF$, which can be defined directly by just unwinding the definitions of flip and compression.

Definition: the **contraction** $T(X)$ of a finite set X is formed by subtracting 1 from all elements, if X doesn't contain 1; removing the largest element, if $X = \{1, 2, ..., k\}$ for some k; subtracting 1 from all elements larger than t, where t is the smallest positive integer not in X, otherwise.

Explicitly:

$$\begin{aligned} T(\varnothing) = T(\{1\}) &= \varnothing \\ T(\{a,b,c,...\}) &= \{a-1, b-1, ...\} \text{ if } a > 1 \\ T(\{1,2,3,...,n\}) &= \{1,2,...,n-1\} \\ T(\{1,2,3,...,m-1,a,b,c,...\}) &= \{1,2,3,...,m-1,a-1,b-1,c-1,...\} \text{ if } a > m. \end{aligned}$$

Define the associated function $y(n)$ by $S(y(n)) = T(S(n))$. In the course of the proof we will show that $y(n)$ coincides with $b(n)$.

Now we are ready to give the proof. The idea is to show that the $x(n)$ and $a(n)$ have the same initial value and the same successive differences. We know that the successive differences are all 0 or 1, so we will actually have shown that the sequences increase at the same points. Having defined $x(n)$ and $y(n)$ correctly, the proof is basically a straightforward, if tedious, exercise. Lemma 4 (and implicitly 1,2, and 3) will be used later, so the reader who skips the proof that the model works should at least look at the lemmas.

We will use some simple notation: the sets of *level* N are those with largest element $N+1$, and the empty set has level -1 (this numbering agrees with our earlier use of the word "level"); $X+1 \quad (X-1)$ denotes the set formed by increasing (decreasing) all elements of a set X by 1; *order-preserving* is meant in the weak sense that if one set equals or precedes another, then this relationship is preserved.

Lemma 1. *C and T both map sets of level N to sets of level $N-1$, for each $N \geq 0$.*

PROOF: Clear from the definitions. □

Lemma 2. *C is surjective and order-preserving.*

PROOF: $C(\{1\} \cup (X+1)) = X$ shows surjectivity. Differences in level or size of sets (and hence relative order) are preserved by C, except that it maps the sets of size 1 and 2 on a given level to the same set on the next lower level (which still preserves order.) That C preserves order among sets of the same level and size is not hard to see from the definition if we recall that these sets are ordered by comparing largest elements first. □

Corollary 1. *$x(1) = 1$; for $n > 1, x(n) - x(n-1) = 0$ or 1.*

Lemma 3. *T is surjective and order-preserving.*

PROOF: $T(X+1) = X$ proves surjectivity. T preserves differences in level or size of sets. Among sets of the same size k and level n: the claim is trivial for $k = n$ and $k = 1$, so assume $1 < k < n$. For these sets it is clear from the definition of our ordering on sets, again because these sets are ordered by comparing largest elements first. □

Corollary 2. $y(1) = 1$*; for* $n > 1, y(n) - y(n-1) = 0$ *or* 1

Lemma 4. $x(n) > x(n-1)$ *iff* $n > 2$ *and* $S(n)$ *contains 1.*

PROOF: $x(1) = x(2) = 1$, by direct calculation, so we may assume $n > 2$. $x(n) > x(n-1)$ iff $S(n)$ is the first set compressing to $S(x(n)) = C(S(n))$, which is equivalent to $1 \in S(n)$. □

Note that this lemma justifies the remark in section 4, that the successive differences of $a(n)$ and $A(n)$ on level N can be read off from R_{N+1} by observing the pattern of 1's and non-1's.

Definition: The *gaps* of a set X are the stretches of consecutive integers in its complement, which are less than the largest element of X. For example, the effect of contraction on a set having gaps is to shorten its first gap by one.

Lemma 5. $x(n) = x(n+1)$ *iff* $S(n)$ *has a gap following its smallest element, for* $n > 2$.

PROOF: Clear, since $x(n) = x(n+1)$ iff $S(n)$ is not the last set compressing to $C(S(n))$. □

Lemma 6. $y(n) = y(n+1)$ *iff* $S(n)$ *contains 1, for* $n > 1$.

PROOF: First observe that the highest set (in our ordering) contracting to a non-empty set J is $J+1$. (Clearly $T(J+1) = J$. If $T(K) = J$ then $|K| \geq |J|$. If $|K| > |J|$ then $K < J+1$ by the definition of our ordering. Otherwise, if $K = \{k_1, ..., k_m\}$, $J = \{j_1, ..., j_m\}$ then $k_i \leq j_i + 1$ (by the definition of T), so $K < J+1$ again by the definition of our ordering.) Then

$$\begin{aligned} 1 \in S(n) &\Rightarrow S(n) < 1 + T(S(n)) \text{ but } T(S(n)) = T(1 + T(S(n))) \\ &\Rightarrow y(n) = y(n+1) \end{aligned}$$

$$\begin{aligned} 1 \notin S(n) &\Rightarrow S(n) \text{ is the last set contracting to } T(S(n)) = S(n) - 1 \\ &\Rightarrow y(n) < y(n+1) \end{aligned}$$

□

Proposition 1. $y(n) = n - x(n-1)$ *for* $n > 1$.

