The Cube Recurrence

Gabriel D. Carroll
Harvard University
Cambridge, Massachusetts 02138
gcarroll@fas.harvard.edu

David Speyer
Department of Mathematics
University of California at Berkeley
Berkeley, California 94720
speyer@math.berkeley.edu

Submitted: Dec 31, 2002; Accepted: Jul 6, 2004; Published: Oct 18, 2004
Keywords: cube recurrence, grove, Gale-Robinson sequence
MR Subject Classifications: 05A15, 05E99, 11B83

Abstract

We construct a combinatorial model that is described by the cube recurrence, a quadratic recurrence relation introduced by Propp, which generates families of Laurent polynomials indexed by points in $\mathbb{Z}^3$. In the process, we prove several conjectures of Propp and of Fomin and Zelevinsky about the structure of these polynomials, and we obtain a combinatorial interpretation for the terms of Gale-Robinson sequences, including the Somos-6 and Somos-7 sequences. We also indicate how the model might be used to obtain some interesting results about perfect matchings of certain bipartite planar graphs.

1 Introduction

Consider a family of rational functions $f_{i,j,k}$, indexed by $(i,j,k) \in \mathbb{Z}^3$ with $k \geq -1$, and given by the initial conditions $f_{i,j,k} = x_{i,j,k}$ (a formal variable) for $k = -1$ and 0, and

$$f_{i,j,k-1}f_{i,j,k+1} = f_{i-1,j,k}f_{i+1,j,k} + f_{i,j-1,k}f_{i,j+1,k} \quad (k \geq 0).$$

This is the *octahedron recurrence*, which has connections with the Hirota equation in physics, with Dodgson's condensation method of evaluating determinants, with alternating-sign matrices, and with domino tilings of Aztec diamonds (see [11], [8], [9], [3], respectively). It turns out that every $f_{i,j,k}$ is a Laurent polynomial in the initial variables $x_{i,j,k}$, i.e. a polynomial in the variables $x_{i,j,k}$ and $x_{i,j,k}^{-1}$. Sergey Fomin and Andrei Zelevinsky, using techniques from the theory of cluster algebras, proved in [4] that the recurrence again generates Laurent polynomials for a large variety of other initial sets (i.e., sets of points (i, j, k) for which we designate $f_{i,j,k} = x_{i,j,k}$). In [10], David Speyer showed further that all such polynomials could be interpreted as enumerating perfect matchings of suitable bipartite planar graphs, generalizing the main result of [3].

In the present paper, we study not the octahedron recurrence but the related *cube recurrence*, suggested by James Propp in [8], in the form $f_{i,j,k} = x_{i,j,k}$ for $i+j+k = -1, 0, 1$, and

$$f_{i,j,k}f_{i-1,j-1,k-1} = f_{i-1,j,k}f_{i,j-1,k-1} + f_{i,j-1,k}f_{i-1,j,k-1} + f_{i,j,k-1}f_{i-1,j-1,k} \quad (i + j + k > 1).$$

Fomin and Zelevinsky showed in [4] that the cube recurrence also generates Laurent polynomials and conjectured that all the coefficients of these polynomials were positive. Propp noticed empirically that each coefficient in these polynomials is in fact equal to 1 and that each variable takes only exponents in the range $-1, \ldots, 4$. Our goal is to construct combinatorial objects that are in bijection with the terms of these Laurent polynomials, under a generalized form of the cube recurrence; Propp's observations, among other interesting results, will then follow directly.

Aside from the aesthetic appeal, an advantage of the combinatorial approach is that it provides a specific and detailed interpretation for the terms of the Laurent polynomials and thus allows deeper results — for example, the total positivity that Fomin and Zelevinsky had been unable to prove using purely algebraic methods. Much more generally, Fomin and Zelevinsky conjectured in [5] that all Laurent polynomials arising from cluster algebras have only positive coefficients. Although it is too early to expect anything like our methods to resolve this more general conjecture, combinatorics does provide fairly powerful methods for exploring the algebra of at least some interesting types of polynomial recurrences. In the present instance, this includes not only the cube recurrence itself but simpler one-dimensional recurrences such as

$$s_n = 3s_{n-1}s_{n-2}/s_{n-3}$$

$$s_n = (s_{n-1}s_{n-5} + s_{n-2}s_{n-4} + s_{n-3}^2)/s_{n-6}$$

$$s_n = (s_{n-1}s_{n-6} + s_{n-2}s_{n-5} + s_{n-3}s_{n-4})/s_{n-7}$$

The first of these one-dimensional recurrences results from taking $s_n = f_{n,0,0}$ in the form of the cube recurrence given above (or more generally $s_n = f_{i,j,k}$ for any i, j, k with $i + j + k = n$), and it is then satisfied by $s_n = 3^{\lfloor n^2/4 \rfloor}$, by an easy induction. The second and third, when we set $s_n = 1$ for $1 \leq n \leq 6$ (resp. $1 \leq n \leq 7$), are the *Somos-6* and *Somos-7 sequences* (see [6]), which are also really specializations of the cube recurrence. These sequences' definitions are simple enough that it would not be surprising to see them naturally crop up elsewhere, and the combinatorial machinery we will develop provides a way of understanding them in detail — an initially unexpected way, to boot. We should point out, however, that the generality of the combinatorial approach is a subtle thing: although the structure of our proof parallels that used by Speyer for the octahedron recurrence, the combinatorial objects we construct are quite different.

It may be worth mentioning that the cube recurrence can be written in a slightly more symmetrical fashion than above: the families of functions $(f_{i,j,k})$ satisfying the cube recurrence can be made to correspond to the families $(g_{i,j,k})$ satisfying the recurrence

$$g_{i,j,k}g_{i-1,j-1,k-1} + g_{i-1,j,k}g_{i,j-1,k-1} + g_{i,j-1,k}g_{i-1,j,k-1} + g_{i,j,k-1}g_{i-1,j-1,k} = 0$$

by taking $g_{i,j,k} = -f_{i,j,k}$ when $i + j + k \equiv 0 \bmod 4$, $g_{i,j,k} = f_{i,j,k}$ otherwise. This latter equation has the advantage of being invariant not only under translation and permutation of coordinates but also under reflections (e.g. the substitution $i \leftarrow -i$). However, we will not make further use of it here.

2 The recurrence

We will consider the polynomials generated by the cube recurrence using various sets of initial conditions. In order to describe these initial conditions, we will need to develop some notation.

Define the *lower cone* of any $(i, j, k) \in \mathbb{Z}^3$ to be

$$C(i, j, k) = \{(i', j', k') \in \mathbb{Z}^3 \mid i' \leq i, j' \leq j, k' \leq k\}.$$

Let $\mathcal{L} \subseteq \mathbb{Z}^3$ be a nonempty subset such that, whenever $(i, j, k) \in \mathcal{L}, C(i, j, k) \subseteq \mathcal{L}$. (Thus, $\mathcal{L}$ is an order-ideal in $\mathbb{Z}^3$ under the standard product ordering.) Let $\mathcal{U} = \mathbb{Z}^3 - \mathcal{L}$, and define the *initial set*

$$\mathcal{I} = \{(i, j, k) \in \mathcal{L} \mid (i + 1, j + 1, k + 1) \in \mathcal{U}\}.$$

To each $(i, j, k) \in \mathcal{I}$ we assign a formal *vertex variable* $x_{i,j,k}$. We also define *edge variables* $a_{j,k}, b_{i,k}, c_{i,j}$ for all $i, j, k \in \mathbb{Z}$; the reason for this terminology will become clear later.

Now let $f_{i,j,k} = x_{i,j,k}$ for $(i, j, k) \in \mathcal{I}$. When $(i, j, k) \in \mathcal{U}$ and $C(i, j, k) \cap \mathcal{U}$ is finite, we recursively define

$$f_{i,j,k} = \frac{b_{i,k}c_{i,j}f_{i-1,j,k}f_{i,j-1,k-1} + c_{i,j}a_{j,k}f_{i,j-1,k}f_{i-1,j,k-1} + a_{j,k}b_{i,k}f_{i,j,k-1}f_{i-1,j-1,k}}{f_{i-1,j-1,k-1}}. \tag{1}$$

(We leave $f_{i,j,k}$ undefined for all other points (i, j, k).)

By successive expansion of the f's on the right side using equation (1) repeatedly, we get that $f_{i,j,k}$ is a well-defined rational function in the vertex and edge variables $\{a_{j',k'}, b_{i',k'}, c_{i',j'}, x_{i',j',k'} \mid (i', j', k') \in C(i, j, k) \cap \mathcal{I}\}$, which takes a positive value when all the variables are set to 1. To verify definedness and positivity, use induction on $|C(i, j, k) \cap \mathcal{U}|$: either $(i-1, j, k) \in \mathcal{I}$, or $(i-1, j, k) \in \mathcal{U}$ and $|C(i-1, j, k) \cap \mathcal{U}| < |C(i, j, k) \cap \mathcal{U}|$; similarly for $(i, j - 1, k - 1)$, and so forth. (We need positivity in the induction hypothesis to ensure that the recurrence never produces a division by 0.) Also, to see that all edge variables appearing in $f_{i,j,k}$ really are of the form $a_{j',k'}, b_{i',k'}$, or $c_{j',k'}$ for some $(i',j',k') \in C(i,j,k) \cap \mathcal{I}$, just notice (for example) that $(i,j,k') \in \mathcal{I}$ for some $k' \le k$: simply choose the maximal k' for which $(i,j,k') \in \mathcal{L}$, as some such k' must exist, by finiteness.

Henceforth, we will only investigate the value of $f_{i,j,k}$ for one particular (i,j,k). Because the definitions of $\mathcal{L},\mathcal{U},\mathcal{I}$ and the recurrence itself are invariant under translation (modulo some relabeling of variables), we may assume that $(i,j,k) = (0,0,0)$. We may also make another simplifying assumption: Let $\mathcal{L}' = \mathcal{L} \cap C(0,0,0)$, and define $\mathcal{U}',\mathcal{I}'$ analogously to $\mathcal{U},\mathcal{I}$. It is easy to check that $\mathcal{I} \cap C(0,0,0) \subseteq \mathcal{I}'$. In particular, this means that running the cube recurrence gives the same value for $f_{0,0,0}$ regardless of whether we use $\mathcal{I}$ or $\mathcal{I}'$ as our initial set, so we can safely replace $\mathcal{L}$ by $\mathcal{L}'$. Therefore, we assume henceforward that $\mathcal{L} \subseteq C(0,0,0)$ and that $C(0,0,0) \cap \mathcal{U}$ is finite, except where clearly stated otherwise.

One further definition related to initial sets will prove helpful. We call a nonnegative integer N a *cutoff* for the initial set $\mathcal{I}$ if it satisfies the two properties:

(i) $(i,j,k) \in \mathcal{I}$ whenever $i+j+k \le -N$ and $\max\{i,j,k\} = 0$;

(ii) $(i,j,k) \notin \mathcal{I}$ when $i+j+k \le -N-3$ and $\max\{i,j,k\} < 0$.

We write both conditions for clarity, but in fact (ii) implies (i): if $i+j+k \le -N$, choose the lowest integer l such that $(i-l,j-l,k-l) \in \mathcal{L}$ (some l exists by finiteness), and then $(i-l+1,j-l+1,k-l+1) \in \mathcal{U}$, so $(i-l,j-l,k-l) \in \mathcal{I}$; we have $\max\{i,j,k\} = 0$ implying $l \ge 0$, but if $l > 0$ then $(i-l,j-l,k-l)$ violates (ii), so $l = 0$, and (i) holds. Also note that if $-N < \min\{i+j+k \mid (i,j,k) \in \mathcal{U} \cap C(0,0,0)\}$, then N is a cutoff — condition (ii) is easy to check. Thus, a cutoff is a measure of the size of the "interesting" region of the initial set $\mathcal{I}$, and finiteness of $\mathcal{U} \cap C(0,0,0)$ ensures that a cutoff always exists.

3 Groves

We will now introduce the combinatorial objects that will form the basis for our understanding of the cube recurrence. We first provide the definition that will be most useful for purposes of subsequent proofs; later, we will offer an alternative representation that may be more practical as a kind of "shorthand." We assume throughout that $\mathcal{L},\mathcal{U},\mathcal{I}$ are as described in Section 2.

One preliminary notion that will prove crucial is that of a *square*. For each point $(i,j,k) \in \mathcal{I}$, we define the following three sets:

$$\begin{aligned}
s_a(i,j,k) &= \{(i,j,k),(i,j-1,k),(i,j,k-1),(i,j-1,k-1)\} \\
s_b(i,j,k) &= \{(i,j,k),(i-1,j,k),(i,j,k-1),(i-1,j,k-1)\} \\
s_c(i,j,k) &= \{(i,j,k),(i-1,j,k),(i,j-1,k),(i-1,j-1,k)\}
\end{aligned}$$

We then define a square to be any set of the form $s_a(i,j,k), s_b(i,j,k)$, or $s_c(i,j,k)$ that is contained in $\mathcal{I}$. Each square can be decomposed into two pairs of points as follows:

$$e_a(i,j,k) = \{(i,j-1,k),(i,j,k-1)\} \qquad e'_a(i,j,k) = \{(i,j,k),(i,j-1,k-1)\}$$

$$
\begin{aligned}
e_b(i,j,k) &= \{(i-1,j,k),(i,j,k-1)\} & \qquad e'_b(i,j,k) &= \{(i,j,k),(i-1,j,k-1)\}\\
e_c(i,j,k) &= \{(i-1,j,k),(i,j-1,k)\} & \qquad e'_c(i,j,k) &= \{(i,j,k),(i-1,j-1,k)\}
\end{aligned}
$$

We refer to $e_a(i,j,k)$ as the *long edge* and $e'_a(i,j,k)$ as the *short edge* associated with $s_a(i,j,k)$, and similarly for the other pairs.

