CONNECT FOUR AND GRAPH DECOMPOSITION

Laurent Evain, Mathias Lederer,
and Bjarke Hammersholt Roune

January 26, 2016

Abstract. We introduce standard decomposition, a natural way of decomposing a labeled graph into a sum of certain labeled subgraphs. We motivate this graph-theoretic concept by relating it to Connect Four decompositions of standard sets. We prove that all standard decompositions can be generated in polynomial time as a function of the combined size of the input and the output. This implies that all Connect Four decompositions can be generated in polynomial time.

Keywords. Standard sets, labeled graphs, polynomial time complexity

Subject classification. 05A17; 05A19; 05C30; 05C78; 68Q25

Acknowledgements

Thanks to Dustin Cartwright, Daniel Erman, Allen Knutson, Diane Maclagan, Jenna Rajchgot, Roy Skjelnes, Greg Smith, Mike Stillman and Elmar Teufl for fruitful discussions. We thank the anonymous referees for helpful suggestions.

1. Introduction

Let G be a directed graph. We say that an integer-valued labeling on the nodes of G is *compatible with the edge relation* if for all edges (a, b), the label of node a is less than or equal to the label of node b. Graphs satisfying that compatibility form the class of *standard graphs*; they are the objects of study of the present paper.

The paper is divided into two parts. In the first part, we study standard graphs and introduce a way of decomposing a standard graph as a sum of *standard components* — these are the standard subgraphs of G whose labels are 0 or 1. Here addition of labeled graphs is defined as addition of the labels. A *standard decomposition* of a standard graph is a multiset of standard components whose sum is the given graph. Standard components may be viewed as the building blocks of a standard graph.

Standard decomposition is not unique — standard graphs in general admit more than one standard decomposition. Figure 1.1 shows a simple example of a standard graph and all its standard decompositions. This raises the question of what the complexity of generating all standard decompositions given a standard graph is. Theorem 1.1 answers this question and it is the main result of the first part of the paper.

Theorem 1.1. *It is possible to generate all the standard decompositions of a standard graph in polynomial time.*

A word for clarity is needed. Example 3.1 shows that the number of decompositions can depend exponentially on the input. For such problems, it is natural to consider the so-called "generating complexity", in which one considers the running time as a function of the combined size of the input and the output. This is our framework in this paper, in particular when we assert the polynomial dependency in Theorem 1.1 above, or in Theorem 1.2 and its corollary below.

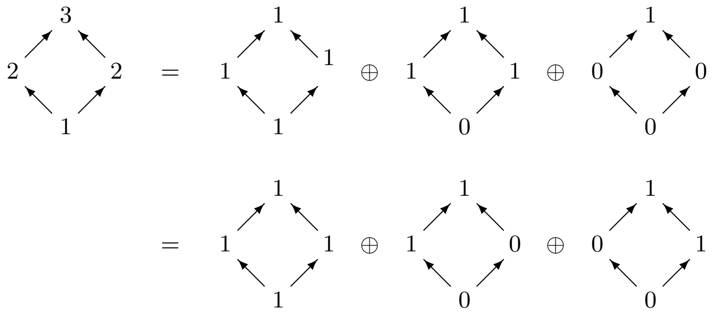

Figure 1.1: A standard graph and its two standard decompositions

In the second part of the paper, we link standard graphs and standard decomposition to a previously studied subject — *Connect Four decomposition of standard sets*. A standard set is an "n-dimensional staircase", and a Connect Four decomposition of a standard set Δ is a set of $n-1$ dimensional standard sets from which Δ can be built by stacking them on top of each other and "letting gravity pull them down". Figure 1.2 shows a simple example of a three dimensional standard set and all its Connect Four decompositions; the graph from Figure 1.1 encodes that same standard set.

Connect Four decomposition is a notion relevant to the study of the Hilbert scheme of points Lederer (2014), which appears as the combinatorial part in many techniques. For instance, it is useful in the study of singularities of plane curves as a tool for the Horace method Hirschowitz (1985), in the context of Gröbner basis theory Eisenbud (1995), Lederer (2008), Lederer (2014), to compute tangent spaces Nakajima (1999, Proposition 7.5), or to produce new counterexamples to Hilbert's fourteenth problem Evain (2005). Handling Connect Four decompositions is what originally prompted the work in this paper. We will show:

THEOREM 1.2. *(i) The two problems,*

(a) computing standard decompositions of labeled graphs, and

(b) computing Connect Four decompositions of finite standard sets,

are equivalent in the sense that for each labeled graph G, there exists a standard set Δ such that the standard decompositions of G are in canonical bijection with the Connect Four decompositions of Δ, and conversely.

(ii) This equivalence preserves polynomial complexity in the sense that for each labeled graph G, we can compute a standard set Δ with graph G in polynomial time, and for each standard set Δ, we can compute its graph $G(\Delta)$ in polynomial time.

COROLLARY 1.3. *It is possible to generate all Connect Four decompositions of a standard set in polynomial time.*

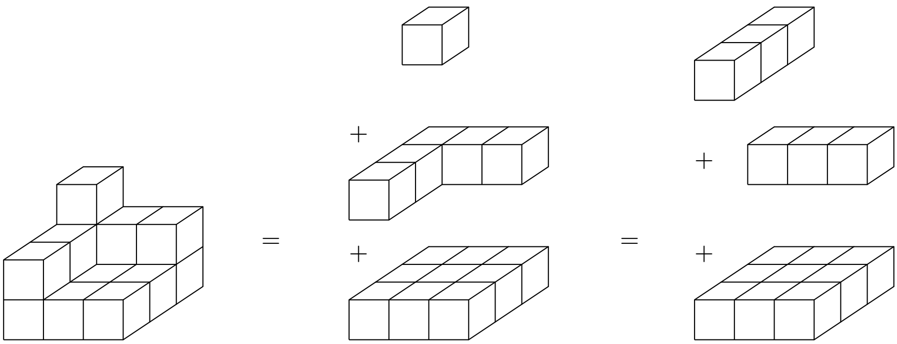

Figure 1.2: A standard set and its two Connect Four decompositions

We conclude our paper with an appendix which links the notions introduced in this paper to other classical tools and problems. First we present a generating function for the number of standard decompositions of a given graph. Then we show that the set of all *Connect Four games* in $\mathbb{N}^d$ of a given size n is in canonical bijection with the set of $(d-1)$-fold iterated partitions of n.

A word on the proofs. We prove Theorem 1.1 by presenting an algorithm that generates all standard decompositions in polynomial time. The algorithm is based on reducing the problem of computing all standard decompositions of G to the problem of computing all standard decompositions of G containing a fixed node v. We then solve that problem in a recursive way. Any choice of the node v results in a correct algorithm, yet we give a specific choice of v that allows the algorithm to generate its output in polynomial time.

The proof of Theorem 1.2 is done in several steps. We first attach a graph $G(\Delta)$ to each standard set Δ such that the standard decompositions of $G(\Delta)$ and the Connect Four decompositions of Δ are in canonical bijection. From $G(\Delta)$ we then define another graph $G'(\Delta)$ that is easier to work with, called the *canonicalized standard graph*, which has the same decompositions (see Proposition 7.5). We show that all labeled graphs arising from standard sets in this way have three specific properties, namely,

- they are standard,
- they are connected, and
- they have a unique node of maximal label.

Let $\mathcal{S}$ be the class of labeled graphs satisfying these conditions. The connectedness assumption in the definition of $\mathcal{S}$ is not essential for the complexity of the graphs from that class, since the standard decompositions of a disjoint union of graphs is the product of the standard decompositions of the individual graphs. We prove in Proposition 8.5 that each connected graph in $\mathcal{S}$ arises from a standard set if, in addition, the relation on the nodes of the graph defined by the edges of the graph is transitive. In Proposition 9.2, we show that for each connected standard graph, there exists a graph in $\mathcal{S}$ such that the standard decompositions of the two graphs are in canonical bijection.

2. Standard graphs and standard components

All graphs under consideration are directed, have finitely many nodes and do not have any parallel edges or loops. Given a graph, let $<$ be the partial preorder on the set of nodes such that $a < b$ if b is reachable from a. The graphs that we consider are labeled in the following sense.

Definition 2.1 (Labeled graph). *A* labeled graph *is a graph G with a finite node set $\mathcal{V}_G$ (possibly empty), an edge set $\mathcal{E}_G \subseteq \mathcal{V}_G \times \mathcal{V}_G$ such that the graph contains no loops (ie. $\forall a \in \mathcal{V}_G, (a,a) \notin \mathcal{E}_G$) and a labeling of nodes $\mathcal{L}_G \colon \mathcal{V}_G \to \mathbb{Z}$. If $(a,b) \in \mathcal{E}_G$, a is called the source of the edge and b is called the target of the edge.*

Definition 2.2 (Subgraph). *A* subgraph *of a labeled graph G is a labeled graph H with a finite node set $\mathcal{V}_H \subset \mathcal{V}_G$, an edge set $\mathcal{E}_H = (\mathcal{V}_H \times \mathcal{V}_H) \cap \mathcal{E}_G$, and a labeling of nodes $\mathcal{L}_H \colon \mathcal{V}_H \to \mathbb{Z}$ with $\mathcal{L}_H \leq \mathcal{L}_G$.*

This definition does not allow parallel edges since the edge set is not a multiset. The constraints on parallel edges and loops are not important to the results of this paper. We impose those conditions for simplicity since loops and parallel edges add nothing interesting to the problem.

Definition 2.3 (Standard graph). *A labeled graph G is* standard *if all labels are non-negative and the labeling is compatible with the partial order on the nodes in the sense that $\mathcal{L}_G(a) \le \mathcal{L}_G(b)$ for all edges $(a, b) \in \mathcal{E}_G$.*

We now introduce the operations of *addition* and *subtraction* on labeled graphs.