PROOF: Induction on $n > 1$. $y(2) = 2 - x(1) = 1$ by direct computation. Since both $y(n)$ and $n - x(n-1)$ grow in steps of 0 or 1, we need only show that for $n > 2$ they jump at the same places. So, we must show that $y(n+1) = y(n)$ iff $x(n) > x(n-1)$. But for $n > 2$, these are equivalent by Lemmas 4 and 6. □

Proposition 2. *$x(n) = a(n)$ for all $n \geq 1$.*

Proof: $x(1) = 1, x(2) = 1, x(3) = 2$ are easy to check by unwinding the definitions. This proves the proposition for $n = 1, 2, 3$. We must prove that for $n > 2$, $x(n)$ satisfies the defining recurrence of $a(n)$:

$$x(n) = x(x(n-1)) + x(n - x(n-1)).$$

The proof is by induction on $n > 2$. We have already checked the base case $n = 3$. Now suppose that $x(n) = x(x(n-1)) + x(n - x(n-1))$. Then we wish to show that $x(n+1) = x(x(n)) + x(n+1-x(n))$. Since $x(n+1) - x(n)$ is either 0 or 1, we will do so by analyzing when $x(n) = x(n+1)$.

Case 1: $x(n) = x(n-1)$. We must show

$$\begin{aligned} x(n+1) - x(n) &= x(x(n)) + x(n+1-x(n)) - x(x(n-1)) - x(n-x(n-1)) \\ &= x(x(n-1)) + x(n+1-x(n-1)) - x(x(n-1)) - x(n-x(n-1)) \\ &= x(n-x(n-1)+1) - x(n-x(n-1)) \\ &= x(y(n)+1) - x(y(n)) \end{aligned}$$

As both sides are either 0 or 1, it suffices to prove:

$$x(n+1) = x(n) \Leftrightarrow x(y(n)+1) = x(y(n))$$

By lemma 4, $S(n)$ doesn't contain 1, so $S(n)$ and $S(y(n)) = T(S(n))$ differ only in translation by 1 (*i.e.* $T(S(n)) = S(n) - 1$). Therefore n satisfies the condition of lemma 5 iff $y(n)$ does, proving the equivalence.

Case 2: $x(n) = x(n-1) + 1$.
Through similar manipulations we see that we must prove:

$$x(n+1) = x(n) \quad \Leftrightarrow \quad x(x(n)) = x(x(n)-1)$$

Since (by lemma 4) $S(n)$ contains 1:

$$\begin{aligned} x(n+1) = x(n) &\Leftrightarrow 2 \notin S(n) \text{ (lemma 5)} \\ &\Leftrightarrow 1 \notin C(S(n)) = S(x(n)) \\ &\Leftrightarrow x(x(n)) = x(x(n)-1) \text{ (lemma 4).} \end{aligned}$$

□

By Propositions 1 and 2, we have:

Proposition 3. *$y(n) = b(n)$ for all n. (by Propositions 1 and 2)*

We conclude that $a(n)$ and $b(n)$ are indeed described by the operations of compression and contraction on finite sets.

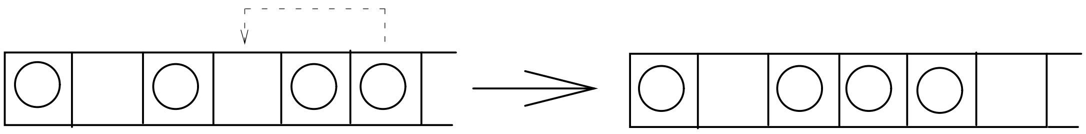

FIGURE 13. "Squashing"

8. Applications

Having proved the set model, we proceed to reap the dividends. We recall that a set is of **level** n if its largest element is $n+1$ (the corresponding numbers are in the range 2^n+1 to 2^{n+1}); the **segment** (n,k) refers to the $\binom{n}{k}$ sets on level n which have k elements, not counting their largest elements (which are all $n+1$). Thus the sets of segment (n,k) correspond to the sequence $R(n,n-k)$, to the $(k+1)^{\text{st}}$ row of T_{n+1}, and to the numbers between $2^n+1+\sum_{t=k+1}^{n}\binom{n}{t}$ and $2^n+\sum_{t=k}^{n}\binom{n}{t}$.

8.1. Compression Revisited. There is a second, more picturesque interpretation of $a(n)$ by sets. To state it we will need another ordering: sort the finite sets of positive integer, first by largest element, then in ascending order of size, then lexicographically. The first few sets in this order are

$$\varnothing < \{1\} < \{2\} < \{1,2\} < \{3\} < \{1,3\} < \{2,3\} < \{1,2,3\} < \ldots$$

Now, consider a semi-infinite one-dimensional board with cells marked {1,2,3,...}. Each set corresponds to an arrangement of counters on this board in an obvious way, by placing a counter in cell k for each element k of the set.

Define the operation of *squashing* as follows: for a finite set represented by counters on the board, pick up the first counter on the right (*i.e.* the one on the highest cell) and move it to the first empty square to the left (if any). If no square to the left is vacant, the counter is discarded. (See, for example, Figure 13.) Or equivalently, push the rightmost counter one step to the left, possibly forcing other counters to move left, and discard any counter that falls off the board.

Then: the $a(n)^{\text{th}}$ set is formed by squashing the n^{th} set!

The proof is to note that the n^{th} set in this enumeration corresponds to the n^{th} in our previous enumeration as follows: given a set X with largest element M, apply the flip map to X, then replace each $x \in X, x < M$ by $(M-x)$. This (involutive) correspondence also interchanges compression and squashing, which proves the equivalence of the two operations.

8.2. The symmetry property of $2a(n)-n$. The symmetry of $A(n) = 2a(n)-n$ can be explained in terms of a more fundamental symmetry between $a(n)$ and $b(n)$. Recall that the flip operation is an order-reversing involution of the sets of each level, given by taking the complement of all but the largest element. As with any operation on sets, there is an associated function F on positive integers: if $2^n < x \leq 2^{n+1}$, $F(x) = 2^{n+1} + 1 - x$.