We now construct an (infinite) graph $\mathcal{G}$ whose vertices are the points in $\mathcal{I}$. The edges of $\mathcal{G}$ are simply the long and short edges of all squares occurring in $\mathcal{I}$. Of course, these are not what one normally thinks of as the "edges" of a square, nor are those edges ordinarily different lengths, but the motivation for our terminology becomes clearer when we project $\mathcal{I}$ onto a plane. An example is shown in Figure 1, where $\mathcal{L} = \{(i,j,k \in C(0,0,0) \mid i+j+k \leq -4\}$. The projection of each square (really, the projection of each square's convex hull) is a parallelogram, shown shaded in the figure. The gray lines represent edges of $\mathcal{G}$: the short edge and long edge associated with a square project to the short and long diagonals of the corresponding parallelogram. Any choice of $\mathcal{I}$ will in fact produce a tiling of the plane by parallelograms; we do not prove this here, since it is secondary to our main concerns, although it will become apparent from techniques to be introduced subsequently.

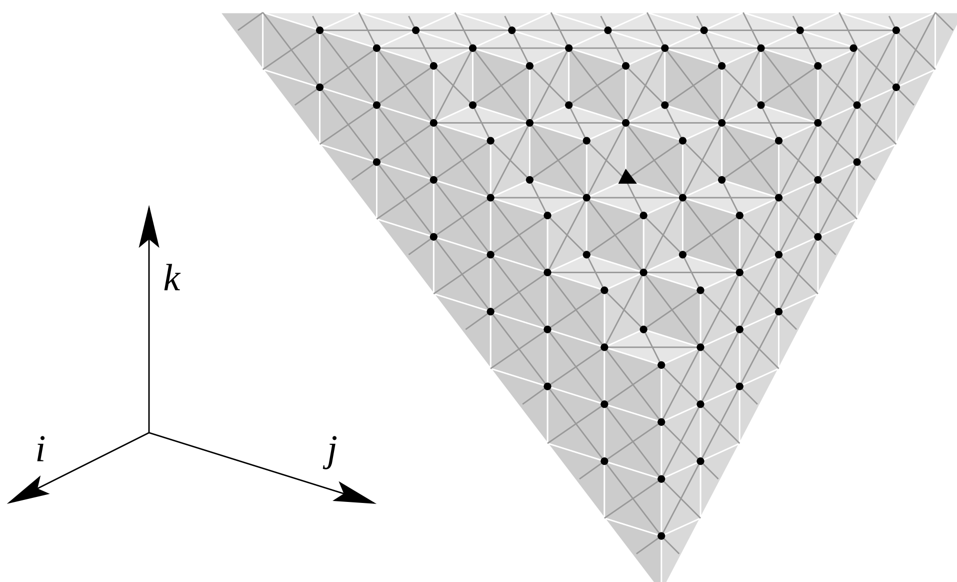

Figure 1: An example $\mathcal{G}$; anchor is $(-2,-2,-2)$

If we project onto the plane $i+j+k=0$, these parallelograms will actually be rhombi. However, we find it visually clearer to project onto a less symmetrical plane of the form $\alpha i + \beta j + \gamma k = 0$ for α, β, γ close to 1. We will also usually mark one of the points of $\mathcal{I}$ (which we call the *anchor*) by a triangle in the figure, rather than a circle, and will indicate its coordinates explicitly, to help the reader get oriented.

We will find it useful to describe the squares in the outer, "uninteresting" regions of $\mathcal{I}$.

Lemma 1 *Suppose N is a cutoff for $\mathcal{I}$. Then each of the following points belongs to four squares of $\mathcal{I}$, as listed, and no other squares:*

(i) If $p < -N$, then

- $(p,0,0) \in s_b(p,0,0), s_b(p+1,0,0), s_c(p,0,0), s_c(p+1,0,0)$;
- $(0,p,0) \in s_c(0,p,0), s_c(0,p+1,0), s_a(0,p,0), s_a(0,p+1,0)$;
- $(0,0,p) \in s_a(0,0,p), s_a(0,0,p+1), s_b(0,0,p), s_b(0,0,p+1)$.

(ii) If $p, q < 0$ and $p+q < -N-1$, then

- $(0,p,q) \in s_a(0,p,q), s_a(0,p+1,q), s_a(0,p,q+1), s_a(0,p+1,q+1)$;
- $(p,0,q) \in s_b(p,0,q), s_b(p+1,0,q), s_b(p,0,q+1), s_b(p+1,0,q+1)$;
- $(p,q,0) \in s_c(p,q,0), s_c(p+1,q,0), s_c(p,q+1,0), s_c(p+1,q+1,0)$.

Proof: We prove the first statement in each triple, as the others are analogous.

(i) It is clear that $(p,0,0)$ belongs to the four specified squares, if they exist — that is, if the specified sets s_b and s_c really are contained in $\mathcal{I}$. But it is straightforward to check that all the points in these sets lie in $\mathcal{I}$ by condition (i) for a cutoff. To see that no other squares contain $(p,0,0)$, consider all the possible sets of the form $s_a(i,j,k)$, $s_b(i,j,k)$, or $s_c(i,j,k)$, for $(i,j,k) \in \mathbb{Z}^3$, that would contain $(p,0,0)$. A priori there are up to twelve such sets (four of each type). But aside from the four specified in the lemma, seven of the others contain a point whose maximum coordinate is $+1$ and so clearly cannot lie in $\mathcal{I}$. The only remaining possibility is $s_a(p,0,0)$. But if $s_a(p,0,0) \subseteq \mathcal{I}$, then $(p,-1,-1) \in \mathcal{I} \Rightarrow (p+1,0,0) \in \mathcal{U}$. Since condition (i) for a cutoff gives $(p+1,0,0) \in \mathcal{I}$, we have a contradiction, and this shows that $(p,0,0)$ belongs to no squares other than the four specified.

(ii) The proof is quite similar. Again it is straightforward to see that the four specified squares are contained in $\mathcal{I}$ and contain $(0,p,q)$. And again we have twelve potential squares containing $(0,p,q)$, but now only four of them — $s_b(1,p,q), s_b(1,p+1,q), s_c(1,p,q), s_c(1,p,q+1)$ — can be eliminated on the grounds of containing a point with coordinate $+1$. However, the remaining four squares that need to be eliminated are $s_b(0,p,q), s_b(0,p+1,q), s_c(0,p,q), s_c(0,p,q+1)$. Each of these would have to contain the point $(-1,p,q)$. But if $(-1,p,q) \in \mathcal{I}$ then $(0,p+1,q+1) \in \mathcal{U}$, contradicting $(0,p+1,q+1) \in \mathcal{I}$ from the cutoff condition. So again we have a contradiction.

□

Now for our main definition; suppose that N is a cutoff for $\mathcal{I}$. We define an *$\mathcal{I}$-grove within radius N* to be a subgraph $G \subseteq \mathcal{G}$ with the following properties:

- (Completeness) the vertex set of G is all of $\mathcal{I}$;
- (Complementarity) for every square, exactly one of its two edges occurs in G;

- (Compactness) for every square all of whose vertices satisfy $i + j + k < -N$, the short edge occurs in G;
- (Connectivity) every component of G contains exactly one of the following sets of vertices, and conversely, each such set is contained in some component:
  - $\{(0, p, q), (p, 0, q)\}, \{(p, q, 0), (0, q, p)\}$, and $\{(q, 0, p), (q, p, 0)\}$ for all p, q with $0 > p > q$ and $p + q \in \{-N - 1, -N - 2\}$;
  - $\{(0, p, p), (p, 0, p), (p, p, 0)\}$ for $2p \in \{-N - 1, -N - 2\}$;
  - $\{(0, 0, q)\}, \{(0, q, 0)\}$, and $\{(q, 0, 0)\}$ for $q \leq -N - 1$.

Loosely speaking, then, a grove consists of a fixed set of edges outside the region $\{(i, j, k) \in \mathcal{I} \mid i + j + k \geq -N\}$, together with a graph inside this region constrained by connectivity conditions among the vertices near the region's boundary. Figure 2 illustrates this structure (here for $N = 4$). The squares illustrated are necessarily present if $N = 4$ is a cutoff; the straight black lines are the short edges forced by compactness, and the wavy lines connect vertices that must belong to the same component in any grove. Figure 3 shows an example of an actual grove within radius 4 on the initial set obtained by taking $\mathcal{L} = \{(i, j, k) \in C(0, 0, 0) \mid i + j + k \leq -4\}$; the thick black lines are just the edges of G.

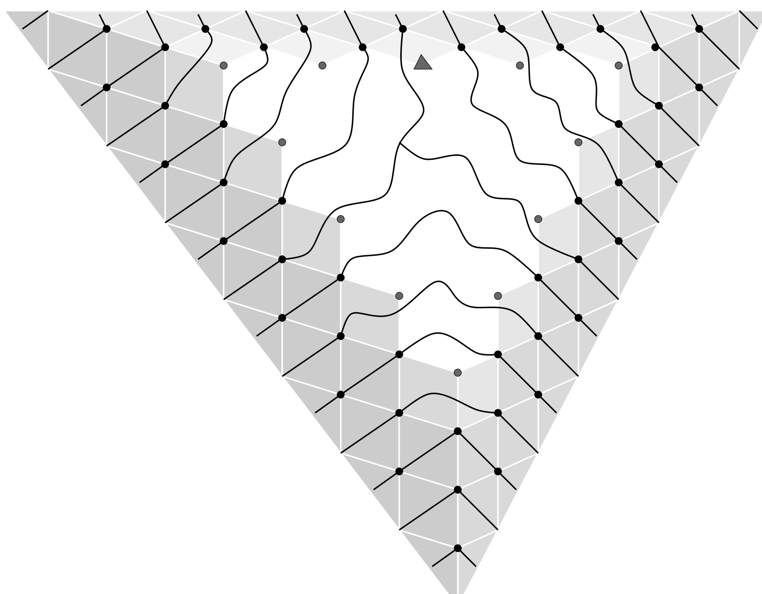

Figure 2: Compactness and connectivity conditions, $N = 4$; anchor is $(-2, -2, 0)$

We can reinterpret the compactness condition to specify exactly which edges of G meet the outer vertices:

Lemma 2 *Let G be an $\mathcal{I}$-grove within radius N. If $i \leq -N - 2$, then $(i, 0, 0)$ lies on the short edges $e'_b(i, 0, 0), e'_c(i, 0, 0)$ and no other edges; if $j, k < 0$ and $j + k \leq -N - 3$,*

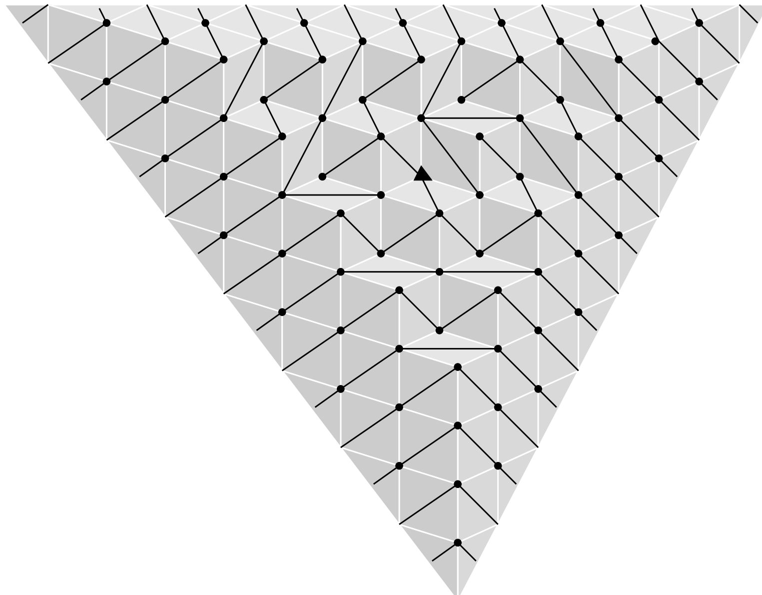

Figure 3: An example of a grove; anchor is $(-2,-2,-2)$

then $(0,j,k)$ *lies on the short edges* $e'_a(0,j+1,k+1), e'_a(0,j,k)$ *and no other edges, and similarly for the other permutations of coordinates.*

Proof: This follows immediately from Lemma 1 and the compactness condition. (See e.g. Figure 2 for an illustration.) □

We will show that the set of $\mathcal{I}$-groves within radius N actually does not depend on the choice of N (as long as N is a cutoff for $\mathcal{I}$). Although it is fairly straightforward to prove this directly from the definition, we will not take the trouble to do so here; instead, we will obtain it as a consequence of our main theorem.

In order to state the theorem, it will be necessary to indicate how groves can be represented algebraically. Given a grove G within radius N, define the corresponding Laurent monomial

$$m(G) = \left(\prod_{e_a(i,j,k) \in E(G)} a_{j,k} \right) \left(\prod_{e_b(i,j,k) \in E(G)} b_{i,k} \right) \left(\prod_{e_c(i,j,k) \in E(G)} c_{i,j} \right) \left(\prod_{(i,j,k) \in \mathcal{I}} x_{i,j,k}^{\deg(i,j,k)-2} \right).$$

The first three products are finite because they are simply products of edge variables corresponding to long edges in G, and the compactness condition ensures that only finitely many long edges appear. (These products elucidate our use of the term "edge variable.") The last product is finite because there are only finitely many (i,j,k) with $i+j+k > -N-3$, and Lemma 2 ensures that all the $(i,j,k) \in \mathcal{I}$ with $i+j+k \le -N-3$ have degree 2, so that the $x_{i,j,k}$ corresponding to these vertices cannot appear in the product (the lemma requires $\max\{i,j,k\} = 0$, but this holds since N is a cutoff).