Definition 2.4 (Addition and subtraction). *Let G and H be labeled graphs. Suppose that $\mathcal{V}_G$ and $\mathcal{V}_H$ are subsets of a common set $\mathcal{V}$. Suppose that for all $(a, b) \in (\mathcal{V}_G \cap \mathcal{V}_H) \times (\mathcal{V}_G \cap \mathcal{V}_H)$, the conditions $(a, b) \in \mathcal{E}_G$ and $(a, b) \in \mathcal{E}_H$ are equivalent. Then $G \oplus H$ is the labeled graph with node set $\mathcal{V}_{G\oplus H} := \mathcal{V}_G \cup \mathcal{V}_H$, edge set $\mathcal{E}_{G\oplus H} := \mathcal{E}_G \cup \mathcal{E}_H$ and labeling*

$$\mathcal{L}_{G\oplus H} := \begin{cases} \mathcal{L}_G & \text{for } v \in \mathcal{V}_G \setminus \mathcal{V}_H, \\ \mathcal{L}_G + \mathcal{L}_H & \text{for } v \in \mathcal{V}_G \cap \mathcal{V}_H, \\ \mathcal{L}_H & \text{for } v \in \mathcal{V}_H \setminus \mathcal{V}_G. \end{cases}$$

We define $G \ominus H$ to have the same node set and edge set as $G \oplus H$, but with labeling

$$\mathcal{L}_{G\ominus H} := \begin{cases} \mathcal{L}_G & \text{for } v \in \mathcal{V}_G \setminus \mathcal{V}_H, \\ \mathcal{L}_G - \mathcal{L}_H & \text{for } v \in \mathcal{V}_G \cap \mathcal{V}_H, \\ -\mathcal{L}_H & \text{for } v \in \mathcal{V}_H \setminus \mathcal{V}_G. \end{cases}$$

The sum of two standard graphs is again a standard graph. This is illustrated in the example from Figure 2.1.

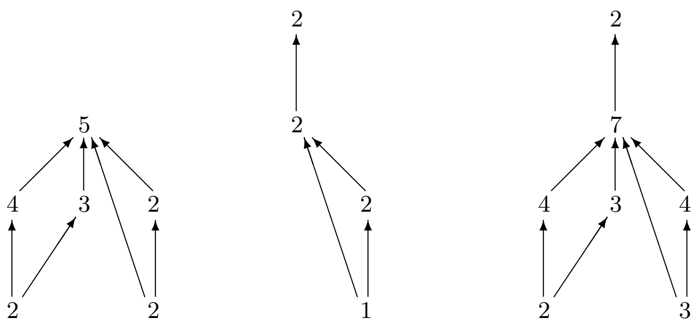

G H $G \oplus H$

Figure 2.1: The sum of two standard graphs

Definition 2.5 (0-1 graph). *A labeled graph is a* 0-1 graph *if all labels are 0 or 1.*

If we take a standard graph and replace all positive labels by 1, then we obtain another standard graph. This is a *standard 0-1 graph* — a graph that is both standard and a 0-1 graph. Some subgraphs H of a standard graph G are standard 0-1 graphs and in some cases we can write G as $H \oplus G'$, where G' is another standard graph. In this case we call H a *standard component* of G.

Definition 2.6 (Standard component). *Let G and H be labeled graphs with the same set of nodes and edges, ie. $\mathcal{V}_G = \mathcal{V}_H$ and $\mathcal{E}_G = \mathcal{E}_H$. Then H is a* standard component *of G if*

(i) H is a standard 0-1 graph;

(ii) $G \ominus H$ is a standard graph; and

(iii) not all labels in H are zero.

We think of standard components of G as the building blocks of G. Our goal is to determine all the ways to build a graph out of such building blocks.

Definition 2.7 (Standard decomposition). *Let G be a labeled graph. A multiset of labeled graphs $\mathcal{H}$ is a* standard decomposition *of G if each $H \in \mathcal{H}$ is a standard component of G and $G = \sum_{H \in \mathcal{H}} H$. We denote the set of standard decompositions of G by $\mathcal{D}(G)$.*

To keep formulas succinct, we use the shorthand notation $\sum \mathcal{H} := \sum_{H \in \mathcal{H}} H$.

A standard 0-1 graph G admits only the standard decomposition $\{G\}$. In particular, the building blocks of a graph are indecomposable; this is why we call them standard *components*. We define standard decompositions to be multisets rather than sets since a standard component can appear multiple times within one decomposition.

Example 2.8. Figure 2.2 shows a standard graph and all its decompositions. ◊

Since the sum of two standard graphs with the same nodes and edges is standard, a labeled graph has a standard decomposition only if it is standard. Proposition 2.10 shows that the converse is also true.

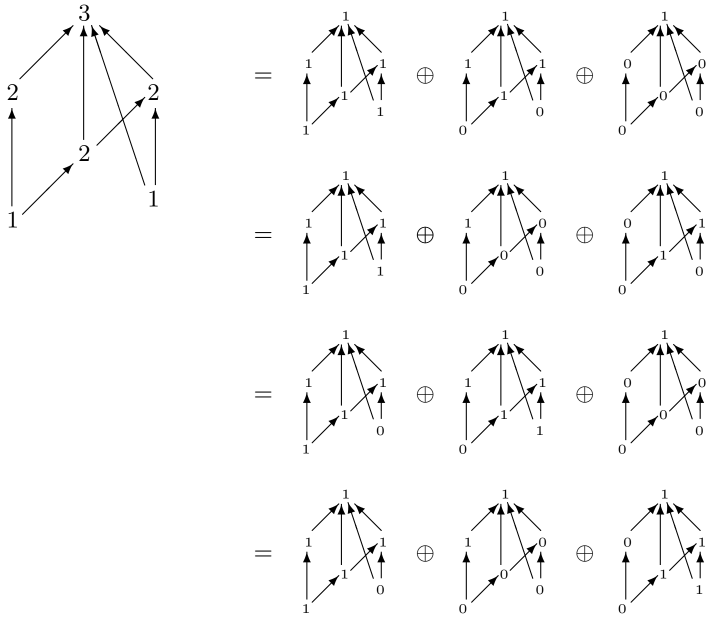

Figure 2.2: All decompositions of a graph.

Definition 2.9 (Maximal standard component). *The* maximal standard component *of a standard graph G is the unique standard component H for which $\mathcal{L}_H(v) = 1$ if, and only if, $\mathcal{L}_G(v) > 0$.*

This component is maximal in the sense that all other standard components are labeled subgraphs of the maximal standard component. Note that the maximal standard component is always a standard component unless we are in the degenerate case where all nodes of G are labeled zero.

Proposition 2.10. *Let $\mathcal{H}$ be a multiset of standard components of a labeled graph G. Then $\mathcal{H}$ can be extended to a standard decomposition of G if, and only if, $G \ominus \sum \mathcal{H}$ is standard. In particular, a standard graph admits a standard decomposition.*

Proof. **Proof of "if":** If $G = \sum \mathcal{H}$ then we are done, so suppose that $G \neq \sum \mathcal{H}$. Let C be the maximal standard component of $G \ominus \sum \mathcal{H}$. Then $G \ominus \sum (\mathcal{H} \cup \{C\})$ is standard. The assertion follows from this by induction.

Proof of "only if": If $\mathcal{H}'$ is a multiset of standard components of G that contains $\mathcal{H}$, and $G \ominus \sum \mathcal{H}$ is not standard, then neither is $G \ominus \sum \mathcal{H}'$, so $\mathcal{H}'$ is not a standard decomposition of G.

The last statement of the proposition follows from the first by taking $\mathcal{H} = \emptyset$.

□

Corollary 2.11. *If H is a standard component of a standard graph G, then H is an element of at least one standard decomposition of G.*

3. Standard node decompositions

We now turn to the topic of the computational complexity of the problem of computing standard decompositions. We start with a simple instructive example.

Example 3.1. Let G_n be the labeled graph defined by

$$\mathcal{V}_{G_n} := \{y, x_1, \ldots, x_n\}, \qquad \mathcal{E}_{G_n} := \{(x_1, y), \ldots, (x_n, y)\},$$

and $\mathcal{L}_{G_n}(x_i) := 1$ for $i = 1, \ldots, n$, while $\mathcal{L}_{G_n}(y) := 2$. There are 2^n standard components of G_n, corresponding to the n independent choices of whether to include or exclude the value at each x_i. The standard decompositions of G_n are pairs of standard components that include complementary subsets of $\{x_1, \ldots, x_n\}$. So G_n has 2^{n-1} standard decompositions while having only $n+1$ nodes. $\Diamond$

Consider the computational problem whose input is a labeled graph G and whose output is the set of standard decompositions $\mathcal{D}(G)$. Recall that $\mathcal{D}(G) \neq \emptyset$ if, and only if, G is standard — however, we will formulate our statements for arbitrary labeled graphs, thus covering also the case where the output is the empty set. Example 3.1 shows that this computation cannot be done in time better than exponential in the worst case since just writing down the output can take exponential time. For problems such as this, it is standard practice to consider an alternative notion of complexity, *generating complexity*, in which we consider the running time as a function of the combined size of the input *and* the output.

We present an algorithm for standard decomposition of graphs that runs in polynomial time in the combined size of input and output. This algorithm is based on the following notion of decomposing a single node of a standard graph.

Definition 3.2 (Standard node decomposition). *Let G be a labeled graph and let v be a node of G. A multiset of standard graphs $\mathcal{H}$ is a* standard v-decomposition *of G if*

(i) *each $H \in \mathcal{H}$ is a standard component of G,*

(ii) *$\mathcal{L}_H(v) = 1$ for all $H \in \mathcal{H}$,*

(iii) *$G \ominus \sum \mathcal{H}$ is standard,*

(iv) *$|\mathcal{H}| = \mathcal{L}_G(v)$.*

We denote the set of standard v-decompositions of G by $\mathcal{D}_v(G)$.