Proposition 4. *$a(n) - b(n)$ is invariant under flip.*

PROOF: Recall that b is conjugate to a with respect to flip: $b = FaF$, and $F^2 = Id$. Hence $(a - b)F = aF - bF = Fb - Fa$. Since flip just reverses each level, for any elements a, b of the same level $Fb - Fa = a - b$. $\square$

Notice that $a(n) - b(n) = a(n) + a(n-1) - n$, which differs from $A(n)$ by at most 1, so this is almost the symmetry observed. Proving the equivalence of the two is simple:

Proposition 5. *$A(n)$ has symmetric humps: $A(2^n + j) = A(2^{n+1} - j)$ for $n \geq 1, 1 \leq j \leq 2^n$*

PROOF: Writing out $A(n)$ in terms of $a(n)$, we wish to prove:

$$2a(2^n + j) - (2^n + j) = 2a(2^{n+1} - j) - (2^{n+1} - j)$$

The symmetry of $a(n) - b(n)$ can be written as:

$$a(2^n + j) + a(2^n + j - 1) - (2^n + j) = a(2^{n+1} + 1 - j) + a(2^{n+1} - j) - (2^{n+1} + 1 - j)$$

Subtracting,

$$a(2^n + j) - a(2^n + j - 1) + a(2^{n+1} + 1 - j) - a(2^{n+1} - j) = 1$$

$$k(p) + k(F(p)) = 1, \text{ where } k(x) = a(x) - a(x-1), p = 2^n + j$$

But this is equivalent (by Lemma 4 above) to the fact that among a subset and its flip, exactly one contains $\{1\}$. $\square$

Proposition 6. *$A(n) > 0$ if n is not a power of 2.*

PROOF: $A(2^k + 1) = 2a(2^k + 1) - (2^k + 1) = 1$. When neither $n - 1$ nor n is a power of 2, $S(a(n))$ $(= C(S(n)))$ has one element less than $S(b(n))$ $(= T(S(n)))$, as $S(n)$ is neither a singleton nor of the form $\{1, 2, ..., k\}$. Thus it comes later in the enumeration, so $a(n) > b(n)$. Then

$$A(n) = (a(n) - b(n)) + (a(n) - a(n-1)) \geq a(n) - b(n) > 0$$

$\square$

8.3. A formula for $a(n)$. It is easy to prove, directly from the definition of our order, that if $S(n) = \{a_1, a_2, ..., a_k, M+1\}$, then

$$n = 1 + 2^M + \sum_{t=k+1}^{t=M} \binom{M}{t} + \sum_{j=1}^{k} \binom{a_j - 1}{j}$$

Since $S(a(n)) = \{a_2 - 1, a_3 - 1, .., a_k - 1, M\}$, this gives a formula for $a(n)$:

$$a(n) = 1 + 2^{M-1} + \sum_{t=k}^{M} \binom{M-1}{k} + \sum_{j=2}^{k} \binom{a_j - 2}{j-1}$$

(For $k = 0, 1$ some of the summations should be interpreted as being empty.) This is equivalent to the algorithm of G. Philips cited in [9], which also has a straightforward proof using the Pascal's triangle representation of R_n.

8.4. $a(n)/n \to 1/2$.

Proposition 7. $|a(n)/n - 1/2| < \frac{C}{\sqrt{\lfloor log_2(n-1) \rfloor}}$ *for large* n.

Proof: We bound $2a(n) - n$, or equivalenty $a(n) - b(n)$. When n or $n-1$ is a power of 2, $a(n) = b(n)$ so there is nothing to prove. Now suppose that n is in binomial segment $(L, K), 0 < K < L$. Then $S(a(n)), S(b(n))$ will be in binomial segments $(L-1, K-1)$ and $(L-1, K)$, respectively. Thus $a(n) - b(n)$ is bounded by the total length of these two binomial segments, $\binom{L-1}{K-1} + \binom{L-1}{K-2} = \binom{L}{K-1}$. This is at most the largest binomial coefficient on the L^{th} row of Pascal's triangle, approximately $\frac{2^L}{\sqrt{\pi L/2}}$ by Stirling's formula[3] [6]. This yields the above estimate with $C = \frac{1}{\sqrt{2\pi}} + \epsilon$, for any $\epsilon > 0$. $\square$

Corollary 3. $a(n)/n \to 1/2$ *as* $n \to \infty$

The corollary was proved by Conway, who asked for the rate of convergence and the last n for which $|a(n)/n - 1/2|$ exceeds $1/20$. The first question is answered, up to a constant, by the above estimate (the best such constant will be determined in the next section). The second was solved by Mallows, based on his analysis of the successive differences of $A(n)$. We will derive a rapid algorithm for answering the second question, with $1/20$ replaced by any given $\epsilon > 0$, in section 10.

[3]The approximation to the middle binomial coefficient, obtained from Stirling's formula, is always an overestimate. This fact is not needed for our result, which is only an asymptotic estimate.

8.5. The doubling inequality. We now prove Conway's doubling inequality $2a(n) \geq a(2n)$. Our proof uses the Pascal's triangle representation of the successive differences of $a(n)$. It is the same as the triangle of plus and minus signs used above for $A(n)$, but with plus and minus replaced by 1 and 0.