Notice also that $m(G)$ uniquely determines G (independently of N), since it states precisely which long edges occur in G. To see this, we need only observe that no two distinct long edges can be represented by the same edge variable. For example, if $a_{j,k}$ represented two edges $e_a(i,j,k)$ and $e_a(i',j,k)$ with $i > i'$, then $(i,j,k) \in s_a(i,j,k) \subseteq \mathcal{I}$, so $(i'+1,j,k) \in \mathcal{L}$ and $(i',j-1,k-1) \notin \mathcal{I}$, contradicting $(i',j-1,k-1) \in s_a(i',j,k) \subseteq \mathcal{I}$.

At this point we are prepared to state our main theorem.

Theorem 1 *Let $\mathcal{L}$,$\mathcal{U}$,$\mathcal{I}$ be as described in Section 2, including the assumptions $\mathcal{L} \subseteq C(0,0,0)$ and $\mathcal{U} \cap C(0,0,0)$ finite. Define $f_{i,j,k}$ as in Section 2, and let N be any cutoff for $\mathcal{I}$. Then*

$$f_{0,0,0} = \sum_G m(G),$$

where the sum is taken over all $\mathcal{I}$-groves within radius N.

The proof is postponed while we discuss other properties of groves. We point out now, however, that there is only one way of decomposing $f_{0,0,0}$ as a sum of Laurent monomials in the variables $\{a_{j,k}, b_{i,k}, c_{i,j}, x_{i,j,k}\}$. Because each monomial $m(G)$ in turn determines G uniquely, we see that we obtain the same set of radius-N groves regardless of the choice of cutoff N. Consequently, in all subsequent discussion (except the proof of Theorem 1 itself), we may drop the N and simply use the term "grove," or "$\mathcal{I}$-grove" if the choice of $\mathcal{I}$ is ambiguous.

Before proceeding, we make one basic observation about the structure of groves.

Theorem 2 *Every grove is acyclic.*

The proof will use one preliminary result. Let $\mathcal{J} = \{(i,j,k) \in \mathcal{I} \mid i+j+k \geq -N-2\}$. We then have

Lemma 3 *The set $\mathcal{J}$ consists of $3\binom{N+3}{2}+1$ points and contains exactly $3\binom{N+2}{2}$ squares.*

The proof is again deferred; it will provide practice for the proof of Theorem 1.

Proof of Theorem 2: Let G be a grove on the initial set $\mathcal{I}$; we begin by proving that H, the induced subgraph on $\mathcal{J}$, is acyclic. We claim that if any two vertices of $\mathcal{J}$ lie in the same component of G, they are connected by a path contained in $\mathcal{J}$. For suppose not; choose two vertices connected by a path (which we may assume to have no repeated vertices) not contained in $\mathcal{J}$. Choose a vertex (i,j,k) of this path with $i+j+k$ minimal; then we must have $i+j+k < -N-2$, and (i,j,k) is not an endpoint of the path. Because N is a cutoff, $\max\{i,j,k\} = 0$; assume without loss of generality $i = 0$ and $j \geq k$. By Lemma 2, if $j < 0$ then the only edges of G meeting this vertex are $e'_a(0,j+1,k+1)$ and $e'_a(0,j,k)$, and the second of these cannot be used in the path (since its endpoint $(0,j-1,k-1)$ would violate minimality). Similarly, if $j = 0$ then the only edges of G at this vertex are $e'_a(0,0,k), e'_b(0,0,k)$, both of which would violate minimality. Hence, at most one edge incident to $(0,j,k)$ may be used in the path, contradicting the assumption of no repeated vertices. The claim follows. Hence, every component of G that contains any vertex of $\mathcal{J}$ induces a single component of H.

Now consider the following classes of vertices in $\mathcal{J}$:

- $\{(0,p,q),(p,0,q)\}, \{(p,q,0),(0,q,p)\}$, and $\{(q,0,p),(q,p,0)\}$ for all p,q with $0 > p > q$ and $p+q \in \{-N-1,-N-2\}$;
- $\{(0,p,p),(p,0,p),(p,p,0)\}$ for $2p \in \{-N-1,-N-2\}$;
- $\{(0,0,q)\}, \{(0,q,0)\}$, and $\{(q,0,0)\}$ for $q \in \{-N-1,-N-2\}$.

By the foregoing and connectivity for G, each class is contained in a single component of H, and certainly no component may contain vertices of more than one class (otherwise the corresponding component of G would, which is impossible). We also claim that every component of H contains one of the above classes of vertices; it suffices to show that the corresponding component of G contains some vertex (i,j,k) with $i+j+k \in \{-N-1,-N-2\}$. Suppose not. The connectivity condition specifies various sets of points (i,j,k), each satisfying $i+j+k \leq -N-1$, such that each component of G contains one of these sets. Hence, our particular component of G contains such a point (i,j,k), and since we have supposed $i+j+k \neq -N-1,-N-2$, we have $i+j+k < -N-2$. But it is impossible for such a vertex to be connected to a vertex of $\mathcal{J}$ by a path not going through any point with sum of coordinates in $\{-N-1,-N-2\}$, since the sum of coordinates changes by at most two at each step along the path. This is a contradiction.

We therefore conclude that the components of H are in bijection with our classes of vertices, of which there are $3N+7$. On the other hand, by Lemma 3 (and the complementarity condition), H has $3\binom{N+3}{2}+1$ vertices and $3\binom{N+2}{2}$ edges, for a minimum of

$$\left(3\binom{N+3}{2}+1\right)-3\binom{N+2}{2}=3N+7$$

components, with equality only if H is acyclic. Equality does hold, so H is acyclic, as claimed.

Now, suppose G contains some cycle. Since the definition of a grove is independent of the choice of cutoff N, we may choose N large enough so that all the vertices of the cycle belong to $\mathcal{J}$. Then the induced subgraph H contains a cycle, and this is a contradiction. Hence, G is acyclic. □

By projecting onto a plane of the form $\alpha i+\beta j+\gamma k=0$, for $\alpha,\beta,\gamma>0$, we can reinterpret groves as graphs on the lattice-like vertex set that we obtain as the projection of $\mathcal{I}$. One can show that the interiors of distinct parallelograms (the convex hulls of projections of squares) cannot overlap; this fact, together with the complementarity requirement, ensures that the resulting graphs are planar. Accepting this, it is possible to give a more intuitive proof of Theorem 2 than the one we have presented above, along the following lines: if a grove G contained a cycle, then in the planar projection, this cycle would enclose some vertices of G; planarity would then prevent the enclosed vertices from being connected to any of the boundary vertices involved in the connectivity condition, giving a contradiction. However, we have not found a rigorous and self-contained way of fleshing out this argument that is more succinct than the counting proof provided here.

The definition of a grove we have presented is somewhat cumbersome. For purposes of empirical investigation, infinite graphs are inconvenient to work with; groves as defined above also contain, in a certain sense, redundant information. We therefore will present — in slightly less detail — a simplification which we have at times found intuitively more useful, in the hope that it will also prove more practical for later investigations.

We can define a point $(i, j, k) \in \mathbb{Z}^3$ to be *even* or *odd* depending on the parity of $i + j + k$. Every square then has two even vertices and two odd vertices, and the edge used in G connects two vertices of the same parity (and so can also be called *even* or *odd*). It follows that, given "half" of a grove, we can uniquely reconstruct the other half: if we know which even edges are used in G, the remaining edges must be precisely the odd edges of those squares whose even edges are not used.

Finally, we can restrict our attention to a finite piece of $\mathcal{I}$, since, whenever a vertex (i, j, k) satisfies $i + j + k \leq -N - 3$, we know precisely which edges are incident to it, by Lemma 2. Putting all these observations together, we define a *simplified grove within radius* N, where N is a cutoff for $\mathcal{I}$ and furthermore is odd, to be a subgraph G' of $\mathcal{G}$ satisfying:

- (Vertex set) the vertex set of G' is $\{(i, j, k) \in \mathcal{I} \mid i + j + k \equiv 0 \bmod 2;\ i + j + k \geq -N - 1\}$;
- (Acyclicity) G' is acyclic;
- (Connectivity) the *even boundary vertices* $\{(i, j, k) \in \mathcal{I} \mid i + j + k = -N - 1;\ \max\{i, j, k\} = 0\}$ can be partitioned into the following sets so that each component of G' contains exactly one set, and conversely, each set is contained in some component:
  - $\{(0, p, q), (p, 0, q)\}, \{(p, q, 0), (0, q, p)\}$, and $\{(q, p, 0), (q, 0, p)\}$ for $0 > p > q$, $p + q = -N - 1$;
  - $\{(0, \frac{-N-1}{2}, \frac{-N-1}{2}), (\frac{-N-1}{2}, 0, \frac{-N-1}{2}), (\frac{-N-1}{2}, \frac{-N-1}{2}, 0)\}$,
  - $\{(0, 0, -N - 1)\}, \{(0, -N - 1, 0)\}$, and $\{(-N - 1, 0, 0)\}$.

For any grove G, the induced subgraph on $\{(i, j, k) \mid i+j+k \equiv 0 \bmod 2;\ i+j+k \geq -N-1\}$ is then a simplified grove. (An example of a grove and the corresponding simplified grove, with $N = 3$, is shown in the top two panels of Figure 4; the vertices circled in the left panel are those belonging to the vertex set of the simplified grove.) The only condition that is not entirely evident is connectivity, which follows from the connectivity condition on G, with a slight bit of subtlety: we must check that any path in the grove G connecting two vertices belonging to G' is in fact entirely contained in the vertex set of G'. However, the requirement that all vertices on the path be even poses no difficulty, as even vertices are connected only to even vertices; and the verification that all vertices (i, j, k) on the path satisfy $i + j + k \geq -N - 1$ is precisely as in the proof of Theorem 2.

We next show the converse of the above — every simplified grove is induced by a unique grove. Given a simplified grove G', we again let $\mathcal{J} = \{(i, j, k) \in \mathcal{I} \mid i + j + k \geq -N - 2\}$. We extend G' to a graph G on all of $\mathcal{I}$ as follows: for every square contained in $\mathcal{J}$ whose even edge does not appear in G', we include the odd edge instead; we then include the

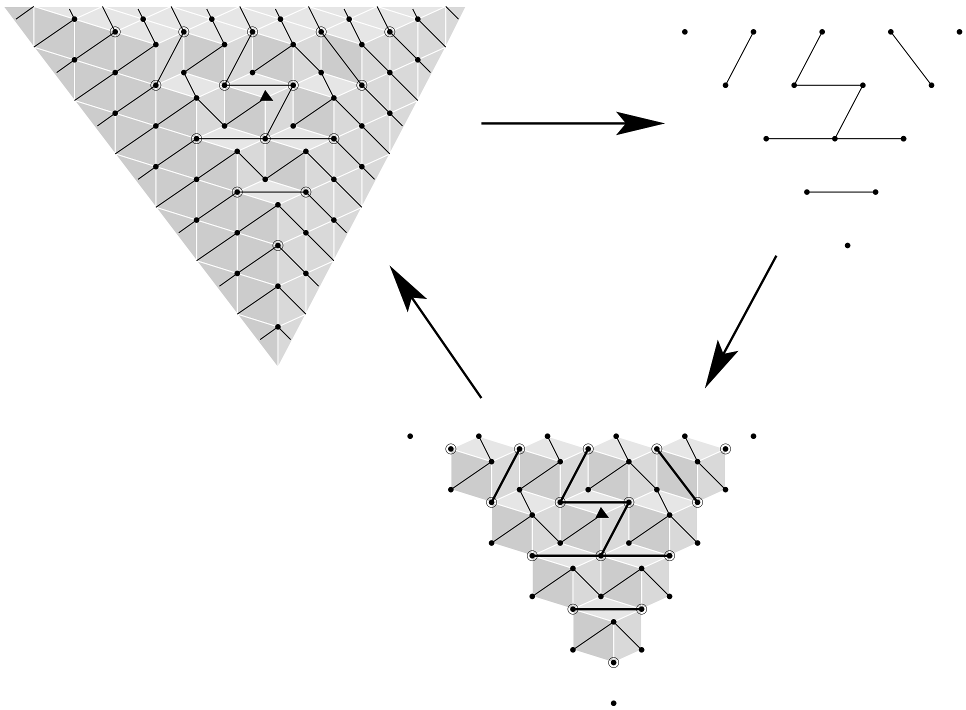

Figure 4: Obtaining a simplified grove from a grove; anchor is $(-1,-1,-1)$

short edges of all other squares in $\mathcal{I}$. In view of the complementarity and compactness conditions, the resulting graph is the only possible grove on $\mathcal{I}$ that induces the simplified grove G'. (This process is also shown in Figure 4, following the two diagonal arrows; the lower panel shows the intermediate graph on $\mathcal{J}$, with the edges belonging to the simplified grove thickened.) We claim that this G is indeed a grove. All the conditions except connectivity are immediate; this last is a bit more involved. To make the argument clearer, Figure 5 compares the connectivity conditions for G' and G, here for $N = 3$. The solid wavy lines represent paths between even boundary vertices, as specified by simplified-grove connectivity; the *odd boundary vertices* $\{(i,j,k) \mid i+j+k = -N-2;\ \max\{i,j,k\} = 0\}$ are circled, and the dashed wavy lines show the connectivities between them that are required for a grove. (Compare to Figure 2, although there we take $N = 4$.)

To prove the connectivity condition for G, we first show that our graph H on the vertex set $\mathcal{J}$, obtained from G' by including the appropriate odd edges, is acyclic. It cannot have any cycles on even vertices, as such a cycle would already have existed in G'. If there is a cycle on the odd vertices, then it is not hard to see that its planar projection must enclose some even vertices of $\mathcal{J}$ but cannot enclose any even boundary vertices. By planarity, then, G' has some vertices that cannot be connected to any even boundary vertices, violating its connectivity requirement. Thus, no cycles have been introduced in $\mathcal{J}$.