Consider a standard graph G with a standard decomposition $\mathcal{H}$ and a node v of G. The submultiset of $\mathcal{H}$ whose elements give v a label of 1 forms a standard v-decomposition of G. Another way of characterizing a standard v-decomposition is that it is a minimal multiset $\mathcal{H}$ of standard components of G such that $G \ominus \sum \mathcal{H}$ gives v the label 0 and such that $\mathcal{H}$ can be extended to a standard decomposition of G.

LEMMA 3.3. *Let G be a labeled graph. Let A be a multiset of standard 0-1 subgraphs of G and let B be a submultiset of A. Then A is a standard decomposition of G if, and only if, $A \setminus B$ is a standard decomposition of $G \ominus \sum B$.*

PROOF. **Proof of "if":** Assume that $A \setminus B$ is a standard decomposition of $G \ominus \sum B$. Then $G \ominus \sum B = \sum (A \setminus B)$ so $G = \sum A$. It only remains to prove that each $a \in A$ is a standard component of G. To prove that, we need to show that $G \ominus a$ is standard. We already know that a is a standard component of $G \ominus \sum B$, which implies that $G \ominus \sum B \ominus a$ is standard. Then $G \ominus a = (G \ominus \sum B \ominus a) \oplus \sum B$ is standard, as it is a a sum of standard graphs with identical node sets.

Proof of "only if": Assume that A is a standard decomposition of G. Then $G = \sum A$ so $G \ominus \sum B = \sum (A \setminus B)$. It only remains to prove that each $a \in A \setminus B$ is a standard component of $G \ominus \sum B$. To prove that we need to show that $G \ominus \sum B \ominus a$ is standard. We already know that $G \ominus \sum A$ has all labels zero, so it is standard. Then $G \ominus \sum B \ominus a = (G \ominus \sum A) \oplus \sum (A \setminus (B \cup \{a\})))$, so $G \ominus \sum B \ominus a$ is standard, as it is a sum of standard graphs with identical node sets. $\square$

Every standard graph has at least one standard decomposition, so Proposition 3.4 implies that if we can generate standard v-decompositions in polynomial time, then we can also generate standard decompositions in polynomial time.

PROPOSITION 3.4. *Let v be a node of a labeled graph G. Then*

$$\mathcal{D}(G) = \left\{ \mathcal{H} \cup \mathcal{H}' \,\middle|\, \mathcal{H} \in \mathcal{D}_v(G), \mathcal{H}' \in \mathcal{D}\left(G \ominus \sum \mathcal{H}\right) \right\},$$

where no decomposition appears twice on the right hand side.

PROOF. **Proof of $\subseteq$:** Let $D \in \mathcal{D}(G)$ and let $\mathcal{H}$ be the submultiset of D whose elements give v the label 1. Then $\mathcal{H} \in \mathcal{D}_v(G)$. It only remains to prove that $D \setminus \mathcal{H} \in \mathcal{D}(G \ominus \sum \mathcal{H})$, which follows from Lemma 3.3.

Proof of $\supseteq$: Let $\mathcal{H} \in \mathcal{D}_v(G)$ and let $\mathcal{H}' \in \mathcal{D}(G \ominus \sum \mathcal{H})$. Then $\mathcal{H}' \cup \mathcal{H}$ is a standard decomposition of G by Lemma 3.3.

Proof of "no duplicates": Let $\mathcal{H}_1, \mathcal{H}_2 \in \mathcal{D}_v(G)$, $\mathcal{H}'_1 \in \mathcal{D}(G \ominus \sum \mathcal{H}_1)$ and $\mathcal{H}'_2 \in \mathcal{D}(G \ominus \sum \mathcal{H}_2)$. Then $\mathcal{H}_1 \cup \mathcal{H}'_1 \neq \mathcal{H}_2 \cup \mathcal{H}'_2$ unless $\mathcal{H}_1 = \mathcal{H}_2$ and $\mathcal{H}'_1 = \mathcal{H}'_2$. Indeed, $\mathcal{H}'_1 \cup \mathcal{H}'_2$ is disjoint from $\mathcal{H}_1 \cup \mathcal{H}_2$, as the elements of

$\mathcal{H}'_1 \cup \mathcal{H}'_2$ give v the label zero, while the elements of $\mathcal{H}_1 \cup \mathcal{H}_2$ give v the label 1. □

4. Generating standard node decompositions

Proposition 3.4 reduces the problem of generating $\mathcal{D}(G)$ in polynomial time to the problem of generating the standard node decomposition $\mathcal{D}_v(G)$ in polynomial time for some freely chosen node v of G. In this section we investigate this problem. Our solution is based on choosing the right node v to decompose.

Consider the set of all standard components of G that give v the label 1. We impose an ordering, $H_1, \ldots, H_k$, on the elements of that set. This ordering can be chosen arbitrarily, but is fixed once and for all. Now let F be any labeled subgraph of G. For each such F and each $i = 1, \ldots, k$, we define

$$\tau(F, i) := \{\mathcal{H} \subseteq \{H_1, \ldots, H_i\} \,|\, \mathcal{H} \in \mathcal{D}_v(F)\}$$

Note that $\mathcal{D}_v(F)$ and $\tau(F, i)$ are both sets of multisets, thus the condition $\mathcal{H} \in \mathcal{D}_v(F)$ for a multiset $\mathcal{H}$ makes sense; If we can compute $\tau(F, i)$ in general then we can also compute $\mathcal{D}_v(G)$ since $\tau(G, k) = \mathcal{D}_v(G)$. In order to compute $\tau(F, i)$, consider the recursive formula

$$\tau(F, i) = \begin{cases} \{\emptyset\} \text{ if all labels of } F \text{ are zero, else} \\ \emptyset \text{ if } F \text{ is not standard or } i = 0, \text{ else} \\ \tau(F, i-1) \cup \{\mathcal{H} \cup \{H_i\} \,|\, \mathcal{H} \in \tau(F \ominus H_i, i)\}. \end{cases} \tag{4.1}$$

This way of writing τ immediately suggests an algorithm based on recursively evaluating the expression. It is a problem with this approach that this algorithm can spend a large amount of computational steps to determine that $\tau(F, i)$ is empty. This is an obstacle to proving that this algorithm generates its output in polynomial time.

We say that a pair (F, i) is *relevant* if $\tau(F, i) \neq \emptyset$, and *irrelevant* otherwise.[1] For making the algorithm generate its output in polynomial time, we need a criterion for detecting irrelevant pairs. We can use such a criterion to quickly eliminate irrelevant pairs in the algorithm.

Proposition 4.2. *Let v be a node of minimal positive label in a labeled graph G. Let $\mathcal{H}$ be a multiset of standard components of G that give v the label 1. Let H be the maximal standard component of G. Assume that $G \ominus \sum \mathcal{H}$ is standard. Let $\mathcal{H}'$ be the union of $\mathcal{H}$ and the multiset containing $\mathcal{L}_G(v) - |\mathcal{H}|$ copies of H. Then $\mathcal{H}'$ is a standard v-decomposition of G.*

Proof. Upon applying the proof of Proposition 2.10 to $\mathcal{H}$, we obtain a standard decomposition $\mathcal{H}'' \supseteq \mathcal{H}$ of G. Since the label of v is minimal among all positive labels appearing in G, the first $\mathcal{L}_G(v) - |\mathcal{H}|$ rounds of the inductive

[1] In particular, we see that it suffices to consider graphs such that $0 \leq \mathcal{L}_F(w) \leq \mathcal{L}_G(w)$ for all nodes w, since (F, i) is irrelevant otherwise.

construction in that proof will use the same maximal standard component H. After that the label of v has become zero, so the maximal standard components used in later rounds of the construction will give v the label zero. So the subset of $\mathcal{H}''$ that gives v the label 1 is precisely $\mathcal{H}'$, which implies that $\mathcal{H}'$ is a standard v-decomposition of G. □

Through choosing wisely the node v and the order of the standard components $H_1, \ldots, H_k$, Proposition 4.3 gives an if-and-only-if criterion for detecting irrelevant pairs.

PROPOSITION 4.3. *Given a labeled graph G, choose v to be a node of minimal positive label, and choose an order on the standard components $H_1, \ldots, H_k$ giving v the label 1 such that H_1 is the maximal standard component of G. Let $\mathcal{H}$ be a multiset whose elements are chosen among the standard components H_i of G, and let $F := G \ominus \sum \mathcal{H}$. Then a pair (F, i) with $1 \leq i \leq k$ is relevant if, and only if, F is standard.*

PROOF. **Proof of "if":** Assume that F is standard. By Proposition 4.2, $G = \sum \mathcal{H} \oplus \sum \mathcal{H}_1 \oplus \sum \mathcal{H}_2$, where $\mathcal{H}_1$ is a multiset containing copies of H_1 and $\mathcal{H}_2$ is a multiset containing standard components of G with label 0 on v. Thus $F = \sum \mathcal{H}_1 \oplus \sum \mathcal{H}_2$, and $\tau(F, i)$ is not empty.

Proof of "only if": This part is obvious. □

5. Generating standard decompositions in polynomial time

Based on the previous two sections, we can now present an algorithm for generating standard decompositions and prove that it runs in polynomial time.

THEOREM 5.1. *The algorithm in Figure 5.1 generates the standard decompositions of a labeled graph in polynomial time.*

The pseudo code for `standardDecompositions` implements the recursive formula from Proposition 3.4. The recursion from Section 4 is implemented in the pseudo code `standardNodeDecompositions`, where the function `Tau` is τ from that section. Line 20 eliminates pairs that are irrelevant according to Proposition 4.3.