If n is the k^{th} term of a certain level, then $2n$ is the $2k^{\text{th}}$ position on the next level. Since the inequality is an equality at powers of 2, it suffices to show that the sum of the first $2k$ positions of a row of the Pascal triangle at most twice the sum of the first k positions of the previous row. Indeed, suppose that the first k positions contain elements of segments $0, 1, ..., j$ of the row, and no other segments. Schematically,

$$N = S_0 \quad S_1 \quad S_2 \quad ... \quad S_{j-1} \quad X$$

where X is a piece of the j^{th} segment. Then, schematically, on the next row,

$$2N = S_0 \quad S_0S_1 \quad S_1S_2 \quad S_2S_3 \quad ... \quad S_{j-2}S_{j-1} \quad S_{j-1}XY$$

where X and Y have the same length and XY is a substring of S_jS_j. (Note that $S_{j-1}X$ always lies within one segment, but the full string $S_{j-1}XY$ may spill over into the next segment S_jS_{j+1}.) Then $2a(n) - a(2n)$ is just the difference in sum between X and Y, and we wish to show that this is non-negative. Intuitively, it clear from the definition of the Pascal's triangle that 1's are prevalent on the left and 0's on the right, and that as one slides a "window" of length $|X|$ from left to right the density of ones among its occupants will diminish. This is not literally true, but the following is enough to prove the result.

Lemma 7. *Let P and Q be two strings of consecutive terms of the same length on a single row of the Pascal's triangle, such that either P begins at the same place as a segment or Q ends at the same place as a segment, and P is earlier than Q. Then the sum of the terms in P is at least the sum of the terms in Q.*

There are then 2 cases:

- (a) XY is a substring of S_j. We are done by the lemma, as X begins S_j.

- (b) XY is a substring of S_jS_j but not of S_j. Suppose $S_j = XY_1$ and $Y = Y_1Y_2$. Note that X has Y_2 as an initial substring: $X = Y_2X_1$ with $|X_1| = |Y_1|$, and Y_1 is the end of segment S_j. Then by the lemma Y_1 has at most the same sum as X_1.

Proof of Lemma: This is certainly true on level 0. Assume the result for the n^{th} row; we'll show the result for the $(n+1)^{\text{st}}$ row.

It suffices to show the result for P contained entirely within one segment. (We can assume that P and Q do not overlap. Say $P = P_0P_1...P_j$ (where each P_i is contained entirely within a segment, and for $i > 0$ begins that segment). Then Q could be divided up into $Q_1Q_2...Q_jQ_0$ (where P_i has the same length as Q_i), and the sum of the elements of P_i would be at least the sum of the elements of Q_i.) Similarly, it suffices to show the result for Q contained entirely within a single segment.

Say P lies in the segment that was the concatenation of segments A and B in the previous row of the Pascal's triangle, and Q lies in the segment that was the concatenation of segments C and D in the previous row. Then P and Q can be seen as strings of consecutive terms of the same length in the n^{th} row, with either P beginning with A or Q ending with D, so by the inductive hypothesis the sum of the terms in P is at least the sum of the terms in Q. □

The same argument shows $a(n) + a(n+1) \geq a(2n+1)$, so both cases can be subsumed in the more general inequality $a(\lfloor \frac{n}{2} \rfloor) + a(\lceil \frac{n}{2} \rceil) \geq a(n)$ (where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are the floor and ceiling functions, respectively). Such inequalities are typical of convex functions. Now, $A(n)$ is not convex (its graph wiggles up and down on a small scale), but on a large scale, the humps of $A(n)$ look very much convex. We will see next that this phenomenon is visible in the *asymptotic* shape of the humps.

9. Asymptotics of $a(n)$

We know that $A(n)$ has positive, symmetric humps between its zeros. Mallows observed that the shapes of these humps approach a limiting form. Specifically, if M_L is the maximum of $A(n)$ on level L, we may normalize the graph of the L^{th} hump to get a graph G_L of width 1 and height 1, by applying the affine map

$$(x, A(x)) \to (x/2^L, A(x)/M_L).$$

Linear interpolation gives a graph G_L which defines a continuous function $G_L : [1,2] \to [0,1]$. Mallows conjectured that these graphs converge to a certain limit curve G, parametrized in terms of the Gaussian distribution (see Figure 12). If true, this would provide essentially complete information about the asymptotics of $a(n)$; for example, if the tangent line from (0,0) to this curve has slope T, then the best possible constant C in our above estimate is $\frac{T}{2\sqrt{2\pi}} + \epsilon$.

We will prove the conjecture in the following form: if $x_n \in [2^n, 2^{n+1}]$, and $x_n/2^n \to s$, then $G_n(x_n) \to G(s)$, where $G : [1,2] \to [0,1]$ is the function whose graph has the parametrization:

$$x(t) = 1 + \frac{1}{\sqrt{\pi}} \int_{-\infty}^{t} e^{-z^2} dz, \qquad y(t) = e^{-t^2}, \quad -\infty < t < \infty$$

These complicated formulas disguise a simple combinatorial statement: on binomial segment (L,k), A is well approximated by $\frac{1}{2}\binom{L}{k}$, and in particular $M_L \approx \frac{1}{2}\binom{L}{L/2}$. Good approximation means that the error is negligible in comparison to M_L. Thus for values of x close to $2^n + \sum_{t=0}^{k} \binom{n}{t}$, $A(x)/M_n$ is approximately $\binom{n}{k}/\binom{n}{n/2}$. The