Now we again apply the component-counting technique. From connectivity for a simplified grove, we can check that G' always has $(3N+5)/2$ components. If we consider the graph H, Lemma 3 tells us that we have $3\binom{N+2}{2}$ edges and $3\binom{N+3}{2}+1$ vertices; since the graph is acyclic, it has

$$\left(3\binom{N+3}{2}+1\right)-3\binom{N+2}{2}=3N+7$$

components. In particular, there are $(3N+7)-(3N+5)/2=(3N+9)/2$ components on the odd vertices.

Now divide the odd boundary vertices into classes according to connectivity for a grove: we have the classes $\{(0,p,q),(p,0,q)\}$, $\{(p,q,0),(0,q,p)\}$, $\{(q,p,0),(q,0,p)\}$ for $0>p>q, p+q=-N-2$, as well as $\{(0,0,-N-2)\},\{(0,-N-2,0)\},\{(-N-2,0,0)\}$. We thus get $(3N+9)/2$ classes. We wish to show that, in our graph G on $\mathcal{I}$, no two odd boundary vertices from different classes can lie in the same component. To see this, we consider the paths of G' connecting the even boundary vertices, together with the extra short edges emanating from them (given by the compactness condition), and project into the plane. These paths divide the plane into sectors, each of which contains the odd boundary vertices of (at most) one class, as in Figure 5. If any two odd boundary vertices from different classes lie in the same component of G, the path connecting them must cross one of the paths on even vertices. By planarity, this can only happen if the intersection happens at an actual vertex of the graph. However, one path uses only even vertices and the other uses only odd vertices — contradiction.

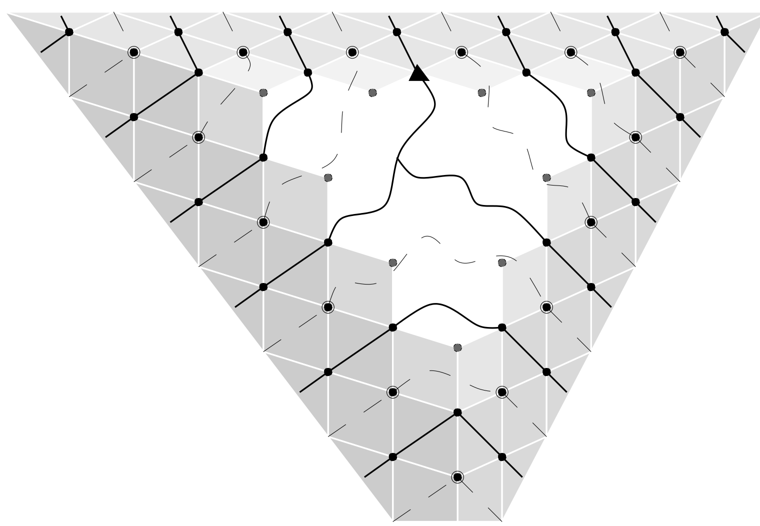

Figure 5: Connectivity restrictions among odd boundary vertices; anchor is $(-2,-2,0)$

We have shown that, for any component of G — and so for any component of H — the odd boundary vertices it contains all belong to the same class. Since we have $(3N+9)/2$ classes, we get at least $(3N+9)/2$ components of H on odd vertices. But we have already shown that equality occurs. This is only possible if, for each class, all its vertices belong to the same component of H, and every component of H contains one such class of vertices. We may now conclude that each class of vertices required for a grove is contained in some component of G: either the class consists of odd boundary vertices, and we have just proven that the vertices in the class are connected; it consists of even boundary vertices, and the conclusion follows from connectivity for G'; or it consists of a single vertex outside of $\mathcal{J}$, and the conclusion is trivial. Conversely, consider any component of G. If it contains any vertex of $\mathcal{J}$, it either contains odd vertices, and then (as we have just seen) it contains an entire class of odd boundary vertices, or it contains an even vertex, and then (by connectivity for G') it contains an entire class of even boundary vertices. Moreover, no two boundary vertices from different classes may be connected within $\mathcal{J}$, and the extra short edges cannot introduce any new such connectivities (by the usual minimal-point argument from the proof of Theorem 2); thus, no component of G contains more than one class of boundary vertices. And if our component of G does not contain any vertex of $\mathcal{J}$, then it is a chain of short edges, containing exactly one vertex of the form $(0,0,q)$, $(0,q,0)$, or $(q,0,0)$ ($q \leq -N-3$), and no vertices of $\mathcal{J}$ — in particular, none of the classes of boundary vertices. Thus, each component of G contains exactly one class of vertices. We now have verified that the grove connectivity condition is met in its entirety.

This completes the proof that every simplified grove is induced by a unique grove. We now have established a bijection between groves and simplified groves within radius N, for any odd cutoff N, so the latter may be used as a sort of shorthand to represent the former. When N is an even cutoff, we can make a completely analogous definition: a simplified grove within radius N is a subgraph G' of $\mathcal{G}$ satisfying

- (Vertex set) the vertex set of G' is $\{(i,j,k) \in \mathcal{I} \mid i+j+k \equiv 1 \bmod 2;\ i+j+k \geq -N-1\}$;
- (Acyclicity) G' is acyclic;
- (Connectivity) the *odd boundary vertices*, $\{(i,j,k) \in \mathcal{I} \mid i+j+k = -N-1;\ \max\{i,j,k\} = 0\}$, can be partitioned into the following sets so that each component of G' contains exactly one set, and each set is contained in some component:
  - $\{(0,p,q),(p,0,q)\}, \{(p,q,0),(0,q,p)\}, \{(q,p,0),(q,0,p)\}$ for $0 > p > q$, $p+q = -N-1$;
  - $\{(0,0,-N-1)\}$, $\{(0,-N-1,0)\}$, and $\{(-N-1,0,0)\}$.

The simplified groves within radius N are again in bijection with the groves within radius N; the proof is completely analogous to the one given for the case of N odd.

4 Some consequences

Before proceeding to the proof of the main theorem, let us continue to assume that it holds and note some consequences. As usual, the hypotheses of Section 2 hold, including the assumptions that $\mathcal{L} \subseteq C(0,0,0)$ and $|\mathcal{U} \cap C(0,0,0)| < \infty$, except where stated otherwise.

It is clear that, as a Laurent polynomial in $\{a_{j,k}, b_{i,k}, c_{i,j}, x_{i,j,k}\}$, $f_{0,0,0}$ has every coefficient equal to 1 — that is, a grove G is uniquely determined by the monomial $m(G)$, as already discussed. In [8], Propp conjectured that, for the initial set $\mathcal{I} = \{(i,j,k) \in \mathbb{Z}^3 \mid -1 \leq i+j+k \leq 1\}$, every coefficient is still 1 if we set all edge variables equal to 1 and simply view each $f_{i,j,k}$ as a Laurent polynomial in the variables $x_{i',j',k'}$. Fomin and Zelevinsky conjectured (in [4]) that the coefficients remain at least nonnegative for general initial sets. We will show that the coefficients are all 1 when all edge variables are set to 1 and the initial set is arbitrary. This is equivalent to the assertion that any grove G (on a given initial set) can be uniquely reconstructed from the degrees of its vertices.

More specifically, we will prove:

Theorem 3 *Suppose $s_a(i_0, j_0, k_0) \subseteq \mathcal{I}$ is a square. Then, for any grove G,*

$$\sum_{\substack{(i,j,k)\in\mathcal{I} \\ j<j_0,\ k<k_0}} (\deg(i,j,k) - 2)$$

equals -1 if the long edge $e_a(i_0, j_0, k_0)$ appears in G, and 0 if it does not appear.

(We have already noted that the sum has only finitely many nonzero terms.) This theorem, and the analogous statements for the other two types of squares, will imply Propp's conjecture for any $\mathcal{I}$.

Proof: Let t be a new formal indeterminate. We apply the following variable substitutions:

$$a_{j_0,k_0} \leftarrow t; \qquad a_{j,k} \leftarrow 1 \text{ otherwise};$$

$$b_{i,k}, c_{i,j} \leftarrow 1;$$

$$x_{i,j,k} \leftarrow t \quad (j < j_0, k < k_0); \qquad x_{i,j,k} \leftarrow 1 \text{ otherwise.}$$

We proceed to compute values of $f_{i,j,k}$ $(i,j,k \in \mathcal{I} \cup \mathcal{U})$ by the recurrence (1); each such value will then be a Laurent polynomial in t, by Theorem 1. We claim that, in fact, $f_{i,j,k}$ is a constant multiple of t for $j < j_0, k < k_0$, and is simply a constant otherwise. (Notice also that these constants are always positive.)

The proof is by induction on the cardinality of $C(i,j,k) \cap \mathcal{U}$. If this cardinality is 0, we have $(i,j,k) \in \mathcal{I}$, and $f_{i,j,k} = x_{i,j,k} = t$ or 1 according to whether or not $j < j_0$ and $k < k_0$. Otherwise, we split into cases:

- If $j < j_0, k < k_0$, then (1) simply says that

$$f_{i,j,k} = (f_{i-1,j,k}f_{i,j-1,k-1} + f_{i,j-1,k}f_{i-1,j,k-1} + f_{i,j,k-1}f_{i-1,j-1,k})/f_{i-1,j-1,k-1}. \qquad (2)$$

  However, either $(i-1,j,k) \in \mathcal{I}$, or $(i-1,j,k) \in \mathcal{U}$ and $|C(i-1,j,k) \cap \mathcal{U}| < |C(i,j,k) \cap \mathcal{U}|$ (this was discussed at the beginning of Section 2), so the induction hypothesis tells us that $f_{i-1,j,k}$ is a constant multiple of t. Similarly, $f_{i,j-1,k-1}, f_{i,j-1,k}$, and so forth are all constant multiples of t, and we conclude that $f_{i,j,k}$ is as well.

- If $j = j_0, k < k_0$, then (1) again reduces to (2) The induction hypothesis now tells us that $f_{i-1,j,k}, f_{i-1,j,k-1}, f_{i,j,k-1}$ are constants, while $f_{i,j-1,k-1}, f_{i,j-1,k}, f_{i-1,j-1,k}$, and $f_{i-1,j-1,k-1}$ are constant multiples of t, and it follows that $f_{i,j,k}$ is also a constant.

- If $j < j_0, k = k_0$, the reasoning is the same as in the previous case.

- If $j = j_0, k = k_0$, then (1) takes the form

$$f_{i,j,k} = (f_{i-1,j,k}f_{i,j-1,k-1} + tf_{i,j-1,k}f_{i-1,j,k-1} + tf_{i,j,k-1}f_{i-1,j-1,k})/f_{i-1,j-1,k-1}.$$

  The induction hypothesis now tells us that $f_{i-1,j,k}, f_{i,j-1,k}, f_{i,j,k-1}, f_{i-1,j,k-1}$, and $f_{i-1,j-1,k}$ are constants, while $f_{i,j-1,k-1}, f_{i-1,j-1,k-1}$ are constant multiples of t. We conclude that $f_{i,j,k}$ is again a constant.

- Finally, if $j > j_0$ or $k > k_0$, then we have (2) once again. All the f's appearing on the right are now constants, so $f_{i,j,k}$ is as well.

This completes the induction. In particular, we know now that $f_{0,0,0}$ is a constant polynomial (because $j_0, k_0 \leq 0$). So, for every grove G, $m(G)$ becomes a constant under our variable substitutions — the total exponent of t is 0. This means that either the sum in the problem statement is -1 and the long edge $e_a(i_0, j_0, k_0)$ (corresponding to a_{j_0,k_0}) does appear, or the sum is 0 and the long edge does not appear, as claimed. □

Another conjecture appearing in [8] is that, in each term of any polynomial generated by the cube recurrence, every $x_{i,j,k}$ has its exponent in the range $\{-1, 0, \dots, 4\}$. In terms of groves, this is equivalent to the statement that every vertex has degree no less than 1 and no more than 6. The lower bound is obvious, since the connectivity condition ensures that there are no isolated vertices (except possibly those of the forms $(0, 0, q), (0, q, 0), (q, 0, 0)$, but these lie on short edges given by the compactness condition). The upper bound holds because each vertex has at most six neighbors in $\mathcal{G}$; this essentially follows from the fact that, in the parallelogram tiling of the plane corresponding to $\mathcal{I}$, no vertex can belong to more than six parallelograms. To make this argument precise, we show that any $(i, j, k) \in \mathcal{I}$ belongs to at most six squares; this is sufficient, since each square contributes only one neighbor to (i, j, k) in $\mathcal{G}$. A priori, (i, j, k) may belong to up to twelve squares: $s_a(i, j, k), s_a(i, j + 1, k), s_a(i, j, k + 1), s_a(i, j + 1, k + 1)$, and similarly for the other two orientations. However, we arrange these squares into pairs such that, in each pair, only one square can actually occur in $\mathcal{I}$. For example, $s_b(i+1, j, k+1)$ and $s_c(i+1, j, k)$ cannot both occur, since these would require $(i + 1, j, k + 1)$ and $(i, j - 1, k)$ (respectively) to be both in $\mathcal{I}$, a contradiction. The full pairing is shown in Table 1.

This shows that every vertex has at most 6 neighbors in $\mathcal{G}$ (and so in G), as claimed.

The preceding results were originally phrased in [8] as algebraic statements about the cube recurrence, but they can now be interpreted as geometric statements about the structure of groves. Another geometric fact worth noting concerns the distribution of the orientations of long edges in any grove.