In reading the pseudo code for `Tau`, note that the first return is of the value $\{\emptyset\}$ while the second is of the value $\emptyset$. Here $\{\emptyset\}$ is a set containing one decomposition while $\emptyset$ is a set containing nothing.

PROOF (Proof of Theorem 5.1 and thus also of Theorem 1.1). Recall that *generating output in polynomial time* means that the algorithm runs in polynomial time in the combined size of input and output — this is the meaning of the word "generate" in this context.

The size of the input and output depend on the representation used. We specify a graph as a list of nodes with labels and a list of edges. We specify

function STANDARDDECOMPOSITIONS(G)
  if all labels of all nodes of G are zero **then**
    return $\{\emptyset\}$
  else
    choose a node $v \in \mathcal{V}_G$ of minimal positive label
    $D \leftarrow$ STANDARDNODEDECOMPOSITIONS(G, v)
    return $\{\mathcal{H} \cup \mathcal{H}' \,|\, \mathcal{H} \in D, \mathcal{H}' \in$ STANDARDDECOMPS$(G \ominus \sum \mathcal{H})\}$
  end if
end function
function STANDARDNODEDECOMPOSITIONS(G, v)
  $H_1 \leftarrow$ the maximal standard component of G
  $H_2, \ldots, H_k \leftarrow$ all other standard components of G that give v the label 1
  $S \leftarrow \{H_1, \ldots, H_k\}$
  return TAU(G, k, S)
end function
function TAU(F, i, S)
  if all labels of F are zero **then**
    return $\{\emptyset\}$
  else
    if F is not a standard graph **then**
      return $\emptyset$
    else
      return TAU$(F, i-1, S) \cup \{\mathcal{H} \cup \{H_i\} \,|\, \mathcal{H} \in$ TAU$(F \ominus H_i, i, S)\}$
    end if
  end if
end function

Figure 5.1: An algorithm for standard decomposition.

the set of decompositions as a list of standard components followed by a list of sets that specify a decomposition by referring back to the list of components. Each standard component is specified by a bit per node indicating whether that node is an element of the standard component.

We assume a model where all labels and indices take up one word of space, rather than the logarithmic number of bits actually necessary to hold these numbers. The only arithmetic operations we perform is subtractions $a - b$ where $a > b$ so this assumption does not weaken the theorem.

`standardNodeDecompositions` is correct: Suppose that we call the function `standardNodeDecompositions` on the pair (G, v). We know that v is a node of minimal positive label in G since `standardDecompositions` always makes calls to `standardNodeDecompositions` with such a v. Also observe that the sequence $H_1, \ldots, H_n$ are ordered to satisfy the precondition of 20. We then see that `standardNodeDecompositions` computes the correct value $\mathcal{D}_v(G)$ since it directly implements the recursive formula from equation (4.1) along with the criterion for irrelevant pairs from Proposition 4.3.

`standardNodeDecompositions` is polynomial: Let G have n nodes and e edges. We do not give pseudo code for generating $H_1, \ldots, H_k$, but it is not difficult to do this in time $O(k(n+e))$ using backtracking. We first need to prove that $k(n+e)$ is polynomial in the size of the output.

Let l be the label of v in G. Every H_i is an element of at least one standard decomposition of G by Corollary 2.11, and each v-decomposition has exactly l elements, so $k \leq ld$ where d is the number of standard v-decompositions of G. So computing $H_1, \ldots, H_k$ can be done in time $O(ld(n+e))$. The size of the input is $\Theta(n+e)$ and the size of the output is $\Theta(ld + kn)$ since it takes l elements of S to specify each of the d decompositions and for each irreducible decomposition we need one bit per node to specify whether it is in the graph or not. Clearly $ld(n+e) = \Omega(ldn^2)$ is bounded above by a polynomial in $ld + kn$, so the time to compute S is polynomial.

It remains to prove that `Tau` takes polynomial time. Each individual call to `Tau`, not counting recursive subcalls, can be done in time $O(n+e)$. We need an upper bound for the number of recursive calls.

Consider a tree T where each recursive call to `Tau` is a node labeled by the parameters (F, i) and where there is an edge from the caller to the callee. The relevant leaves of T give rise to one distinct node decomposition per leaf so d, the number of v-decompositions of G, is also the number of relevant leaves of T. Let r be the number of irrelevant leaves of T — these do not give rise to a v-decomposition. Since T is a binary tree we see that there are $r + d - 1$ internal nodes in T. We need an upper bound for r.

Since Proposition 4.3 is an if-and-only-if criterion for irrelevant pairs, we see that the sub-tree rooted at any internal node contains a relevant pair. This implies that the sibling of an irrelevant leaf A is a root of a sub-tree that contains some relevant leaf B. Let f be the mapping $A \mapsto B$. If $f(A) = B$ then the parent of A is on the path from the root of T to B. All the relevant leaves are at depth k or less, so f can map at most k irrelevant leaves to each relevant leaf. This implies that $r \leq dk$.

We have seen that there are d relevant leaves, at most dk irrelevant leaves and therefore also at most $d+dk$ internal nodes in T, which is a total of at most $2d+2dk$ nodes. So the time taken by all recursive calls to `Tau` is $O(dk(n+e))$. Recall that the input size is $\Theta(n+e)$ and the output size is $\Theta(ld+kn)$. Clearly $dk(n+e)$ is dominated by a polynomial in $(n+e)+(ld+kn)$. This proves that `standardNodeDecompositions` generates $\mathcal{D}_v(G)$ in polynomial time.

`standardDecompositions` is correct: We have already done the correctness proof since `standardDecompositions` directly implements the recursive formula for $\mathcal{D}(G)$ from Proposition 3.4.

`standardDecompositions` is polynomial: We have seen that each call to `standardNodeDecompositions` generates its own output in polynomial time. Consider a tree T where each recursive call to `standardDecompositions` is a node with an edge from the caller to the callee. Let q be the number of leaves of T. Every leaf contributes at least one distinct decomposition to the output, so q is a lower bound on the number of decompositions of G. The multiset of node decompositions computed by all the calls to `standardNodeDecompositions` is in bijection with the edges of T. All trees have more nodes than edges and more leaves than internal nodes so the combined time to compute all the node decompositions is dominated by a polynomial in $q(n+e)$ where $n+e$ is the input size for the original input which is an upper bound on the size of any graph produced during the computation.

Line 7 could a priori seem to require too much time by going through all the elements of D. However, we can charge this work to each of the children of that node that are produced in this way which clears up the problem. As trees have more leaves than internal nodes the total number of nodes of T is less than $2q$. This proves that the total time to compute $\mathcal{D}(G)$ is bounded by a polynomial in $w(n+e)$ where w is the number of decompositions and $\Theta(n+e)$ is the size of the input. □

We can extract some bounds on the number of node decompositions from the arguments just given.

PROPOSITION 5.2. *Let v be a node of a standard graph G. Let $l := \mathcal{L}_G(v)$. If G has k standard components that give v label 1, then there are between $\frac{k}{l}$ and $\binom{k+l-1}{l}$ standard v-decompositions of G, and these two numbers coincide when $l = 1$. If v is a node of minimal positive label in G, then there are at least k standard v-decompositions of G.*

PROOF. Every v-decomposition of G has exactly l elements, and the elements of each such multiset are chosen among the k standard components that give v the label 1, so there cannot be more than $\binom{k+l-1}{l}$ standard v-decompositions.

Every one of the k standard components giving v label 1 can be extended to a standard decomposition of G by Corollary 2.11 and therefore also to a standard v-decomposition. We get the minimal number of standard v-decompositions when each of these extensions are unique. As each standard v-decomposition has l elements, that implies the existence of at least $\frac{k}{l}$ standard v-decompositions.

If v is a label of minimal positive label, then each standard component H that gives v the label 1 can be extended to a v-decomposition using only the maximal standard component by Proposition 4.2. So there are at least k standard v-decompositions in this case. □

Here are examples in which the bounds from the proposition are sharp.

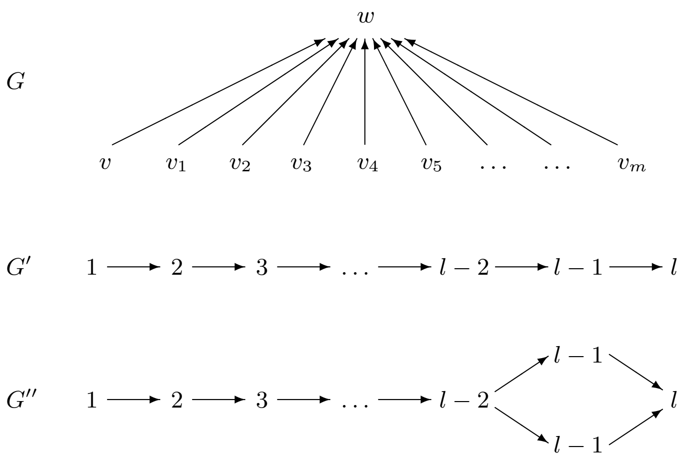

Figure 5.2: Three graphs leading to sharp bounds in Proposition 5.2

Example 5.3. Consider the graph G from Figure 5.2, whose labels we will presently specify, and the graphs G' and G'' from the same figure, whose labels are specified in the picture.