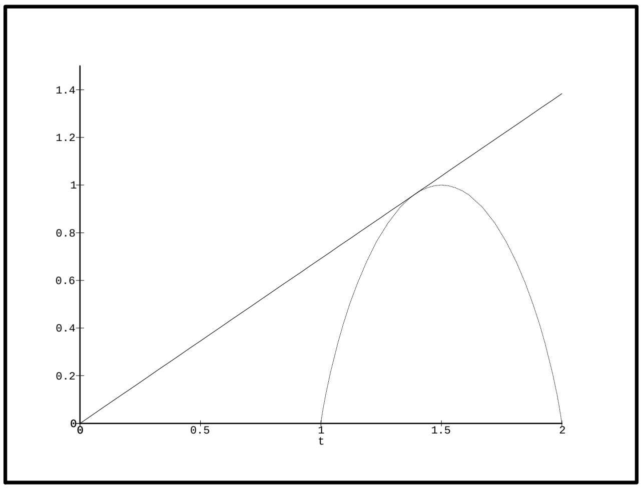

FIGURE 14. Asymptotic shape of the humps of $A(n)$

formula for the limit curve then follows from the asymptotics of binomial coefficients [6]:

$$\binom{n}{n/2} \approx \frac{2^n}{\sqrt{\pi n/2}}$$

$$\binom{n}{n/2+p\sqrt{n}} \approx e^{-2p^2}\binom{n}{n/2}$$

$$\frac{1}{2^n}\sum_{0\le t\le n/2+p\sqrt{n}}\binom{n}{t} \approx \frac{1}{\sqrt{\pi}}\int_{-\infty}^{\sqrt{2}p} e^{-t^2}dt$$

Theorem 3. *Suppose $0 \le k \le L$, and $K = min(k, L-k)$. Then on segment (L,k), the minimum value of A is $\binom{L-1}{K-1}+1$ (and 1 less than this when $k \ge L/2$). The maximum value is $\binom{L-1}{K}+\sum_{t=0}^{K-1}\frac{1}{t+1}\binom{2t}{t}$.*

PROOF: It suffices to derive the upper and lower bounds when $k \ge L/2$ and derive the rest from these using the symmetry of $A(n)$. Therefore, we will assume that $k \le L/2$ and estimate A on segment $(L, L-k)$. Suppose that $S(n) = \{a_1, ..., a_k, L+1\}$ belongs to this segment. Then $S(a(n)) = \{a_2-1, ..., a_k-1, L\}$. The formula of section 8, (ii), easily gives

$$2a(n) - n = Z - (a_1 - 1) + \sum_{i>1}[\binom{a_i-2}{i-1} - \binom{a_i-2}{i}]$$

where Z depends on L and k, but not on the a_i. Here a_i must lie in the range $\{i, ..., L-k+i\}$. In this range the minimum value of the term involving a_i $(i > 1)$ occurs when a_i is as large as possible, $a_i = L - k + i$, and the maximum when a_i is either $2i$ or $2i-1$ (the maximum occurs twice). Because $n \geq 2k$, we can assign values to the a_i to maximize, or minimize, all these terms simultaneously. The minimum occurs when $a_i = L-k+i$ (that is, for the last set of the segment), and the maximum when $a_1 = 1, a_i = 2i$ or $2i - 1, i > 2$. The corresponding extremal values of A are the ones stated, as is not hard to check by direct computation from the formula of section 8. □

Corollary 4. *The maximum value of $a(n)$ on segment (L, k) is attained 2^{K-1} times.*

Proof: There are two maximizing choices for each a_i, $i > 1$. □

This curious phenomenon is visible in the graphs of $A(n)$ for small n. Also, surprisingly, the Catalan numbers $\frac{1}{n+1}\binom{2n}{n}$ make an appearance in the upper bound.

Corollary 5. *If $L \geq 2$, then on segment (L, k), $\left|\frac{A(n)-\frac{1}{2}\binom{L}{k}}{M_L}\right| \leq \frac{3}{L}$ where M_L is the maximum of $A(n)$ on level L.*

Proof: The quantity to be estimated involves 2 parts: a difference of consecutive binomial coefficients, and a sum of Catalan numbers which is bounded by $\frac{2}{k}\binom{2k-2}{k-1}$, $k \leq n/2$. The first part is largest for $k = \lfloor\frac{n-1}{2}\rfloor$, the second for $k = \lfloor\frac{n}{2}\rfloor$, which gives the result after some simple manipulations. □

It follows that A is determined closely enough to imply the claimed limiting distribution. It's easy to check by differentiating the parametric equations that the curve is convex.

10. Estimating $a(n)/n$

The original problem that Conway posed about $a(n)$ was to find $N(\epsilon)$, the last value of n such that $|\frac{a(n)}{n} - \frac{1}{2}| > \epsilon$, for $\epsilon = \frac{1}{20}$. For this and other large values of ϵ, it is possible to use explicit bounds to reduce the problem to a direct search. The size of the answer grows as $2^{\epsilon^{-2}}$, so the brute force approach is not feasible in general. Mallows found that $N(\frac{1}{20}) = 1489$ and $N(\frac{1}{40}) = 6083008742$. Here, we outline a memory-efficient and rapidly convergent algorithm to find such numbers. It is easy to program, and in particular gives an independent verification of the $N(\epsilon)$ reported in [9].[4]

Let m_n be the maximum of $a(t)/t$ on level n.

[4]In [11], p. 59, the value of $N(1/20)$ is incorrectly given as 3173375556. Curiously, this number is an almost-correct answer to a different question: it is the last candidate value of $N(1/40)$ on level 31.

Lemma 8. *The sequence* (m_n) *is decreasing.*

PROOF: We will prove that for each n, $m_{n+1} \le m_n$. Suppose that for $2^n < t \le 2^{n+1}, a(t) \le m_n t$. Since $a(t)/t$ is minimized at powers of 2, this inequality will also hold in the slightly larger range $2^n \le t \le 2^{n+1}$. Note that for t on level $n+1$, both $a(t-1)$ and $b(t) = t - a(t-1)$ are between 2^n and 2^{n+1}. Thus,

$$\begin{aligned} a(t) &= a(a(t-1)) + a(t - a(t-1)) \\ &\le m_n a(t-1) + m_n(t - a(t-1)) \\ &\le m_n t \end{aligned}$$

□

Let T be the slope of the tangent line from $(0,0)$ to the curve G defined in the previous section (see Figure 12).