Theorem 4 *Suppose that a grove G has n_a long edges of the form $e_a(i, j, k)$, n_b of the form $e_b(i, j, k)$, and n_c of the form $e_c(i, j, k)$. Then the numbers $n_a + n_b - n_c, n_b + n_c - n_a, n_c + n_a - n_b$ are all nonnegative and even.*

If (i,j,k) belongs to these squares	we get this contradiction in $\mathcal{I}$
$s_b(i+1,j,k+1), s_c(i+1,j,k)$	$(i+1,j,k+1),(i,j-1,k)$
$s_a(i,j,k), s_b(i+1,j,k)$	$(i+1,j,k),(i,j-1,k-1)$
$s_c(i+1,j+1,k), s_a(i,j+1,k)$	$(i+1,j+1,k),(i,j,k-1)$
$s_b(i,j,k), s_c(i,j+1,k)$	$(i,j+1,k),(i-1,j,k-1)$
$s_a(i,j+1,k+1), s_b(i,j,k+1)$	$(i,j+1,k+1),(i-1,j,k)$
$s_c(i,j,k), s_a(i,j,k+1)$	$(i,j,k+1),(i-1,j-1,k)$

Table 1: Proof that every vertex has degree ≤ 6

Proof: The proof is similar to that used for Theorem 3. We view n_a, n_b, n_c as functions of G; it suffices to prove the result for $n_a(G)+n_b(G)-n_c(G)$, as the other cases are analogous. We let t be a formal indeterminate and make the variable substitutions

$$\begin{aligned} a_{j,k}, b_{i,k} &\leftarrow t; \\ c_{i,j} &\leftarrow 1/t; \\ x_{i,j,k} &\leftarrow 1. \end{aligned}$$

We again compute values of $f_{i,j,k}$ by recurrence (1); we know from Theorem 1 that each $f_{i,j,k}$ is a Laurent polynomial in t with positive coefficients. We claim that $f_{i,j,k}$ is in fact a polynomial in t^2 whose constant term is positive.

The proof is by induction on $|C(i,j,k)\cap\mathcal{U}|$, as usual. If this cardinality is 0, then $(i,j,k)\in\mathcal{I}$, so $f_{i,j,k}=x_{i,j,k}=1$, and the result holds. Otherwise, the recurrence (1) tells us that

$$f_{i,j,k}f_{i-1,j-1,k-1} = f_{i-1,j,k}f_{i,j-1,k-1} + f_{i,j-1,k}f_{i-1,j,k-1} + t^2 f_{i,j,k-1}f_{i-1,j-1,k}.$$

The induction hypothesis ensures that the right side is a polynomial in t^2 with positive constant term (since each factor is). We also know that $f_{i-1,j-1,k-1}$ is a polynomial in t^2 with positive constant term. Let d be the minimal exponent for which the coefficient of t^d in $f_{i,j,k}$ is nonzero; then the lowest power of t appearing on the right side must also be t^d, so $d=0$. This assures us that $f_{i,j,k}$ is a polynomial in t with positive constant term. However, since $f_{i-1,j,k}$ and so forth are actually polynomials in t^2, $f_{i,j,k}$ is a rational function of t^2; hence, it is a polynomial in t^2, and the induction step holds.

In particular, $\sum_G m(G) = f_{0,0,0}$ is a polynomial in t^2. On the other hand, for each grove G, it is apparent that $m(G)=t^{n_a(G)+n_b(G)-n_c(G)}$. It follows that $n_a(G)+n_b(G)-n_c(G)$ is nonnegative and even. $\square$

This observation allows us to provide a complete combinatorial proof of the Laurent property of the cube recurrence as stated (and proved, using techniques developed in the theory of cluster algebras) by Fomin and Zelevinsky in [4]. Their version is essentially as follows:

Theorem 5 *Let $\mathcal{L}\subseteq\mathbb{Z}^3$ such that whenever $(i,j,k)\in\mathcal{L}$, $C(i,j,k)\subseteq\mathcal{L}$. Let $\mathcal{U}=\mathbb{Z}^3-\mathcal{L}, \mathcal{I}=\{(i,j,k)\in\mathcal{L}\mid(i+1,j+1,k+1)\in\mathcal{U}\}$, and suppose $C(i,j,k)\cap\mathcal{U}$ is finite* *for each* $(i,j,k) \in \mathcal{U}$*. Also, let* α, β, γ *be formal indeterminates. Define* $f_{i,j,k} = x_{i,j,k}$ *for* $(i,j,k) \in \mathcal{I}$ *and*

$$f_{i,j,k} = \frac{\alpha f_{i-1,j,k} f_{i,j-1,k-1} + \beta f_{i,j-1,k} f_{i-1,j,k-1} + \gamma f_{i,j,k-1} f_{i-1,j-1,k}}{f_{i-1,j-1,k-1}}$$

recursively for $(i,j,k) \in \mathcal{U}$*. Then each* $f_{i,j,k}$ *is a Laurent polynomial in the* $x_{i,j,k}$ *with coefficients in* $\mathbb{Z}[\alpha, \beta, \gamma]$*.*

Proof: By our usual translation and intersection argument, we may replace $\mathcal{L}$ by $\mathcal{L} \cap C(0,0,0)$ and define $f_{i,j,k}$ only for $(i,j,k) \in C(0,0,0)$, and it suffices to show that $f_{0,0,0}$ is a Laurent polynomial in the $x_{i,j,k}$ with coefficients in $\mathbb{Z}[\alpha,\beta,\gamma]$. Working over $\mathbb{Q}(\sqrt{\alpha}, \sqrt{\beta}, \sqrt{\gamma})$, let $a_{j,k} = \sqrt{\beta\gamma/\alpha}, b_{i,k} = \sqrt{\gamma\alpha/\beta}, c_{i,j} = \sqrt{\alpha\beta/\gamma}$ for all i,j,k; the recurrence (1) then assumes the desired form. By Theorem 1, we know that

$$f_{0,0,0} = \sum_G \sqrt{\beta\gamma/\alpha}^{\,n_a(G)} \sqrt{\gamma\alpha/\beta}^{\,n_b(G)} \sqrt{\alpha\beta/\gamma}^{\,n_c(G)} \cdot X(G)$$

(where $X(G)$ is some Laurent monomial in the variables $x_{i,j,k}$, with coefficient 1)

$$= \sum_G \alpha^{\frac{n_b(G)+n_c(G)-n_a(G)}{2}} \beta^{\frac{n_c(G)+n_a(G)-n_b(G)}{2}} \gamma^{\frac{n_a(G)+n_b(G)-n_c(G)}{2}} \cdot X(G).$$

The result now follows from Theorem 4. □

So far we have been studying the cube recurrence in general, but for purposes of concreteness it is helpful to study the recurrence in the context of specific initial sets. The question that Propp originally asked in [8] was, given the initial conditions $f_{i,j,k} = x_{i,j,k}$ $(-1 \le i+j+k \le 1)$ and the recurrence without edge variables

$$f_{i,j,k} = \frac{f_{i-1,j,k} f_{i,j-1,k-1} + f_{i,j-1,k} f_{i-1,j,k-1} + f_{i,j,k-1} f_{i-1,j-1,k}}{f_{i-1,j-1,k-1}} \quad (i+j+k > 1),$$

how to interpret combinatorially the terms of the Laurent polynomial $f_{i,j,k}$ for $i+j+k = n > 1$. Propp observed, by setting every $x_{i',j',k'} = 1$ and using an easy induction, that there are $3^{\lfloor n^2/4 \rfloor}$ such terms; the question was then what relevant objects have the property that there are $3^{\lfloor n^2/4 \rfloor}$ objects of order n. We now have the means to answer this question. By translation and $C(0,0,0)$-intersection, the problem is equivalent to describing $f_{0,0,0}$ given the initial set obtained from $\mathcal{L} = \{(i,j,k) \in C(0,0,0) \mid i+j+k \le 1-n\}$. This $\mathcal{I}$ is

$$\{(i,j,k) \in C(0,0,0) \mid -1-n \le i+j+k \le 1-n; \text{ or } i+j+k < -1-n,\ \max\{i,j,k\} = 0\}.$$

Thus, the terms of the polynomial correspond to groves on this initial set, which we call *standard groves of order* n. A typical graph $\mathcal{G}$ for such an initial set is the one shown in Figure 1, there for $n = 5$, and the grove shown in Figure 3 is one of the $3^{\lfloor 5^2/4 \rfloor} = 729$ standard groves of order 5. Standard groves have a further convenient property: when

$N = n - 1$ is used as the cutoff, the simplified grove corresponding to a standard grove G of order n consists simply of the long edges of G, as the reader may check (Figure 4 provides an example).

Michael Kleber (cited in [4]) proposed using the initial conditions $f_{i,j,k} = x_{i,j,k}$ for points satisfying $\min\{i,j,k\} = 0$, and computing $f_{i,j,k}$ (again without edge variables) for $i,j,k > 0$. Under our usual series of manipulations, finding this $f_{i,j,k}$ is equivalent to finding $f_{0,0,0}$ with initial set induced by

$$\mathcal{L} = C(0,0,0) - \{(i',j',k') \mid i' > -i, j' > -j, k' > -k\}.$$

A typical graph $\mathcal{G}$ on this initial set is shown in Figure 6 (here for $(i,j,k) = (2,3,2)$), together with a grove on this initial set.

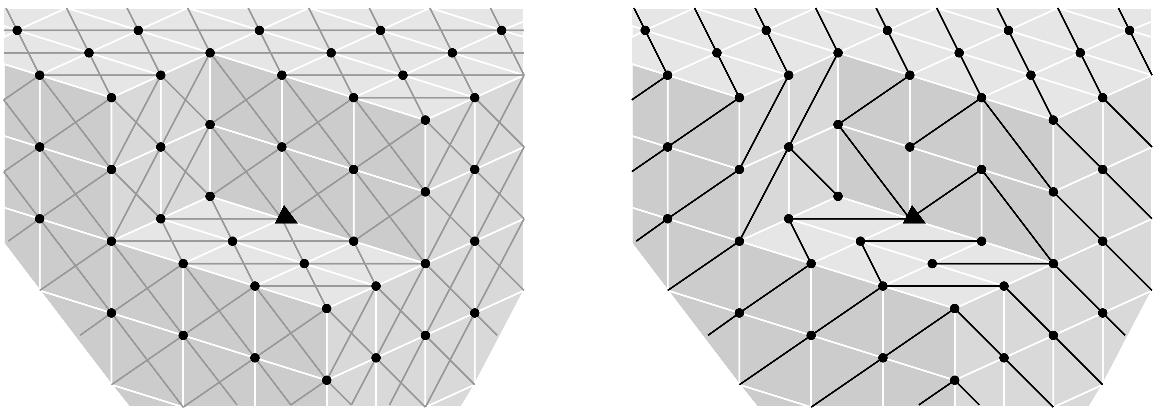

Figure 6: Michael Kleber's initial set and a sample grove; anchor is $(-2,-2,-2)$

Another interesting special case arises in connection with the Gale-Robinson Sequence Theorem, conjectured in [6] and proven by Fomin and Zelevinsky in [4] in the following form. (The reduction of the Gale-Robinson recurrence to the cube recurrence, also used in [4], was first suggested by Propp.)

Theorem 6 (Gale-Robinson Sequence Theorem) *Let p,q,r be positive integers, let $n = p + q + r$, and let α, β, γ be formal indeterminates. Suppose that the sequence $y_0, y_1, y_2, \ldots$ satisfies the Gale-Robinson recurrence:*

$$y_{l+n} = \frac{\alpha y_{l+p} y_{l+n-p} + \beta y_{l+q} y_{l+n-q} + \gamma y_{l+r} y_{l+n-r}}{y_l}$$

for $l \geq 0$. Then each y_l is given by a Laurent polynomial in the initial terms $y_0, \ldots, y_{n-1}$ with coefficients in $\mathbb{Z}[\alpha, \beta, \gamma]$.

Indeed, for fixed l, we can set $\mathcal{L} = \{(i,j,k) \in \mathbb{Z}^3 \mid pi + qj + rk < n - l\}$, yielding $\mathcal{I} = \{(i,j,k) \in \mathbb{Z}^3 \mid -l \leq pi + qj + rk < n - l\}$. After setting $x_{i,j,k} = y_{pi+qj+rk+l}$ ($(i,j,k) \in \mathcal{I}$) and applying the form of the recurrence stated in Theorem 5, we then obtain $f_{i,j,k} =$

$y_{pi+qj+rk+l}$ for $(i,j,k) \in \mathcal{U}$ by induction. The desired Laurentness of $y_l = f_{0,0,0}$ is then a direct consequence of Theorem 5. Moreover, if y_l is viewed as a Laurent polynomial in $y_0, \ldots, y_{n-1}, \alpha, \beta, \gamma$, then all the coefficients are nonnegative; this is immediate from Theorem 1 (since y_l is a sum of monomials, each with coefficient 1) and again addresses a conjecture in [4].

If we apply our usual $C(0,0,0)$-intersection, we obtain a combinatorial interpretation for the terms of any Gale-Robinson sequence in terms of groves. In particular, the Somos-6 and Somos-7 sequences described in the introduction (which are Gale-Robinson sequences with $(p,q,r) = (3,1,2)$, $(4,1,2)$ respectively; see [6] for further exposition) can be interpreted as simply counting groves on various initial sets, resulting in a new proof that the terms of these sequences are all integers. Figure 7 shows the graph $\mathcal{G}$ for the term y_{11} of the Somos-7 sequence.