- Choose the labels such that $\mathcal{L}_G(v_i) \geq \mathcal{L}_G(v)$ for all i and $\mathcal{L}_G(w) \geq \mathcal{L}_G(v) + \sum_i \mathcal{L}_G(v_i)$. Then for each multiset of standard graphs $\mathcal{H}$ satisfying conditions (1), (2) and (4) from Definition 3.2, condition (3) is automatically satisfied. The number k from the proposition depends on the choice of the labels, but in any case, G has $\binom{k+l-1}{l}$ standard v-decompositions. The upper bound for the number of standard v-decompositions is therefore sharp.
- Define $\mathcal{L}_G(v) := 1$, $\mathcal{L}_G(v_i) := 1$ for all i and $\mathcal{L}_G(w) := 2$. Then the standard components that give v the label 1 correspond to the power set of $\{v_1, \ldots, v_m\}$, whose cardinality is 2^m. All standard components that

give v the label 1 lead to standard v-decompositions of G. The lower bound for the number of standard v-decompositions is therefore sharp.

- The graph G' provides another example of sharpness of the lower bound, this time with $l > 1$. We define v as the node of label l. As in the proposition, we denote by k the number of standard components of G' that give v the label 1. Since v is labeled 1 in every standard component, k is just the number of components of G'. Likewise, a standard v-decomposition of G' is just a standard decomposition of G'. Obviously $k = l$, and there exists precisely one standard v-decomposition.
- Also in the graph G'', we define v as the node of label l. This graph has the property that the lower bound is sharp while, unlike in the previous example, there exists more than one standard v-decomposition. Note that the fraction $\frac{k}{l} = \frac{2l-1}{l}$ is not an integer, but $\lceil \frac{k}{l} \rceil = 2$.

◊

We leave the question open whether there exist k and l as in the proposition such that $\frac{k}{l} > 2$ and there exists a graph G such that the lower bound from the proposition is sharp.

6. From standard sets to standard graphs

In the remaining three sections, we investigate the relation between standard decomposition of labeled graphs and another combinatorial problem called *Connect Four decomposition*. In the end we show that the two problems are equivalent.

A *standard set*, or *staircase*, is a subset $\Delta \subseteq \mathbb{N}^d$ whose complement $C := \mathbb{N}^d \setminus \Delta$ satisfies $C + \mathbb{N}^d = C$. We are only going to consider standard sets of finite cardinalities. Standard sets in $\mathbb{N}$ are just intervals starting at 0; in $\mathbb{N}^2$, they can be identified with partitions, or with Young diagrams[2]; in $\mathbb{N}^3$, they are also known as *plane partitions*; in $\mathbb{N}^d$ for $d > 3$, they are also known as *solid partitions*. Standard sets in $\mathbb{N}^d$ canonically correspond to monomial ideals in the polynomial ring $k[x_1, \dots, x_d]$. See Figure 6.1 for examples in dimensions 1, 2, and 3.

Let $d \geq 2$. Consider the projection to the first $d-1$ components, $q^d : \mathbb{N}^d \to \mathbb{N}^{d-1} : \beta \mapsto (\beta_1, \dots, \beta_{d-1})$ and its complementary projection, $q_d : \mathbb{N}^d \to \mathbb{N} : (\beta_1, \dots, \beta_d) \mapsto \beta_d$ For each standard set Δ, we have the equality

$$\Delta = \left\{ \beta \in \mathbb{N}^d \middle| \, q_d(\beta) < \left| (q^d)^{-1}(q^d(\beta)) \cap \Delta \right| \right\},$$

where $|.|$ denotes the cardinal. Thus, the integer $\left|(q^d)^{-1}(q^d(\beta)) \cap \Delta\right|$ appearing on the right-hand side is the cardinality of the fiber of the projection $q^d : \Delta \to$

[2] in the French notation

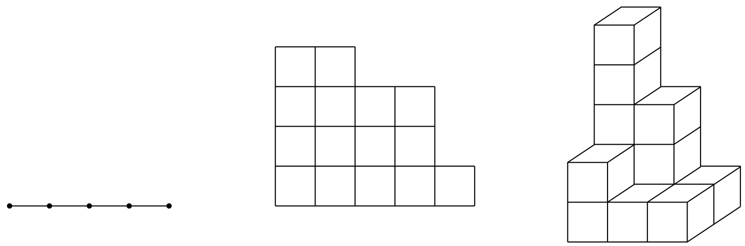

Figure 6.1: Standard sets in dimensions 1, 2 and 3

$\mathbb{N}^{d-1}$ over the point $\gamma := q^d(\beta)$. We call that quantity the *height* of Δ over γ. The equation displayed above implies that the datum of standard set Δ is equivalent to the datum of the projection $\Delta' := q^d(\Delta)$, which is a standard set in $\mathbb{N}^{d-1}$, and the datum of the heights over all $\gamma \in \Delta'$. The heights satisfy a compatibility condition: Upon denoting by h_γ the height over $\gamma \in \Delta'$, we see that $h_{\gamma+e_i} \le h_\gamma$ for all standard basis elements $e_i \in \mathbb{N}^{d-1}$ and all $\gamma \in \Delta'$ such that also $\gamma + e_i \in \Delta'$. These observations motivate the following definition:

Definition 6.1 (Standard graph of a standard set). *Let $\Delta \subseteq \mathbb{N}^d$ be a finite standard set. We define the* standard graph of Δ, *denoted by $G(\Delta)$, by setting*

$$\begin{aligned} \mathcal{V}_{G(\Delta)} &:= q^d(\Delta), \\ \mathcal{E}_{G(\Delta)} &:= \{\, (\gamma', \gamma) \,|\, \gamma' = \gamma + e_i \text{ for some } i\} \\ \mathcal{L}_{G(\Delta)}(\gamma) &:= \left|(q^d)^{-1}(\gamma) \cap \Delta\right|. \end{aligned}$$

The discussion leading to the definition proves that $G(\Delta)$ is indeed a standard graph. The transition from a standard set to its standard graph is illustrated in the first two pictures in Figure 7.1.

Addition of standard graphs has a counterpart on standard sets, called *C4 addition.*

Definition 6.2 (C4 sum). *Let Δ_1 and Δ_2 be two finite standard sets in $\mathbb{N}^d$. We define the* Connect Four sum, *or* C4 sum *of Δ_1 and Δ_2 by*

$$\Delta_1 + \Delta_2 := \left\{ \beta \in \mathbb{N}^d \;\middle|\; \begin{array}{c} q_d(\beta) < \left|(q^d)^{-1}\big(q^d(\beta)\big) \cap \Delta_1\right| \\ + \left|(q^d)^{-1}\big(q^d(\beta)\big) \cap \Delta_2\right| \end{array} \right\}.$$

So for determining the C4 sum of Δ_1 and Δ_2, we define Δ' to be the union of $q^d(\Delta_1)$ and $q^d(\Delta_2)$ and, for all $\gamma \in \Delta'$, h_γ to be the sum of the heights over γ of Δ_1 and Δ_2.[3] Then Δ is characterized by its projection Δ' and the heights h_γ.

[3] We say that the height of Δ_i over γ is zero if $\gamma \notin q^d(\Delta_i)$.

Here is a more graphic way of thinking about the C4 sum: Place Δ_1 and Δ_2 somewhere on the d-axis in $\mathbb{N}^d$ such that they do not intersect, subsequently drop the cubes along the d-axis, until they get stacked above each other on the $1, 2, \dots, (d-1)$-hyperplane. The result is the standard set $\Delta_1 + \Delta_2$. Figure 6.2 illustrates that process in two examples. The figure also explains the analogy to the eponymous game *Connect Four*.

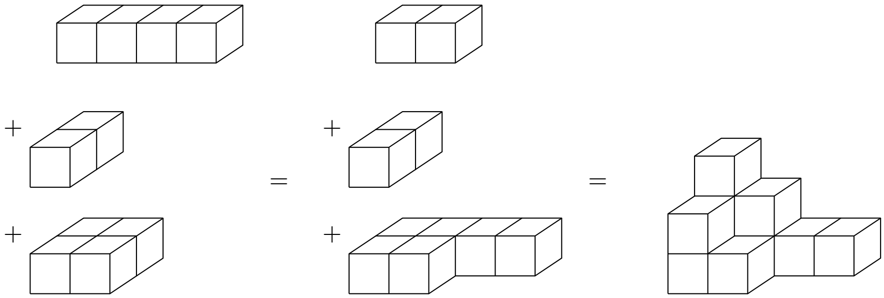

Figure 6.2: C4 sums of 2-dimensional standard sets yielding a 3-dimensional standard set

It is easy to see that

- $\Delta_1 + \Delta_2$ is a standard set;
- its cardinality is the sum of the cardinalities of Δ_1 and Δ_2;
- C4 addition is associative and commutative, and $\emptyset$ is its neutral element;
- $G(\Delta_1 + \Delta_2) = G(\Delta_1) \oplus G(\Delta_2)$.

The last item confirms that C4 addition of standard set is indeed the counterpart of addition of standard graphs. Here is the counterpart of standard decomposition of standard graphs.

Definition 6.3 (C4 decomposition). *Let $\Delta \subseteq \mathbb{N}^d$ be a finite standard set. A* C4 decomposition of Δ *is a multiset $\{\Delta_1, \dots, \Delta_h\}$ of standard sets in $\mathbb{N}^{d-1}$ whose C4 sum equals Δ. Here we understand each Δ_i to be a standard set in $\mathbb{N}^d$ via the embedding $\mathbb{N}^{d-1} \hookrightarrow \mathbb{N}^d : \gamma \mapsto (\gamma, 0)$.*

Figure 6.2 shows C4 decompositions of the standard set in $\mathbb{N}^3$ on the right hand side into two (multi)sets of standard set in $\mathbb{N}^2$. Note, however, that the three-dimensional standard set of that example has more C4 decompositions than the two shown in the figure.