Lemma 9. $m_n\sqrt{n} \to \frac{T}{2\sqrt{2\pi}}$ *as* $n \to \infty$

PROOF: Clearly T is the maximum value of $G(s)/s, 1 \le s \le 2$. The lemma then follows from the very definition of G: if $x \approx s2^n$, then $A(x) \approx G(s)M_n$, where $M_n \approx \frac{1}{2}\binom{n}{n/2}$. □

Direct computation yields $T = 0.691905$, $C = \frac{T}{2\sqrt{2\pi}} = 0.138015$.[5] It follows that the level of $N(\epsilon)$ is about $C\epsilon^{-2}$. This turns out to be quite an accurate estimate of the critical level (see data below). Our strategy is to use this estimate, combined with a good algorithm to search a fixed level, to find L such that level $L+1$ contains no candidates for $N(\epsilon)$ but level L does; the value found on level L will then be $N(\epsilon)$. So, we proceed to describe the algorithm to search one level. First we need a technical lemma.

Let $P(n,k)$ denote the string of plus and minus signs in the $(n,k)^{\text{th}}$ entry of the Pascal triangle which gives the differences of $A(n)$. Reading this string left to right can be interpreted as a summation in an obvious way.

Lemma 10. *The sum of* $P(n,k)$ *is given by* $S(n,k)$ *with* $S(n,k) = \binom{n-1}{k} - \binom{n-1}{k-1}$. *The largest partial sum of* $P(n,k)$ *is given by* $M(n,k)$ *with* $M(n,k) = \binom{n-1}{K} - \binom{n-1}{k-1} + \sum_{t=0}^{K-1} \frac{1}{t+1}\binom{2t}{t}$, $K = min(k, n-k)$.

PROOF: An easy induction on n proves the first formula. The second follows from the fact that $M(n,k) + \binom{n-1}{k-1}$ is the maximum of A on segment (n,k). □

Here is the basic idea of our algorithm; it allows us to use the Pascal's triangle machinery to obtain good local upper bounds on $A(n)/n$. Suppose that we know

[5] In [9], C is obtained as the slope of the tangent line, since the rescaling used there leads to a curve parametrized by $(2x(t), \frac{1}{\sqrt{2\pi}}y(t))$

the value $A(p)$, and that the successive differences of $A(n)$ in the range $p < n \leq q$ are known. Suppose that we also know that the largest partial sum of the successive differences in this range is M. Then, since $A(p) \leq p$, and the differences are all ± 1, in the interval under consideration we have

$$\frac{A(n)}{n} \leq \frac{A(p)+M}{p+M'}, \quad M' = max(1, M)$$

We can use this method on segment (L,k). The initial value $A(p)$ will just be $\binom{L-1}{k-1}$, and $M = M(n,k)$. This gives an upper bound on segment (L,k). Computing these upper bounds on all segments of level L, we eliminate some segments as possible locations for $N(\epsilon)$. What about the segments not ruled out? As the successive differences on segment (L,k) are formed by concatenating the differences for $(L-1, k-1)$ and $(L-1,k)$, each surviving segment can be broken into two subsegments, and our upper bound formula can be applied to each. Iterating this process, we get an efficient method of finding the last candidate value for $N(\epsilon)$ on a given level, or proving that none exists.

Here is a more formal description of the algorithm:

(1) Define $U(n,k,f,p) = \frac{f+M(n,k)}{p+max(1,M(n,k))}$
(2) Input: a positive integer L (the level), and $\epsilon > 0$. The algorithm will determine the last n on level L such that $\frac{A(n)}{n} > \epsilon$, or prove that none exists.
(3) Create an ordered list V of integer vectors:

$$\begin{aligned} V &= v_0, v_1, \ldots, v_{L-1} \\ v_k &= (L, k, 2^L + \sum_{t=0}^{k-1} \binom{L}{t}, \binom{L-1}{k-1}) \end{aligned}$$

(4) Eliminate from V all vectors v such that $U(v) \leq \epsilon$
(5) If V is empty, stop; no such n exists on level L.
(6) If the last term of V is of the form $v = (m, 0, f, p)$ or $v = (m, m, f, p)$, and $U(v) > \epsilon$ stop; the answer is $n = p+1$. Reason: the upper bound is exact on intervals of length 1, and we have reached such an interval where the upper bound exceeds ϵ.
(7) If the last term of V is $v_t = (n,k,f,p)$, replace it by the two terms $v_t = (n-1, k-1, f, p)$ and $v_{t+1} = (n-1, k, f + S(n-1, k-1), p + \binom{n-1}{k-1})$. If either of these vectors satisfies $U(v) \leq \epsilon$, eliminate it.
(8) Go to step 4.

The space required by this algorithm is very reasonable: about $2L^2$ bits. In other words, if the answer has N digits, then the algorithm needs storage space on the order of N^2. The proof is simple and elegant, so we give it here. Essentially the only storage required is for V, the list of integer vectors. Each vector in V is a quadruple of integers (x, y, a, b) where $0 \leq y \leq x \leq L$ and $0 \leq a, b < 2^{L+1}$, which requires about $2L$ bits of storage. Now we will show that the number of vectors on the list never exceeds L. Assign to each vector (x, y, a, b) the non-negative label $x - y$. Then it is easy to show that the labels form a strictly decreasing sequence of integers between 1 and L, each time the algorithm loops.