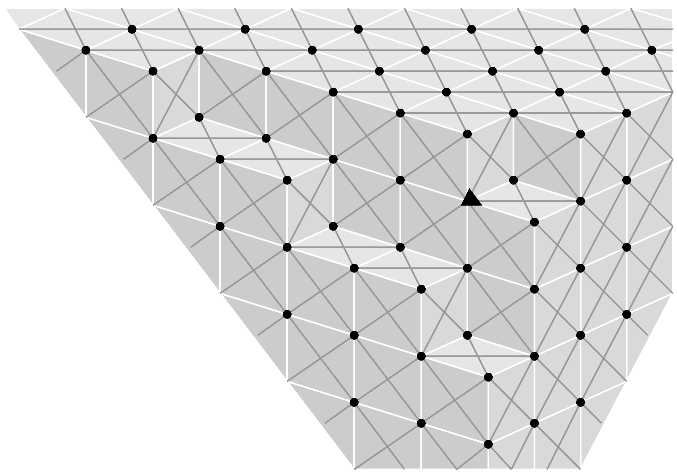

Figure 7: $\mathcal{G}$ for a Gale-Robinson sequence term; anchor is $(-1,-1,-1)$

One more specialization that merits investigation is obtained by taking an arbitrary initial set and setting $a_{j,k} = t, b_{i,k} = t, c_{i,j} = 1/t$, as in the proof of Theorem 4. Then the cube recurrence takes the form

$$f_{i,j,k} = \frac{f_{i-1,j,k}f_{i,j-1,k-1} + f_{i,j-1,k}f_{i-1,j,k-1} + t^2 f_{i,j,k-1}f_{i-1,j-1,k}}{f_{i-1,j-1,k-1}}.$$

In each monomial $m(G)$ coming from any polynomial $f_{i,j,k}$, the total exponent of t resulting from this substitution is $n_a + n_b - n_c$, where n_a, n_b, n_c are defined in terms of G as specified in Theorem 4. Since this quantity is always nonnegative, we conclude that t only appears to nonnegative powers in any $f_{i,j,k}$. Therefore, we can legitimately substitute $t = 0$, and the resulting Laurent polynomials (in the $x_{i,j,k}$) are given by the recurrence

$$f_{i,j,k} = \frac{f_{i-1,j,k}f_{i,j-1,k-1} + f_{i,j-1,k}f_{i-1,j,k-1}}{f_{i-1,j-1,k-1}}. \qquad (3)$$

The terms in $f_{0,0,0}$ then correspond to precisely those groves in which the total number of long edges of the forms $e_a(i,j,k)$ and $e_b(i,j,k)$ equals the number of long edges of the form $e_c(i,j,k)$. On the other hand, by setting $g_{x,y,z} = f_{(x+y-z)/2,(x-y+z)/2,(x-y-z)/2}$ for $(x,y,z) \in \mathbb{Z}^3, x+y+z \equiv 0 \bmod 2$, we obtain from (3) the recurrence

$$g_{x,y,z} = \frac{g_{x-1,y-1,z}g_{x-1,y+1,z} + g_{x-1,y,z-1}g_{x-1,y,z+1}}{g_{x-2,y,z}}.$$

This is the octahedron recurrence; as shown by Speyer in [10], the resulting Laurent polynomials can be interpreted as enumerating the perfect matchings of certain planar bipartite graphs, with the graphs determined by the shape of $\mathcal{I}$. For suitable initial sets, the corresponding graphs include the Aztec diamond graphs (whose perfect matchings correspond to the domino tilings of Aztec diamonds — see [7]) and, more generally, the pine-cone graphs of [1]. Consequently, we have a bijection between the set of perfect matchings of a graph (determined by $\mathcal{I}$) and a particular subset of the groves on the initial set $\mathcal{I}$. Addressing this correspondence here would take us rather afield, but we hope to discuss interesting consequences in a future paper.

5 The main proof

We now turn to the proof of Theorem 1 and, with it, that of Lemma 3. We maintain the assumptions $\mathcal{L} \subseteq C(0,0,0)$ and $|\mathcal{U} \cap C(0,0,0)| < \infty$. Our basic technique will be induction on $|\mathcal{U} \cap C(0,0,0)|$: we hold N fixed and observe what happens as the initial set varies. We begin with the observation that any initial set has a "local minimum."

Lemma 4 *Suppose that* $(0,0,0) \notin \mathcal{I}$. *Then there exist* $i,j,k \leq 0$ *such that* $(i-1,j,k)$, $(i,j-1,k), (i,j,k-1), (i,j-1,k-1), (i-1,j,k-1), (i-1,j-1,k), (i-1,j-1,k-1) \in \mathcal{I}$ *(and so* $(i,j,k) \in \mathcal{U}$*).*

Proof: By finiteness, we can choose $(i,j,k) \in \mathcal{U} \cap C(0,0,0)$ with $i+j+k$ minimal; we claim that these values of i,j,k suffice. For example, minimality ensures that $(i-1,j,k) \in \mathcal{L}$, while $(i,j+1,k+1) \notin \mathcal{L}$ (otherwise we would have $(i,j,k) \in \mathcal{L}$). Therefore, $(i-1,j,k) \in \mathcal{I}$. By similar reasoning, all of the other specified points lie in $\mathcal{I}$. $\square$

Our next task will be to prove Lemma 3. We first recall the statement.

Lemma 3 *Let* N *be a cutoff for* $\mathcal{I}$*, and let* $\mathcal{J} = \{(i,j,k) \in \mathcal{I} \mid i+j+k \geq -N-2\}$. *Then the set* $\mathcal{J}$ *consists of* $3\binom{N+3}{2}+1$ *points and contains exactly* $3\binom{N+2}{2}$ *squares.*

Proof: Fix N; we use induction on $|\mathcal{U} \cap C(0,0,0)|$. If this cardinality is 0, then $\mathcal{L} = C(0,0,0)$, $\mathcal{I} = \{(i,j,k) \in \mathbb{Z}^3 \mid \max\{i,j,k\} = 0\}$, and $\mathcal{J} = \{(i,j,k) \mid \max\{i,j,k\} = 0,\ i+j+k \geq -N-2\}$. Each point of the form $(0,j,k), j+k \geq -N$, gives rise to a square $s_a(0,j,k)$, and there are $\binom{N+2}{2}$ such points. Similarly, there are $\binom{N+2}{2}$ squares of each of the other two types, for a total of $3\binom{N+2}{2}$. Also, it is straightforward to verify that there are $3\binom{N+3}{2}+1$ points in $\mathcal{J}$.

Now suppose $\mathcal{U} \cap C(0,0,0) \neq \emptyset$, so that $(0,0,0) \notin \mathcal{I}$. Choose (i,j,k) as given by Lemma 4. Let $\mathcal{L}' = \mathcal{L} \cup \{(i,j,k)\}$; the lemma ensures that $\mathcal{L}'$ still meets the requirements we have imposed on $\mathcal{L}$. Then define $\mathcal{U}', \mathcal{I}', \mathcal{J}'$ by analogy with $\mathcal{U}, \mathcal{I}, \mathcal{J}$, so that $\mathcal{U}' = \mathcal{U} - \{(i,j,k)\}, \mathcal{I}' = (\mathcal{I} \cup \{(i,j,k)\}) - \{(i-1,j-1,k-1)\}, \mathcal{J}' = (\mathcal{J} \cup \{(i,j,k)\}) - \{(i-1,j-1,k-1)\}$. We have $|\mathcal{U}' \cap C(0,0,0)| = |\mathcal{U} \cap C(0,0,0)| - 1$, and N is still a cutoff for $\mathcal{I}'$, so the induction hypothesis tells us that $\mathcal{J}'$ consists of $3\binom{N+3}{2} + 1$ points and contains $3\binom{N+2}{2}$ squares. However, $\mathcal{I}$ is obtained from $\mathcal{I}'$ (and $\mathcal{J}$ is obtained from $\mathcal{J}'$) by replacing (i,j,k) with $(i-1,j-1,k-1)$. (Figure 8 shows the effect of the replacement in plane projection.) This replacement preserves the number of points. Moreover, it also preserves the number of squares: since the three squares $s_a(i,j,k), s_b(i,j,k), s_c(i,j,k)$ are destroyed, and the three squares $s_a(i-1,j,k), s_b(i,j-1,k), s_c(i,j,k-1)$ are created, we need only check that (i,j,k) and $(i-1,j-1,k-1)$ cannot belong to any other squares. This holds by an argument much like the one in Table 1 — all the other possible squares to which (i,j,k) could belong would contain $(i+1,j,k)$, $(i,j+1,k)$, or $(i,j,k+1)$, which are blocked from occurring in $\mathcal{I}$ by $(i,j-1,k-1)$, $(i-1,j,k-1)$, $(i-1,j-1,k)$, respectively; all the other squares to which $(i-1,j-1,k-1)$ would belong would contain $(i-2,j-1,k-1)$, $(i-1,j-2,k-1)$, or $(i-1,j-1,k-2)$, which are blocked from $\mathcal{I}$ by $(i-1,j,k)$, $(i,j-1,k)$, $(i,j,k-1)$, respectively.

So $\mathcal{J}$ has the same numbers of points and of squares as $\mathcal{J}'$, completing the induction. $\square$

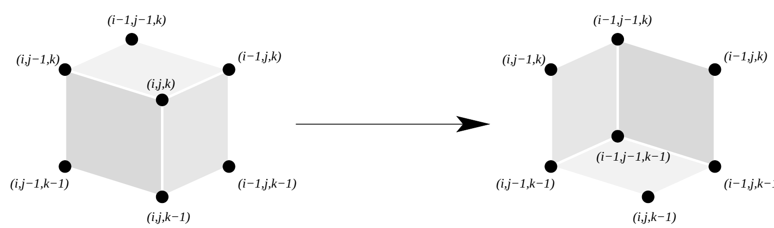

Figure 8: From $\mathcal{I}'$ to $\mathcal{I}$

This same process also can be used to show that every initial set really does correspond to a parallelogram tiling of the plane.

The proof of Theorem 1 is an application of the same technique as that of Lemma 3. The concept is analogous to that of the "urban renewal" proof in [10]: we show that varying the initial set in a controlled manner is tantamount to successively applying variable substitutions in $f_{0,0,0}$, and we interpret these substitutions combinatorially.

Proof of Theorem 1: We again fix N and induct on $|\mathcal{U} \cap C(0,0,0)|$. If the intersection is empty, then $\mathcal{I} = \{(i,j,k) \mid \max\{i,j,k\} = 0\}$. It is straightforward to check that the graph on $\mathcal{I}$ consisting of all short edges (shown in Figure 9) is a grove within radius N. The vertex $(0,0,0)$ has degree 3, all other vertices have degree 2, and no long edges occur; hence, the grove's corresponding monomial is $x_{0,0,0} = f_{0,0,0}$. We claim that there are no other $\mathcal{I}$-groves within radius N; the base case of the induction will then follow.

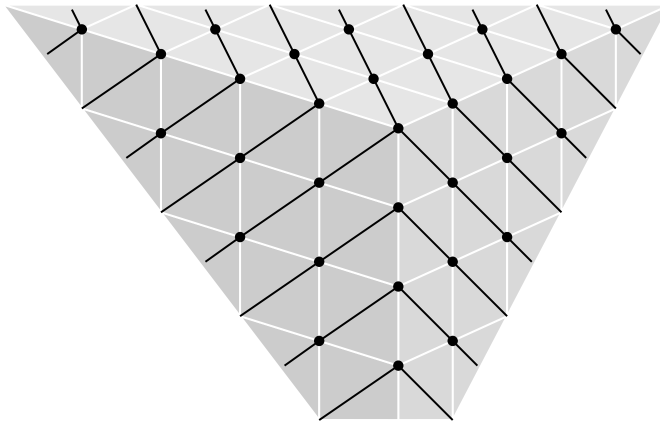

Figure 9: The unique grove for the base case

To see this, notice that the squares in $\mathcal{I}$ are precisely those of the forms $s_a(0,j,k)$, $s_b(i,0,k)$, and $s_c(i,j,0)$ $(i,j,k \leq 0)$. Suppose there exists a grove in which some long edge of the form $e_a(0,j,k)$ appears, and consider such an edge with $j+k$ minimal (only finitely many long edges occur). Then we know by compactness that $j+k \geq -N$, and each of the squares $s_a(0,j-l-1,k-l), s_a(0,j-l,k-l-1)$ $(l \geq 0)$ contributes its short edge. In particular, we have the infinite chain of edges

$$\begin{array}{ccccccc} \cdots & \leftrightarrow & (0,j-3,k-2) & \leftrightarrow & (0,j-2,k-1) & \leftrightarrow & (0,j-1,k) \\ & & & & & & \updownarrow \\ \cdots & \leftrightarrow & (0,j-2,k-3) & \leftrightarrow & (0,j-1,k-2) & \leftrightarrow & (0,j,k-1) \end{array}$$

However, if we put $l = \lceil (j+k+N)/2 \rceil$, then the connectivity condition tells us that $(0,j-l-1,k-l)$ should not be connected to $(0,j-l,k-l-1)$. Thus, we have a contradiction. So no grove (of any radius) can use any long edge of the form $e_a(0,j,k)$, and similarly the other two long edge types cannot be used. Hence, the only possible $\mathcal{I}$-grove is that consisting exclusively of short edges. The base case follows.

Now suppose $\mathcal{U} \cap C(0,0,0)$ is not empty. Let (i,j,k) be as given by Lemma 4, and define $\mathcal{L}'$,$\mathcal{U}'$,$\mathcal{I}'$ as in the proof of Lemma 3; we again observe that N remains a cutoff for $\mathcal{I}'$. Let $f'_{i',j',k'}$, for $(i',j',k') \in (\mathcal{I}' \cup \mathcal{U}') \cap C(0,0,0)$, be the Laurent polynomials (in variables $\{a_{j',k'}, b_{i',k'}, c_{i',j'}, x_{i',j',k'} \mid (i',j',k') \in \mathcal{I}'\}$) generated by the cube recurrence from the initial set $\mathcal{I}'$. Because $|\mathcal{U}' \cap C(0,0,0)| = |\mathcal{U} \cap C(0,0,0)| - 1$, we know by induction that $f'_{0,0,0} = \sum m(G')$, where the sum is over all $\mathcal{I}'$-groves G' within radius N. Let $g = \sum m(G)$, where the sum is over all $\mathcal{I}$-groves G within radius N. We wish to show that $f_{0,0,0} = g$.