The last coordinate plays a special role in the sense that projection q^d forgets the last coordinate. The constructions could obviously be done with an other projection. However, the number of decompositions depends on the chosen projection. For instance, if there are N decompositions of $\Delta \subset \mathbb{N}^d$, if we consider the injection $f : \mathbb{N}^d \to \mathbb{N}^{d+1}$, $(\beta_1, \ldots, \beta_d) \mapsto (0, \beta_1, \ldots, \beta_d)$, then $f(\Delta)$ still has N decompositions with our definition of q^{d+1}. However, if we had defined q^{d+1} to be the projection $(0, \beta_1, \ldots, \beta_d) \mapsto (\beta_1, \ldots, \beta_d)$, then $f(\Delta)$ would have had only one decomposition.

The following proposition is the first step of four in proving that C4 decomposition and standard decomposition of labeled graphs are equivalent.

Proposition 6.4. *Let $\Delta \subseteq \mathbb{N}^d$ be a finite standard set. Then the C4 decompositions of Δ and the standard decompositions of $G(\Delta)$ are in canonical bijection.*

Proof. Let $\{\Delta_1, \ldots, \Delta_h\}$ be a C4 decomposition of Δ. Consider, for $j = 1, \ldots, h$, the graph H_j whose nodes and edges are identical to the nodes and edges of $G(\Delta)$ and whose labeling is given by

$$\mathcal{L}_{H_j}(\gamma) = \begin{cases} 1 \text{ if } \gamma \in H_j \\ 0 \text{ else.} \end{cases}$$

In other words, we think of Δ_j, which is a priori a standard set in $\mathbb{N}^{d-1}$, as being a standard set in $\mathbb{N}^d$, as we do in Definition 6.3, and define $H_j := G(\Delta_j)$. Then H_j is obviously a standard 0-1 graph. The fact that $\{\Delta_1, \ldots, \Delta_h\}$ is a C4 decomposition of Δ implies that $\mathcal{H} := \{H_1, \ldots, H_h\}$ is a standard decomposition of $G(\Delta)$.

Conversely, let $\mathcal{H}$ be a standard decomposition of $G(\Delta)$. Recall that the node set of $G(\Delta)$ is $\Delta' := q^d(\Delta)$, which is a standard set in $\mathbb{N}^{d-1}$. For every $H \in \mathcal{H}$, we define $\Delta(H)$ to be the set of all $\gamma \in \Delta'$ with $\mathcal{L}_H(\gamma) = 1$. The definition of $\mathcal{E}_{G(\Delta)}$, together with the fact that H is a standard graph, shows that $\Delta(H) \subseteq \mathbb{N}^{d-1}$ is a standard set contained in Δ'. The fact that $\mathcal{H}$ is a standard decomposition of $G(\Delta)$ means that for each $\gamma \in \Delta'$, the labels of all nodes γ, which are 0 or 1, sum up to the height h_γ. This means that C4 sum of the corresponding multiset $\{\Delta(H) \,|\, H \in \mathcal{H}\}$ equals Δ, so that multiset is a C4 decomposition of Δ.

The two constructions are readily seen to be mutual inverses. □

7. Canonicalization for graphs of standard sets

The graph of a given standard set will in general contain many nodes of identical label connected by an edge. However, edges between nodes of the same label are irrelevant for computing the standard decomposition of that graph and we can get rid of those redundancies to speed up the computations. Similarly, nodes with label zero do not impact the computation of the standard decomposition.

By passing from a graph to its canonicalization (constructed below), we shall keep only the information necessary for the computation. The canonicalization process is the second step of four in proving that C4 decomposition and standard decomposition of labeled graphs are equivalent.

DEFINITION 7.1 (Canonical labeled graph). *A labeled graph G is* canonical *if*

(i) G is standard;

(ii) all labels are positive;

(iii) $\mathcal{L}_G(a) < \mathcal{L}_G(b)$ for all edges $(a,b) \in \mathcal{E}_G$.

Remark that if G has a cycle, i.e. a sequence of nodes $(a_1, a_2, \ldots, a_k)$ with $k \geq 2$, $a_k = a_1$ and $(a_i, a_{i+1}) \in \mathcal{E}_G$ for all i, then the labels of the nodes $(a_1, \ldots, a_{k-1})$ are equal since $\mathcal{L}_G(a_1) \leq \mathcal{L}_G(a_2) \leq \cdots \leq \mathcal{L}_G(a_k) = \mathcal{L}_G(a_1)$. In particular, a canonical graph has no cycle.

DEFINITION 7.2. *If G is a labeled graph, we denote by G_0 the subgraph of G with set of nodes $\mathcal{V}_{G_0} = \{a \in \mathcal{V}_G, \mathcal{L}_G(a) > 0\}$ and labels $\mathcal{L}_{G_0}(a) = \mathcal{L}_G(a)$ for all $a \in \mathcal{V}_{G_0}$.*
Let $\mathcal{R}$ be an equivalence relation on the set of nodes $\mathcal{V}_G$. Suppose that for every pair of nodes $(a,b) \in \mathcal{V}_G \times \mathcal{V}_G$ with $a\mathcal{R}b$, $\mathcal{L}_G(a) = \mathcal{L}_G(b)$. Then we denote by $G/\mathcal{R}$ the labeled graph defined by :

(i) $\mathcal{V}_{G/\mathcal{R}}$ is the quotient of $\mathcal{V}_G$ by the equivalence relation

(ii) $\forall a \in \mathcal{V}_{G/\mathcal{R}}$, $\mathcal{L}_{G/\mathcal{R}}(a) = \mathcal{L}_G(x)$ where x is any representative in the class of a.

(iii) $\forall (a,b) \in \mathcal{V}_{G/\mathcal{R}} \times \mathcal{V}_{G/\mathcal{R}}$ with $a \neq b$, $(a,b) \in \mathcal{E}_{G/\mathcal{R}}$ if and only if there exists $x \in a, y \in b$ representatives such that $(x,y) \in \mathcal{E}_G$.

We denote by $G/\simeq$ the quotient graph of G where $\simeq$ is the equivalence relation on $\mathcal{V}_G$ generated by the relation $a \simeq b$ if $(a,b) \in \mathcal{E}_G$ and $\mathcal{L}_G(a) = \mathcal{L}_G(b)$.

PROPOSITION 7.3. *Let G be a labeled graph. Then the graph $G_0/\simeq$ is canonical. It is called the canonicalization of G. The standard decompositions of G and its canonicalization $G_0/\simeq$ are in one-to-one correspondence.*

PROOF. All the labels in G_0 are positive by definition, and all the nodes with the same label connected by an edge have been contracted into a single point in $G_0/\simeq$. Thus $G_0/\simeq$ is canonical.
Let $G = \sum \mathcal{H}$ be a decomposition in standard components, where $\mathcal{H}$ is a multiset of labeled graphs. We denote by $\mathcal{H}_0 = \{H', H \in \mathcal{H}\}$ where H' is the labeled subgraph of H defined by $\mathcal{V}_{H'} = \mathcal{V}_H \cap \mathcal{V}_{G_0}$ and for every node a, $\mathcal{L}_{H'}(a) = \mathcal{L}_H(a)$.

Then $G_0 = \sum \mathcal{H}_0$ is a canonical decomposition and obviously all decompositions are obtained in such a way since G_0 and the labeled graphs of $\mathcal{H}_0$ can be extended by adding the nodes in $\mathcal{V}_G \setminus \mathcal{V}_{G_0}$ with label 0 and considering the same edges as in G. The bijection between the decompositions of G and G_0 follows.
Let H be a standard component of G and let (a, b) be an edge of G such that $\mathcal{L}_G(a) = \mathcal{L}_G(b)$. Then $\mathcal{L}_H(a) = 1$ if, and only if, $\mathcal{L}_H(b) = 1$. In other words, if two nodes a, b of G are equivalent for the relation $\simeq_G$ identifying connected nodes with the labels on G, these nodes are equivalent for the relation $\simeq_H$. In particular, it makes sense to consider $\mathcal{H}/\simeq_G = \{H/\simeq_G, H \in \mathcal{H}\}$. The decomposition $G = \sum \mathcal{H}$ yields a decomposition $G/\simeq_G = \sum \mathcal{H}/\simeq_G$. All the decompositions of $G/\simeq_G$ are obtained from this procedure.
Summing up, for any G, there is a canonical one-to-one correspondence between the components of G and the components of G' when $G' = G_0$ or $G' = G/\simeq$. It follows that there is a canonical identification between the decompositions of G and $G_0/\simeq$. □

We obtain the following explicit description of the canonicalization of a standard set Δ.

Definition 7.4 (Definition of $G'(\Delta)$).

- *We say that a non empty subset B of $\mathbb{N}^{d-1}$ is* connected *if for all $\gamma, \gamma' \in B$ with $\gamma \neq \gamma'$, there exists a sequence (γ_j) in B starting at $\gamma_0 = \gamma$ and ending at $\gamma_n = \gamma'$ such that for all j, we either have $\gamma_{j+1} = \gamma_j + e_i$ or $\gamma_j = \gamma_{j+1} + e_i$ for some $i \in \{1, \ldots, d-1\}$.*
- *A connected component of $A \subseteq \mathbb{N}^{d-1}$ is a connected $B \subseteq A$, maximal with respect to inclusion.*
- *Let $\Delta \subseteq \mathbb{N}^d$ be a standard set, $h := \max(q_d(\Delta))$ its height, and $\Delta' := q^d(\Delta)$ its projection. For $a = 1, \ldots, h$, we define the a-th* isohypse *as*

$$\Delta^a := \left\{\gamma \in \Delta' \,\middle|\, \left|(q^d)^{-1}(\gamma) \cap \Delta\right| = a\right\},$$

the set of all points in the projection of height a.