As for the runtime, we do not know if it is bounded by a polynomial in L, although for our purposes it was adequate. Implemented in Mathematica on a Sun 4 computer, it took under 5 minutes to compile the following data.

ϵ	Predicted Level	Critical Level	$N(\epsilon)$
1/20	8	10	1489
1/30	18	19	758765
1/40	31	32	6083008742
1/50	48	49	809308036481621
1/60	69	70	1684539346496977501739
1/70	94	95	55738373698123373661810220400
1/80	122	123	15088841875190938484828948428612052839

11. Generalizations of $a(n)$

Some generalizations of $a(n)$ were proposed by D. Newman [8]. They are defined by setting the first $N + 1$ values to 1, and computing the rest of the values from the now-familiar recurrence $a(n) = a(a(n-1)) + a(n - a(n-1))$. For $N = 0$ we get $a(n) = n$, and $N = 1$ yields Conway's $a(n)$. Newman observed that for each N, there is a sequence of special values playing the same role as the powers of 2 in our situation. For $N = 2$, these are the Fibonacci numbers, and in general the special values satisfy the recurrence $E_n = E_{n-1} + E_{n-N}$.

It is possible to generalize our interpretation of $a(n)$ to these higher analogues. In particular, we can explain the special values E_n and the recurrence they satisfy. We state only the results, since the proofs are completely parallel to the ones above.

We may associate a graph to the sequence as above, and from it derive frequency counts Q_n.

Define trees T_n as follows: $T_1 = 1, T_2 = 2, ... T_N = N$; for $n > N$, T_n has n at the top, connected to copies of $T_1, T_2, ..., T_{n-N}$ on the next row.

Define the sequence R_n as the result of reading tree T_n, starting from the bottom row, and reading each row left to right. Alternatively, $R(1) = 1, R(2) = 2, ..., R(N) = N$, and for $n > N$, $\quad R(n) = \mathsf{e}(R(n-N))n$.

Define $R_n^* = R_n$ with the last term increased by 1. Then $Q(0) = N$, and for $n > 0$, $Q(n)$ is the concatenation of N consecutive sequences R_i^*: $\quad Q_n = R^*((n-1)N + 1) ... R^*(nN)$.

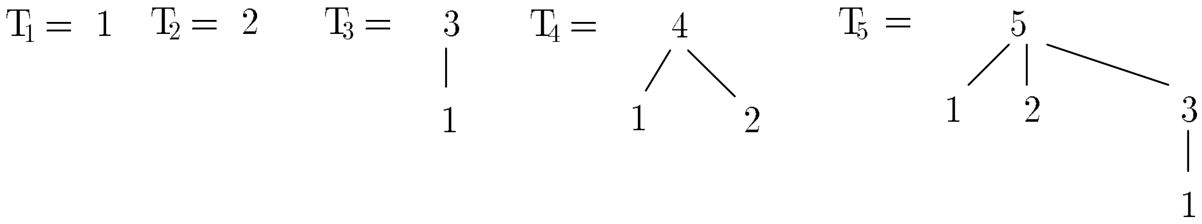

FIGURE 15. Trees for the order-2 $a(n)$ recurrence

The set model is again derived from the trees T_n. It uses the same ordering as above, but applied only to sets where any two distinct elements differ by at least N. Thus for $N = 2$ the first few sets are

$$\varnothing < \{1\} < \{2\} < \{1,3\} < \{3\} < \{1,4\} < \{2,4\}$$

$$< \{1,3,5\} < \{1,5\} < \{2,5\} < \{3,5\} < \{5\} < \ldots$$

The operation corresponding to $a(n)$ is: remove the smallest element of $S(n)$ (if $S(n)$ is not a singleton), then decrease the remaining elements by N. This operation maps $\{k\}$ to $\{k-N\}$; letting E_k be the positions of $\{k\}$ in the enumeration of sets, we find that $E_n = E_{n-1} + E_{n-N}$, and $a(E_n) = a(E_{n-N})$. This explains Newman's observation (albeit sketchily).

To prove that this rule for $a(n)$ works, we again need a similar rule for $b(n) = n - a(n-1)$. There is no longer a direct correspondence between the frequency counts of $a(n)$ and $b(n)$, so finding the operation for the latter function takes some more thought in this case. It turns out that $b(n)$ corresponds to the following operation: if $S(n)$ doesn't contain 1, decrease all elements by 1; if $S(n)$ is saturated, *i.e.* it contains 1 and consecutive elements always differ by exactly N, then replace the last two elements $\{M-N, M\}$ by $\{M-1\}$; otherwise reduce the first gap by 1, as in the case $N = 1$.

The reduced frequency counts for the N^{th} order sequence may be obtained from the Pascal's triangle of the first-order sequence. Starting from the Pascal triangle defined in section 4, we may obtain R_n by reading the terms at positions (t,k) on the line $Nt - (N-1)k = n-1$. This shows that the sums along slope $(1-N)/N$ diagonals of the numerical Pascal's triangle are the numbers E_n; the cases $N = 2$ is a well-known property of the Fibonacci sequence.

Asymptotically, $a(n) \approx \lambda^{-N} n$, where λ is the largest root of the characteristic polynomial of the sequence E_n, *i.e.* $\lambda^N = \lambda^{N-1} + 1, \quad \lambda > 1$. The graph of $\delta(n) = \lambda^N a(n) - n$ between consecutive special values, after a suitable rescaling, converges to a curve explicitly parametrized by quadratic exponentials similar to the case $N = 1$. The corresponding combinatorial assertion is that on the segment corresponding to position (n, k) of the Pascal triangle, $\delta(n) \approx \lambda^{-N}\binom{n}{k}$. The role of the central binomial coefficient is played by $\binom{n}{\frac{n}{\lambda}}$.