On the other hand, we can readily see that $f_{0,0,0}$ is obtained from $f'_{0,0,0}$ by the variable substitution

$$x_{i,j,k} \leftarrow \frac{b_{i,k}c_{i,j}x_{i-1,j,k}x_{i,j-1,k-1} + c_{i,j}a_{j,k}x_{i,j-1,k}x_{i-1,j,k-1} + a_{j,k}b_{i,k}x_{i,j,k-1}x_{i-1,j-1,k}}{x_{i-1,j-1,k-1}}. \qquad (4)$$

Indeed, $f_{i',j',k'}$ is obtained from $f'_{i',j',k'}$ via this substitution, for any $(i',j',k') \in (\mathcal{I}' \cup \mathcal{U}') \cap C(0,0,0)$: it holds for $(i',j',k') = (i,j,k)$ because the left side of equation 4 *is* $f'_{i,j,k}$ and the right side is $f_{i,j,k}$; it holds for any other $(i',j',k') \in \mathcal{I}'$ trivially (because $f_{i',j',k'} = x_{i',j',k'} = f'_{i',j',k'}$), and then it holds for $(i',j',k') \in \mathcal{U}'$ by induction, since both $f_{i',j',k'}$ and $f'_{i',j',k'}$ are generated using the same recurrence (1). More simply, substitution (4) exactly captures the algebraic effect of changing our initial set from $\mathcal{I}'$ to $\mathcal{I}$.

We construct a correspondence between $\mathcal{I}'$-groves and $\mathcal{I}$-groves (within radius N); this correspondence is sometimes one-to-one, sometimes one-to-three, and sometimes three-to-one. It is defined as follows: given an $\mathcal{I}'$-grove, we consider the edges used in the three squares containing (i,j,k) and replace them by any of the corresponding sets of edges in the three squares containing $(i-1,j-1,k-1)$, as shown in Figure 10; the rest of the grove is left intact. The reverse operation — turning $\mathcal{I}$-groves into $\mathcal{I}'$-groves — is defined similarly. In view of the connectivity constraint, we see that every possible grove on either initial set does fall into one of the cases covered by Figure 10 (using the long edges of all three squares shown would leave the vertex $(i-1,j-1,k-1)$ or (i,j,k) isolated). It is evident that our operation preserves the compactness condition, once we observe that $i+j+k > -N$. Moreover, since the correspondence preserves all connectivity relations among vertices other than (i,j,k) and $(i-1,j-1,k-1)$, it quickly follows that the connectivity condition is preserved as well, and the other conditions are trivial. Thus, our operation does take $\mathcal{I}'$-groves to $\mathcal{I}$-groves, and vice versa.

Now let us express this correspondence algebraically. Consider any $\mathcal{I}'$-grove G'. Since (i,j,k) only belongs to three squares there, it has degree $3, 2$, or 1 (never 0). If it has degree 3, then $x_{i,j,k}$ has exponent 1 in $m(G')$. There are three corresponding $\mathcal{I}$-groves (case (i) in Figure 10), which we call G_1, G_2, G_3, in the order in which they appear in the figure. In G_1, vertex $(i-1,j-1,k-1)$ has degree 1; vertices $(i,j,k-1)$ and $(i-1,j-1,k)$ each have degree 1 greater than in G'; and new long edges $e_a(i-1,j,k)$ and $e_b(i,j-1,k)$ are used. All other vertices and edges are the same in G_1 as in G'. Thus,

$$m(G_1) = \frac{m(G')}{x_{i,j,k}} \cdot \frac{a_{j,k} b_{i,k} x_{i,j,k-1} x_{i-1,j-1,k}}{x_{i-1,j-1,k-1}}.$$

Performing similar analyses for G_2 and G_3 gives that $m(G_1) + m(G_2) + m(G_3)$ equals

$$\frac{m(G')}{x_{i,j,k}} \cdot \frac{b_{i,k} c_{i,j} x_{i-1,j,k} x_{i,j-1,k-1} + a_{j,k} c_{i,j} x_{i,j-1,k} x_{i-1,j,k-1} + a_{j,k} b_{i,k} x_{i,j,k-1} x_{i-1,j-1,k}}{x_{i-1,j-1,k-1}}.$$

Since $x_{i,j,k}$ appears in $m(G')$ with exponent 1, we see that $m(G_1) + m(G_2) + m(G_3)$ is obtained from $m(G')$ by the substitution (4).

Now suppose (i,j,k) has degree 2 in G'; thus, $x_{i,j,k}$ does not occur in $m(G')$. From cases (ii) of Figure 10, we see that there is one corresponding $\mathcal{I}$-grove G, and it is easy to check that $m(G) = m(G')$: (i,j,k), $(i-1,j-1,k-1)$ both have degree 2, no other vertices change degrees, and the one long edge that disappears is replaced by another long edge represented by the same variable. Since $x_{i,j,k}$ does not occur in $m(G')$, we have that $m(G)$ is obtained (trivially) from $m(G')$ by the substitution (4).

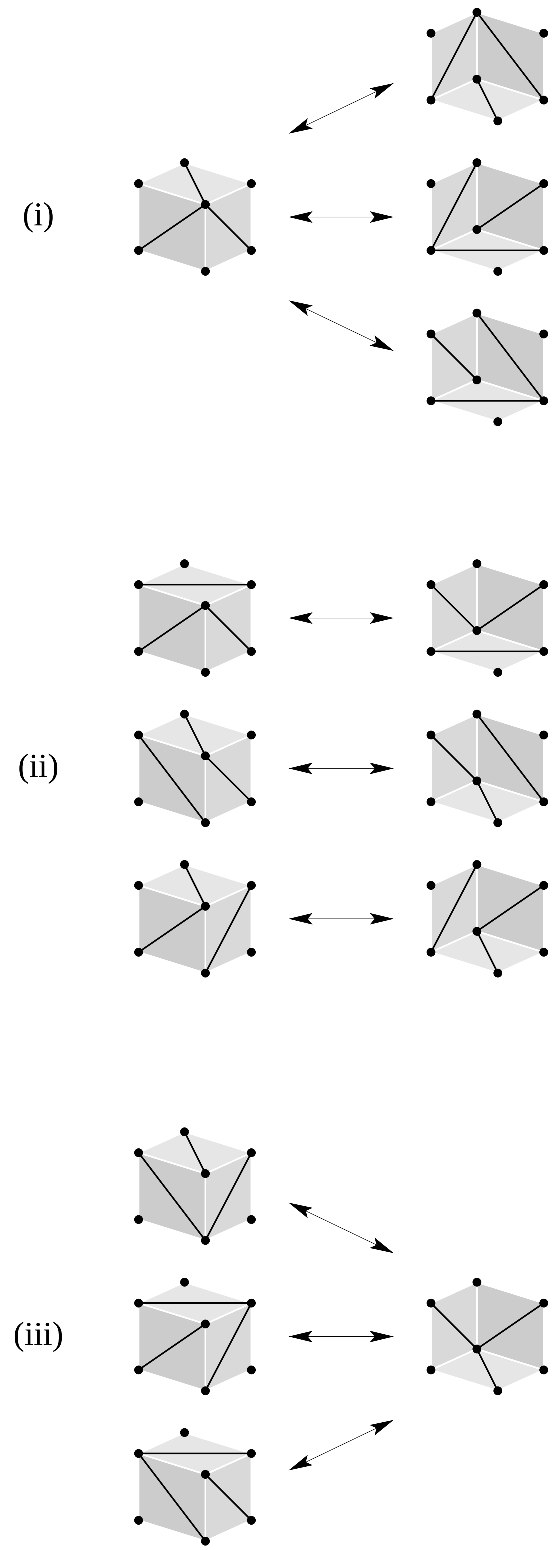

Figure 10: The correspondence between $\mathcal{I}'$-groves (left) and $\mathcal{I}$-groves (right)

Finally, if (i,j,k) has degree 1, then G' belongs to a triple $\{G'_1, G'_2, G'_3\}$ of $\mathcal{I}'$-groves, all of which correspond to the same $\mathcal{I}$-grove G, as shown in case (iii) of Figure 10. (We know that G'_1, G'_2, G'_3 really are all $\mathcal{I}'$-groves because they are obtainable via our correspondence from G, which in turn is obtainable from our original $\mathcal{I}'$-grove G'.) An analysis similar to that used in case (i) shows that

$$m(G'_1)/m(G) = a_{j,k}b_{i,k}x_{i,j,k-1}x_{i-1,j-1,k}/x_{i-1,j-1,k-1}x_{i,j,k}.$$

Similar computations with G'_2 and G'_3 give that $m(G'_1)+m(G'_2)+m(G'_3)$ equals

$$\frac{m(G)}{x_{i,j,k}} \cdot \frac{b_{i,k}c_{i,j}x_{i-1,j,k}x_{i,j-1,k-1} + c_{i,j}a_{j,k}x_{i,j-1,k}x_{i-1,j,k-1} + a_{j,k}b_{i,k}x_{i,j,k-1}x_{i-1,j-1,k}}{x_{i-1,j-1,k-1}},$$

and, using the fact that $x_{i,j,k}$ has exponent -1 in each of $m(G'_1), m(G'_2), m(G'_3)$, we conclude that $m(G)$ is obtained from $m(G'_1)+m(G'_2)+m(G'_3)$ by the same substitution (4).

Thus, summing over all $\mathcal{I}'$-groves and all $\mathcal{I}$-groves within radius N, we see that g is obtained from $f'_{0,0,0}$ by applying (4). Since this same substitution produces $f_{0,0,0}$ from $f'_{0,0,0}$, we have $f_{0,0,0} = g$, and the induction is complete. □

It is worth noting that a speedy combinatorial proof of Theorem 2 can be obtained by the same inductive process we used to prove Theorem 1 — acyclicity is preserved when an $\mathcal{I}'$-grove is replaced by a corresponding $\mathcal{I}$-grove. (Theorem 3 can also be proven this way.) We provided the noninductive proof earlier in part because it elucidates better how acyclicity follows directly from the other properties of a grove.

Readers familiar with delta-wye reduction of graphs, as described e.g. in [2], may notice a resemblance between the correspondence described in Figure 10 and delta-wye moves. Specifically: given $\mathcal{I}, \mathcal{I}'$, and (i,j,k) as in the proof of Theorem 1, suppose $\hat{\mathcal{G}}$ is the subgraph of $\mathcal{G}$ induced by the vertices whose sum of coordinates has the same parity as $i+j+k$, and suppose $\hat{\mathcal{G}}'$ is the subgraph of $\mathcal{G}'$ induced by the vertices of this parity. Then $\hat{\mathcal{G}}$ is obtained from $\hat{\mathcal{G}}'$ by a wye-delta move (wherein a vertex of degree 3 is deleted and three new edges are added joining this vertex's former neighbors).

In [2], Charles Colbourn, J. Scott Provan, and Dirk Vertigan describe a method for using delta-wye reduction to enumerate weighted spanning trees of a graph. Specifically, given an arbitrary finite graph G, one can associate to each edge e of G an "inclusion weight" i_e and an "exclusion weight" o_e, and then define

$$W(G) = \sum_T \Bigg(\prod_{e \in E(T)} i_e \Bigg) \Bigg(\prod_{e \notin E(T)} o_e \Bigg),$$

where the sum is over all spanning trees T of G. Colbourn, Provan, and Vertigan describe how the edge weights on a graph G' obtained from G by a wye-delta or delta-wye move may be assigned so that $W(G') = W(G)$. Although we were unaware of their work at the time we found the proof described above, it turns out that the variable substitution (4) is essentially identical to their weighting scheme. This should not be a huge surprise; after all, groves, being basically spanning forests constrained by certain global connectivity conditions, locally "look the same" as spanning trees and so should show similar susceptibility to local operations. To make our weighting scheme for groves fit into the framework of Colbourn, Provan, and Vertigan, which omits vertex weights, we must specify that when a square $s_a(i', j', k')$ contributes its long edge to $\hat{\mathcal{G}}$ (or $\hat{\mathcal{G}}'$), the inclusion weight for this edge is $a_{j',k'}x_{i',j'-1,k'}x_{i',j',k'-1}$ and the exclusion weight is $x_{i',j',k'}x_{i',j'-1,k'-1}$, and when the square contributes its short edge to $\hat{\mathcal{G}}$ (or $\hat{\mathcal{G}}'$), the exclusion weight for that short edge is $a_{j',k'}x_{i',j'-1,k'}x_{i',j',k'-1}$ and the inclusion weight is $x_{i',j',k'}x_{i',j'-1,k'-1}$ — and similarly for the other types of squares. We also need to restrict ourselves to a finite collection of squares dictated by the cutoff (so that the products involved in weighting groves are finite) and include an overall scale factor to make sure each $x_{i',j',k'}$ appears with the correct exponent. When all of this is properly handled, we have edge-weighting systems for $\hat{\mathcal{G}}$ and $\hat{\mathcal{G}}'$, and we also have a second weighting system for $\hat{\mathcal{G}}$ obtained by starting with our weighting system for $\hat{\mathcal{G}}'$ and applying the changes prescribed by Colbourn, Provan, and Vertigan. It turns out that the two weighting systems for $\hat{\mathcal{G}}$ are identical if and only if $x_{i,j,k}$ is given by the substitution (4). There is thus an intimate link between the cube recurrence and delta-wye reduction. We will not spell out all the details here, as it would require some space to develop all the necessary notation, but the interested reader is advised to obtain the paper [2] and verify that their edge-weighting scheme is tantamount to (4) in the case of groves.

6 Further questions

We now have one proof of Theorem 1. However, in [10], Speyer gives two proofs of the analogous statement for the octahedron recurrence (in terms of perfect matchings of graphs). One proof examines a Laurent polynomial with fixed indices under varying initial sets, as we have done; the second proof uses a fixed initial set and evaluates successive polynomials, making use the method of graphical condensation introduced by Eric Kuo in [7]. Briefly stated, the method operates as follows: Suppose that we wish to determine the number of perfect matchings $m(G)$ for some graph G. We consider certain pairs $\{A, B\}$ of subgraphs of G. When a perfect matching of A and a perfect matching of B are superimposed, the result is a multiset of edges of G, which we here will call a *superimposition from* $\{A, B\}$. Graphical condensation involves finding a subgraph H of G, and several (ordinarily two) further pairs of subgraphs, with the property that each superimposition from $\{G, H\}$ can also be expressed as a superimposition from exactly one of the other pairs, possibly together with some fixed set of additional edges associated with the pair (independent of the superimposition). If every superimposition has the same number of decompositions into a matching of G and a matching of H as decompositions on another pair of subgraphs, then it follows that $m(G)m(H) = \sum_{\{A,B\}} m(A)m(B)$ (here the sum is over our pairs). By placing formal-variable weights on the vertices and edges and associating to each matching and appropriate Laurent monomial in these weights, we can obtain a similar statement for matching-counting polynomials.