- *We define the graph $G'(\Delta)$ by*

$$\begin{aligned}\mathcal{V}_{G'(\Delta)} &:= \{\textit{connected components of } \Delta^a \,|\, a = 1, \ldots, h\};\\ \mathcal{E}_{G'(\Delta)} &:= \{(B, C) \,|\, \exists \gamma' \in B, \gamma \in C : \gamma' = \gamma + e_i \textit{ for some } i\}\\ \mathcal{L}_{G'(\Delta)}(C) &:= a, \textit{ when } C \textit{ is a connected component of } \Delta^a.\end{aligned}$$

The transition from Δ to $G(\Delta)$ and to $G'(\Delta)$ is illustrated in Figure 7.1.

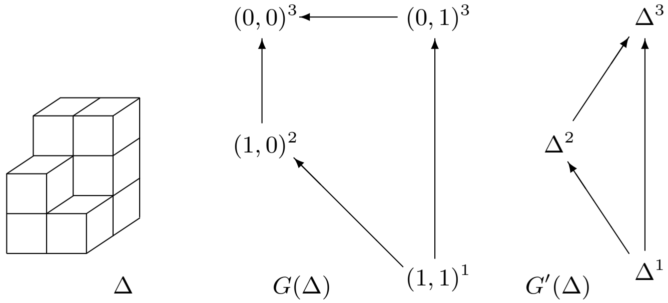

Figure 7.1: A standard set of height 3, its graph, and its canonicalized graph defined with the projection q^3 on a horizontal plane. A point $(a, b, c) \in \Delta$ is by convention represented by a 3-dimensional box of dimension $(1, 1, 1)$ centered on (a, b, c). The graph $G(\Delta)$ has 4 nodes corresponding to the 4 columns of boxes. The labels on the nodes (corresponding to the height of the columns) are denoted by superscripts. The two nodes generated by columns of height 3 have been identified in the canonicalization.

Proposition 7.5. *Let $\Delta \subseteq \mathbb{N}^d$ be a finite standard set. Then $G'(\Delta)$, as defined above, is the canonicalization of the standard graph of Δ.*

Proof. All the labels considered in $G(\Delta)$ are positive so no suppression of node is required. The equivalence relation $\simeq$ is generated by the identification of two nodes with the same label connected by an edge. Thus two nodes are identified in the equivalence relation iff they are connected by a chain of nodes of the same label. This is exactly the identification performed in the definition of $G'(\Delta)$. □

8. From standard graphs with unique maximal nodes to standard sets

For each standard set Δ, the canonicalized graph $G'(\Delta)$ is connected since it is constructed as a quotient of the connected graph $G(\Delta)$ and that taking the quotient preserves the connectedness. Moreover, $G'(\Delta)$ contains a unique node of maximal label, namely, the highest isohypse Δ^h. This graph thus lies in the class $\mathcal{S}$ defined in the Introduction. Example 8.1 and Proposition 8.2 show that graphs in $\mathcal{S}$ may or may not arise from standard sets.

Example 8.1. Figure 8.1 shows a standard graph which arises as the standard graph of a standard set in $\mathbb{N}^4$, namely,

$$\Delta = \left\{ \begin{array}{c} (0,0,0,0),(0,0,0,1),(0,0,0,2), \\ (1,0,0,0),(1,0,0,1),(0,1,0,0),(0,1,0,1), \\ (0,0,1,0),(0,0,1,1),(0,1,1,0),(0,1,1,1), \\ (1,1,0,0),(1,0,1,0) \end{array} \right\}.$$

The picture on the right hand side of that figure shows Δ^3, Δ^2 and $\Delta^1 \subseteq \mathbb{N}^3$. $\Diamond$

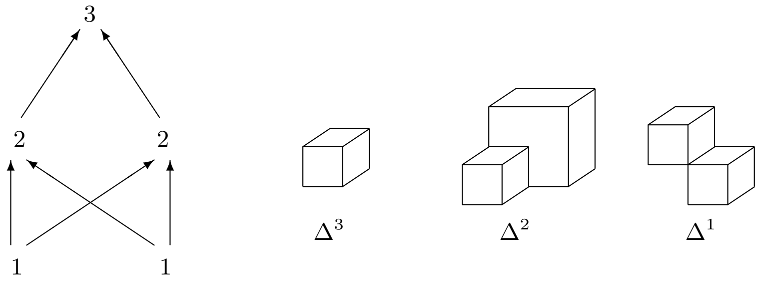

Figure 8.1: A standard graph arising from a standard set in $\mathbb{N}^4$

Proposition 8.2. *The graph shown in Figure 8.2 does not arise as the standard graph of a standard set.*

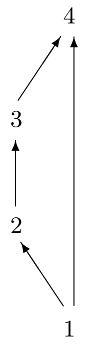

Figure 8.2: A graph not arising from a standard set

Proof. Assume that $\Delta \subseteq \mathbb{N}^d$ is a standard set whose standard graph is the given graph G. In particular, the nodes of G are the isohypses Δ^i, for $i =$

$1, 2, 3, 4$. We claim that there exists an element $\beta \in \Delta^1$ and $i, j \in \{1, \dots, d-1\}$ such that $\beta - e_i \in \Delta^2$ and $\beta - e_j \in \Delta^4$. This will finish the proof, since $\beta - e_i - e_j$ will then lie in Δ. But $\beta - e_i - e_j$ can lie in neither Δ^1 nor Δ^2 nor Δ^3, since either of these inclusions would contradict the standard set property of Δ. However, an inclusion $\beta - e_i - e_j \in \Delta^4$ would force an edge from node Δ^2 to node Δ^4 in the standard graph of Δ, which isn't there.

So we have to prove the above assertion. There exists elements $\sigma \in \Delta^4$ and $\tau \in \Delta^2$ and a sequence $(\gamma_k)_{k=0}^N$ such that

- its subsequence $(\gamma_k)_{k=1}^{N-1}$ lies in Δ^1,
- its starting point γ_0 is σ,
- its end point γ_N is τ, and
- it has the property that for all k, $\gamma_{k+1} = \gamma_k \pm e_i$ for some i.

Take σ, τ and (γ_k) sharing these properties such that, in addition, N, the length of the sequence (γ_k) is minimal. If $N = 2$, then $\beta := \gamma_1$ is of the desired shape. We now assume that $N > 2$, and are going to show that this assumption leads to a contradiction. For doing so, we prove three claims concerning the sequence (γ_k). The first claim is that for all $k < N$,

$$\gamma_k = \sigma + \sum_{i \in I_k} e_i \tag{8.3}$$

for some multiset of indices I_k. Note that $\gamma_k \in \Delta^1$ for all k in question. For $k = 0, 1$, equation ((8.3)) is evident. We assume that the equation holds for k and prove it to hold for $k+1$. Suppose that $\gamma_{k+1} = \sigma + \sum_{i \in I_k} e_i - e_j$ for some $j \notin I_k$. Then, since Δ is standard, $\sigma' := \sigma - e_j \in \Delta^4$. Consider the sequence $(\gamma'_l)_{l=0}^{N-1}$, where

$$\gamma'_l := \begin{cases} \gamma_l - e_j & \text{for } l < k, \\ \gamma_{l+1} & \text{for } l \geq k. \end{cases}$$

This sequence is one element shorter than the original sequence (γ_k). Like the original sequence, it starts in Δ^4 and ends in Δ^2. A priori the elements γ'_m, for $m = 1, \dots, k-1$, may lie in Δ^1, Δ^2, Δ^3 or Δ^4.

- If all of them lie in Δ^1, the sequence (γ'_l) contradicts the minimality of N.
- If $\gamma'_m \in \Delta^2$, the sequence $(\gamma''_p)_{p=0}^{m+1}$, where

  $$\gamma''_p := \begin{cases} \gamma_p & \text{for } l \leq m, \\ \gamma'_m & \text{for } p = m+1 \end{cases}$$

  contradicts the minimality of N.

- If $\gamma'_m \in \Delta^3$, using the absence of elements in $\Delta^2 2$ from the last case and the fact that Δ is standard, we obtain an edge from node Δ^1 to node Δ^3, which isn't there.
- If $\gamma'_m \in \Delta^4$, we consider the largest index M such that $\gamma'_M \in \Delta^4$ and consider the subsequence $\gamma'_M, \dots, \gamma'_N$. The condition $\gamma_i \in \Delta^1$ implies that $\gamma'_i \in \Delta^1 \cup \Delta^2 \cup \Delta^4$, since there is no node from Δ^1 to Δ^3. Thus the first term of $\gamma'_M, \dots, \gamma'_N$ lies in Δ^4, and the other terms in $\Delta^1 \cup \Delta^2$. The subsequence $\gamma'_M, \dots, \gamma'_{M'}$, where $M' \geq M$ is the smallest index with $\gamma'_{M'} \in \Delta^2$, contradicts the minimality of N.

This finishes the proof of the first claim.

Our second claim is that $I_k \subseteq I_{k+1}$ for all sets appearing in ((8.3)). This is true for $I_0 \subseteq I_1$; moreover, since $\gamma_{k+1} = \gamma_k \pm e_i$, our first claim shows that either $I_k \subseteq I_{k+1}$ or $I_k \supseteq I_{k+1}$ holds. Let m be the smallest index such that $I_m \supseteq I_{m+1}$. Then the sequence $(\gamma_k)_{k=0}^{m+1}$ is obtained by adding to σ a number of e_i, one by one, and finally subtracting one of them, say e_j. We obtain a shorter sequence $(\gamma'_k)_{k=0}^{m-1}$ by adding to σ the same sequence of e_i as above, but leaving out e_j. The same arguments as the ones from the four bulleted items above then lead to a contradiction. This finishes the proof of the second claim.