There seems to be no obvious generalization of $A(n)$ which enjoys the same nice properties. Clearly we would like an integer-valued function which measures the distance between $a(n)$ and its linear approximation, that is something approximately proportional to $\lambda^N a(n) - n$. One candidate is $h(a(n)) - n$, where $h(t)$ is an analogue of doubling, formed by expressing t as a sum of E_i's and replacing each E_i by E_{i+N}.[6] This function vanishes at all E_i and seems to be non-negative elsewhere, although we don't have a proof. It is not clear what sort of symmetry to expect, if any. If $h(n)$ does express the right generalization of doubling, then the expected generalization of Conway's inequality would be $h(a(n)) \geq a(h(n))$.

Variants of $a(n)$ with different settings of the initial values give similar, though less elegant, sequences. These sequences are also controlled by R_n-type frequency counts. For example, setting the first 2 values equal to 1 and the next 3 values equal to 2, we obtain an $a(n)$ whose level n frequency count is a sequence of length $3 \times 2^{n-1}$ formed by concatenating: 1; Q_{n-1} with last element increased by 1; and T_n read from the top down, with the last element omitted. This version of $a(n)$ enjoys a double flip-flop property, in that the frequency count can be broken into two pieces, each of which is flip-flop.

12. Further Questions

Why does the recurrence $a(n) = a(a(n-1)) + a(n - a(n-1))$ lead to sequences with such elaborate combinatorial structure? Experimentation shows that the same type of structure persists if the sequence is "seeded" with different initial values. We suspect that there is a more general structure at work. There are some clues pointing to a connection with combinatorial games:

- Recurrences of the form $W_n = W_{n-1} + W_{j(n)} \quad (j(n) < n)$ arise in the winning strategies of generalized one-pile Nim [5]. To each positive, increasing integer function $f(n)$ there corresponds, via the game, such a sequence. The special values of k^{th}-order $a(n)$ correspond to the function $f(n) = kn$. Do generalized Nim games yield good generalizations of $a(n)$? The associated set models would deal with sets such that, if t is in the set, and $W_t = W_{t-1} + W_k$, then the next largest element in the set is at most k.
- The coin-pushing interpretation of $a(n)$ is reminiscent of the "chips and strips" games in [4]. Does $a(n)$ encode properties of games of this sort, such as remoteness, winning positions, or G-values?
- The trees T_n are defined in the same way as finite ordinals, where the n'th is the union of all its predecessors. Conway has defined some generalized ordinals

[6] This is an excellent approximation as well, since there exists a γ between 0 and 1 such that the error $|\lambda^N n - h(n)| < n^\gamma$.

using two-player games [3]. Can sequences satisfying the $a(n)$ recurrence be classified in terms of these structures?

A related question is how to deduce the recurrence, given the set ordering and the compression function. Are there other natural examples of an ordering of finite sets and a compression function leading to such nonlinear, nonlocal recurrences?

It's easy to experiment numerically with $a(n)$-like sequences on a computer (or even by hand), and we can't help but think that a bit of careful observation will unravel the mystery. We hope we have provoked the reader into exploring these fascinating sequences.

Acknowledgements. We are indebted to Steve Tanny (U. of Toronto), Colin Mallows (AT&T Bell Laboratories), John H. Conway (Princeton U.), and especially Dan Kleitman (M.I.T.), for their comments on various stages of this work. We also thank Ian Goldberg (U. of Waterloo) and Hugh Thomas (U. of Toronto) for conversations about an earlier version of this article.

References

1. Robin J. Chapman, On a Problem of Conway, 1988, unpublished.
2. B.W. Conolly, Fibonacci and Meta-Fibonacci sequences, in S. Vajda (ed.), *Fibonacci & Lucas Numbers and the Golden Section: Theory and Application*, E. Horwood Ltd. 1989, 127-139
3. John H. Conway, *On Numbers and Games*, Academic Press, New York 1976.
4. John H. Conway, Richard K. Guy and Elwyn R. Berkelamp, *Winning Ways, for Your Mathematical Plays*, Academic Press, New York 1982.
5. Robert J. Epp and Thomas S. Ferguson, "A Note on Take Away Games", *Fibonacci Quarterly* 18(1980), 300-303.
6. William Feller, *An Introduction to Probability Theory and its Applications, v. 1.*, John Wiley & Sons, Inc., New York 1957.
7. Douglas R. Hofstadter, letters to Colin Mallows and John H. Conway, September 2, 1988.
8. Dan Kleitman, Solution to AMM Problem 3274 (Another Fibonacci Generalization, proposed by David Newman), *American Mathematical Monthly* 98(1991), 958-959.
9. Colin Mallows, Conway's Challenge Sequence, *American Mathematical Monthly* 98(1991), 5-20.
10. Colin Mallows, Letter to the Editor, *American Mathematical Monthly* 99(1992), 547-548.
11. Manfred R. Schroeder, *Fractals, Chaos, Power Laws: Minutes from an Infinite Universe*, W.H. Freeman, New York. 1990.
12. Allen J. Schwenk, "Take-away games", *Fibonacci Quarterly* 8(1970), 225-234, 241.
13. Mario Szegedy, *Conway-Mallows sequence versus Queue game*, unpublished. Dec. 16, 1991.

Department of Mathematics, Harvard University
E-mail address: kubo@math.harvard.edu ravi@math.harvard.edu