One would expect an analogous statement to hold true for groves: loosely speaking, by superimposing the edges of a grove G with those of another grove on a different initial set (allowing for some translation), we obtain a multiset of edges; this superimposition should roughly decompose again into groves on some other two initial sets, plus some extra edges. Indeed, because *every* $f_{i,j,k}$ counts groves on the initial set induced by a translation of $\mathcal{L} \cap C(i,j,k)$, the cube recurrence (1) implies that the grove-condensation statement, once properly formulated, must be true. This statement is rather vague, but an example is shown in Figure 11, here for the initial set induced by $\mathcal{L} = \{(i,j,k) \in C(0,0,0) \mid i+j+k \leq -4\}$. For visual simplicity, the short edges are omitted; the figure shows a multiset of long edges, together with two decompositions into groves on one pair of translated initial sets (left) and two decompositions on another pair (right). Notice that the decompositions on the right side of the figure lack the upper-right edge and the bottom edge; these are the fixed additional edges for this pair of initial sets.

However, we have not found a proof of Theorem 1 along these lines. The difficulty seems to lie in characterizing the objects obtained by combining two groves on different initial sets. A superimposition of perfect matchings of a graph and a subgraph is a relatively nice object — it is a multiset of edges such that certain vertices each belong to exactly two edges and all other vertices belong to exactly one; consequently, it can be characterized as a vertex-disjoint union of cycles, doubled edges, paths, and isolated vertices. In contrast, no simple description of the superimposition of two groves seems readily available, nor is it apparent what the procedure should be for decomposing superimpositions back into groves.

There is at least one specific difference between the hypothesized Kuo condensation for groves and the case of perfect matchings. In the usual form of Kuo condensation for matchings, each superimposition from $\{G, H\}$ also decomposes on just one of the other two pairs of subgraphs. We originally conjectured, by analogy, that each superimposition of two groves can only be decomposed on one of the three relevant other pairs of initial sets. Preliminary investigations have found this conjecture to be false. This makes the question of how Kuo condensation might work for groves all the more interesting, but deeper exploration of the topic is beyond the scope of the present paper.

Another open area of research concerning groves relates to the connection between the cube and octahedron recurrences discussed at the end of Section 4. What does the correspondence between perfect matchings and groves tell us about matchings — or about groves? In particular, the known Kuo condensation algorithm for perfect matchings may shed light on the analogue for groves.

One more interesting question relates to the form of the cube recurrence originally proposed by Propp. The polynomials generated by the octahedron recurrence, with the standard initial set $\{(i,j,k) \in \mathbb{Z}^3 \mid i+j+k \equiv 0 \bmod 2; k \in \{-1,0\}\}$, can be seen not only as enumerating matchings of graphs but also as enumerating compatible pairs of alternating-sign matrices: in each monomial, the exponents with which the variables $x_{i,j,k}$ appear correspond to the entries of the matrices. (See [9] for details.) In particular, if we set $x_{i,j,k} = 1$ whenever $k = -1$, then the terms of the resulting polynomials (in the variables $x_{i,j,0}$) correspond precisely to single alternating-sign matrices. As Propp

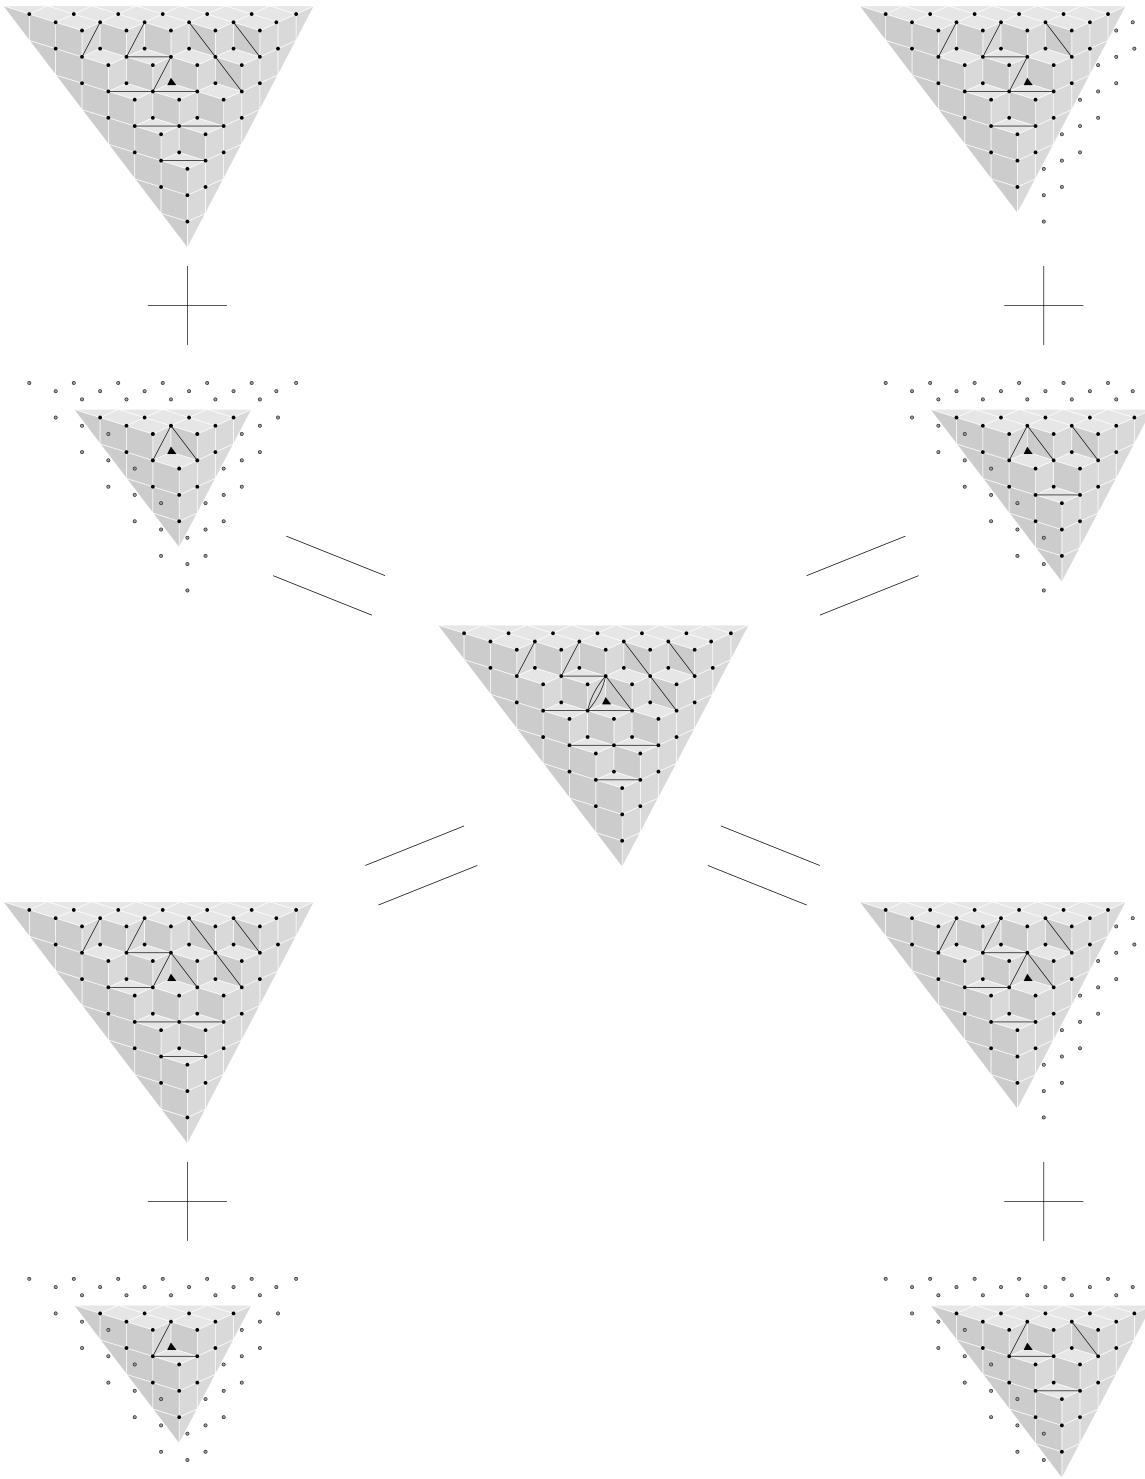

Figure 11: Decompositions of a two-grove superimposition; anchor is $(-2,-2,-2)$ in central picture

observed in [8], something analogous seems to be taking place with the cube recurrence for the standard initial set $\mathcal{I} = \{(i,j,k) \in \mathbb{Z}^3 \mid -1 \leq i+j+k \leq 1\}$. Specifically, fix $x_{i,j,k} = 1$ for $i+j+k = -1, 0$, and also set each edge variable equal to 1. Then any $f_{i,j,k}$ is a Laurent polynomial in $\{x_{i',j',k'} \mid i'+j'+k' = 1;\ i' \leq i, j' \leq j, k' \leq k\}$. For any monomial, the exponents of these $x_{i',j',k'}$, which must be $-1, 0$, or 1 (this is easy to check, since, after translation and intersection, each such point only belongs to three squares), form a sort of "alternating-sign triangle." An example is shown in Figure 12: at left is the grove for a particular monomial; at right is the alternating-sign triangle for the same monomial, formed by the exponents of the $x_{i',j',k'}$ corresponding to the circled vertices. The question therefore arises as to whether there exists some nice description for the set of

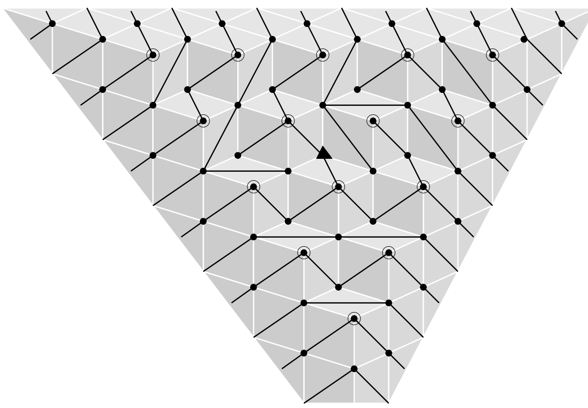

$$\begin{array}{ccccccccc}
0 & & 0 & & 0 & & 1 & & 0 \\
& 0 & & 1 & & -1 & & 0 & \\
& & 0 & & 1 & & 1 & & \\
& & & 0 & & 0 & & & \\
& & & & 0 & & & &
\end{array}$$

Figure 12: A standard grove of order 4 and the corresponding alternating-sign triangle; anchor is $(-2,-2,-2)$

alternating-sign triangles of a given size, analogous to that for alternating-sign matrices. The prospect of developing a theory of alternating-sign triangles is a tantalizing one.

The analogy with perfect matchings of graphs thus provides a wellspring of inspiration for further questions about groves. We hope that the results herein established will constitute the foundation for fruitful future study.

7 Acknowledgments

We would like to thank the other members of the REACH (Research Experiences in Algebraic Combinatorics at Harvard) group for assistance in the early stages of investigating the cube recurrence and for comments on drafts of this paper, as well as the anonymous referees for their extremely helpful comments in the final revision. We would also like to thank the National Security Agency and the National Science Foundation for their financial support of our research, as well as the mathematics departments of Harvard University and the University of Wisconsin-Madison for administrative assistance. Most of all, we thank Jim Propp, the director of REACH, for introducing us to the cube recurrence and for guiding and directing our work on the subject.

References

[1] M. Bousquet-Mélou. "Perfect matchings for Gale-Robinson sequences" (joint work with J. Propp and J. West), Séminaire Lotharingien de Combinatoire **49** (Ellwangen, October 2002).

[2] Charles J. Colbourn, J. Scott Provan, and Dirk Vertigan. "A new approach to solving three combinatorial enumeration problems on planar graphs," *Discrete Applied Mathematics* **60** (1995), 119-129.

[3] N. Elkies, G. Kuperberg, M. Larsen, and J. Propp. "Alternating-Sign Matrices and Domino Tilings," *Journal of Algebraic Combinatorics* **1** (1992), 111-132 and 219-234.

[4] S. Fomin and A. Zelevinsky. "The Laurent Phenomenon," *Advances in Applied Mathematics* **28** (2002), 119-144.

[5] S. Fomin and A. Zelevinsky. "Cluster Algebras I: Foundations," *Journal of the American Mathematical Society* **15** (2002), 497-529.

[6] D. Gale. "Mathematical Entertainments: The Strange and Surprising Saga of the Somos Sequences," *Mathematical Intelligencer* **13** (1991), 40-42.

[7] E. Kuo. "Applications of graphical condensation for enumerating matchings and tilings," *Theoretical Computer Science* **319** (10 June 2004), 29-57.

[8] J. Propp. "The Many Faces of Alternating-Sign Matrices," *Discrete Mathematics and Theoretical Computer Science Proceedings* **AA (DM-CCG)** (2001), 43-58.

[9] D. Robbins and H. Rumsey. "Determinants and Alternating Sign Matrices," *Advances in Mathematics* **62** (1986), 169-184.

[10] D. Speyer. "Perfect Matchings and the Octahedron Recurrence," preprint, available at `http://www.arxiv.org/abs/math.CO/0402452` .

[11] A. Zabrodin. "A survey of Hirota's difference equations," *Theoretical Mathematical Physics* **113** (1997), 1347-1392.