Our third claim is that τ, the final member of our sequence (γ_k), takes the shape

$$\tau = \gamma_N = \gamma_{N-1} - e_j,$$

for some $e_j \notin I_{N-1}$. The complementary cases include $\gamma_N = \gamma_{N-1} - e_j$ for some $e_j \in I_{N-1}$, which immediately contradicts minimality of N, and $\gamma_N = \gamma_{N-1} + e_j$ for some e_j. In the latter case, the inclusion $\gamma_N = \tau \in \Delta^2$ shows that γ_{N-1} would also lie in Δ^2, a contradiction. This finishes the proof of the third claim.

The sequence $(\gamma_k)_{k=0}^N$ is therefore obtained by adding to σ a number of e_i, one by one, and finally subtracting some e_j which is not found among the e_i previously added. We denote by e_u the last element from the sequence of e_i which we add, that is, the one element which we add for passing from γ_{N-2} to γ_{N-1}. Consider $\rho := \tau - e_u$. Then ρ may lie in Δ^2, Δ^3 or Δ^4.

- If $\tau' \in \Delta^2$, the sequence $(\gamma'_l)_{l=0}^{N-1}$, where

  $$\gamma'_l := \begin{cases} \gamma_l & \text{for } l \leq N-2, \\ \tau' & \text{for } l = N-1 \end{cases}$$

  contradicts the minimality of N.
- If $\tau' \in \Delta^3$, we obtain an edge from node Δ^1 to node Δ^3, which isn't there.
- If $\tau' \in \Delta^4$, we obtain an edge from node Δ^2 to node Δ^4, which isn't there.

So we have disproved the assumption that $N > 2$. The proposition follows. $\square$

The graphs in Figure 8.1 and Figure 8.2 define relations on their respective node sets which both fail to be transitive. So one might not guess that transitivity of graphs in $\mathcal{S}$ is crucial for such graphs to arise from standard sets. That, however, is indeed true, as we shall see in Proposition 8.5 below. Let us first establish that passing from a graph to its transitive closure has no impact on standard decompositions.

LEMMA 8.4. *Let G be a standard graph and $\overline{G}$ its transitive closure. Then the standard decompositions of G and $\overline{G}$ are in canonical bijection.*

PROOF. Given a standard decomposition $\mathcal{H}$ of G, replace every member H by its transitive closure $\overline{H}$. The resulting multiset $\overline{\mathcal{H}}$ is a standard decomposition of $\overline{G}$. Given a standard decomposition $\mathcal{K}$ of $\overline{G}$, we delete from every member K all edges that appear in $\overline{G}$ but not in G, and call the resulting graph K°. The resulting multiset $\mathcal{K}^\circ$ is a standard decomposition of G. The maps $\mathcal{H} \mapsto \overline{\mathcal{H}}$ and $\mathcal{K} \mapsto \mathcal{K}^\circ$ are mutual inverses. $\square$

The following proposition is the third step of four in proving that C4 decomposition and standard decomposition of labeled graphs are equivalent.

PROPOSITION 8.5. *Let G be a canonical, connected and transitive standard graph containing a unique node of maximal label. Then there exists a standard set $\Delta \subseteq \mathbb{N}^d$, for some $d \geq 1$, whose canonicalized standard graph is G.*

PROOF. Upon using the terminology of Definition 7.4, we denote by $G'(\Delta)$ the canonicalized standard graph of a standard set Δ. We prove the proposition by two nested inductions, the outer over the number of nodes of G, and the inner over the number of edges of G. The base case of the outer induction is trivial. As for the outer induction step, let G be a given connected and transitive standard graph containing a unique node v_h of maximal label, h. Let v_0 be a node of minimal label. We remove from G the node v_0, along with all edges whose source is v_0. We call the graph thus obtained G_0. Then G_0 is also canonical, connected and transitive. Canonicity and transitivity are obvious. As for connectedness, we note that each node in G other than the node v_0 is the starting point of a sequence of edges ending up in v_h, which sequence does not pass through v_0 by minimality of v_0 and canonicity of G. Moreover, when replacing G by G_0, we do not change the labels of the remaining nodes. Thus G_0 contains a unique node of maximal label. We may therefore assume that there exists a standard set $\Delta_0 \subseteq \mathbb{N}^d$, for some d, such that $G'(\Delta_0) = G_0$.

For establishing the outer induction step, we shall put the node v_0 back into the graph. Transitivity of G implies that this graph contains an edge from v_0 to v_h. Let G_1 be the (transitive) graph that arises from G_0 by adding the one node v_0 and the one edge (v_0, v_h). We now construct a standard set Δ_1 such that $G'(\Delta_1) = G_1$.

Consider the embedding $\iota : \mathbb{N}^d \hookrightarrow \mathbb{N}^{d+1} : \beta \mapsto (0, \beta)$. The transition from Δ_0 to $\iota(\Delta_0)$ does not affect the standard graph of Δ_0. We may therefore assume that $\Delta_0 \subseteq \mathbb{N}^d$ is contained in the hyperplane $\{\beta_1 = 0\}$ of $\mathbb{N}^d$. The node $v_h \in G_0$ corresponds to the isohypse $(\Delta_0)^h$. Let $h_0 < h$ be the label of v_0. We may assume that $v_0 > 1$. The set

$$\begin{aligned} \Delta_1 &:= \Delta_0 \cup M_1, \text{ where} \\ M_1 &:= \{(1, 0, \ldots, 0, \beta_d) \,|\, 0 \leq \beta_d \leq h_0 - 1\} \end{aligned} \tag{8.6}$$

is standard. See Figure 8.3 for a visualization of the transition from Δ_0 to Δ_1. For $a \neq h_0$, the isohypses $(\Delta_0)^a$ and $(\Delta_1)^a$ are identical. The isohypse $(\Delta_1)^{h_0}$ is $(\Delta_0)^{h_0} \cup q^d(M_1) = (\Delta_0)^{h_0} \cup \{e_1\}$. When passing to $G'(\Delta_1)$, we see that this graph arises from $G'(\Delta_0)$ by adding the one node $q^d(M_1)$ and the one edge connecting that new node and $(\Delta_1)^h$. This establishes the outer induction step, and at the same time the inner induction basis.

As for the inner induction step, we may assume to have a transitive graph G_1

- with the same nodes and the same labels as G,
- and a distinguished node v_0
- such that all edges but those with source v_0 agree in G and G_1,

along with a standard set $\Delta_1 \subseteq \mathbb{N}^d$ such that $G'(\Delta_1) = G_1$. Let v_1 be a node of G such that (v_0, v_1) is an edge in G, but our original graph G contains no chain of edges from v_0 to v_1 of length more than 1. We may assume that $\mathcal{L}_G(v_0) < \mathcal{L}_G(v_1)$. Denote by G_2 the graph that arises from G_1 by adding the edge (v_0, v_1). We will prove the existence of a standard set Δ_2 such that $G'(\Delta_2) = G_2$. This will establish the inner induction step, and finish the proof of the proposition.

Analogously as above, we assume that $\Delta_1 \subseteq \mathbb{N}^d$ is contained in the hyperplane $\{\beta_1 = 0\}$ of $\mathbb{N}^d$. The choice of v_1 implies that G_2 is again transitive. For $i = 0, 1$, the node $v_i \in G_1$ corresponds to a connected component C_i of $(\Delta_1)^{h_i}$, where h_i is the label of v_i. The set

$$\begin{aligned} \Delta_{1\frac{1}{2}} &:= \Delta_1 \cup M_{\frac{1}{2}}, \text{ where} \\ M_{1\frac{1}{2}} &:= \Big(\cup_{\alpha \in \mathbb{N}^d} \left((q^d)^{-1}(C_1) \cap \Delta_1 + e_1 - \alpha\right)\Big) \cap \mathbb{N}^d \end{aligned} \tag{8.7}$$

is standard. See the first two pictures in Figure 8.4 for a visualization of the transition from Δ_1 to $\Delta_{1\frac{1}{2}}$: We create a copy of the set $(q^d)^{-1}(C_1) \cap \Delta_1$ in the hyperplane $\{\beta_1 = 1\}$ of $\mathbb{N}^d$ and subsequently pass to the smallest standard set containing both Δ_1 and that copy. Transitivity of G_1 implies that $G'(\Delta_{1\frac{1}{2}}) = G'(\Delta_1)$. Indeed, for all heights $a \neq h_1$, the connected components of $(\Delta_{1\frac{1}{2}})^a$ are identical to of the connected components of $(\Delta_1)^a$. For height h_1, the same is true for those connected components of $(\Delta_{1\frac{1}{2}})^{h_1}$ that do not project to C_1 under q^d. The connected component C_1 of $(\Delta_1)^{h_1}$, however, has a much larger counterpart in $\Delta_{1\frac{1}{2}}$, namely, the union of C_1 and the set $q^d(M_{1\frac{1}{2}})$. As for edges in $G'(\Delta_{1\frac{1}{2}})$ emerging from node $C_1 \cup q^d(M_{1\frac{1}{2}})$, the presence of $M_{1\frac{1}{2}}$ obviously leads to new adjacencies in connected components of isohypses of $\Delta_{1\frac{1}{2}}$. But transitivity of G_1 guarantees that none of those adjacencies lead to an edge in $G'(\Delta_{1\frac{1}{2}})$ that does exist in $G'(\Delta_1)$. So the graphs $G'(\Delta_1)$ and $G'(\Delta_{1\frac{1}{2}})$ are identical.

However, we do not want another standard set with the same canonicalized graph, but rather a graph with one additional edge. We obtain that edge by applying the same trick once more, defining

$$\begin{aligned} \Delta_2 &:= \Delta_1 \cup M_{1\frac{1}{2}} \cup M_2, \text{ where} \\ M_2 &:= \Big(\cup_{\alpha\in\mathbb{N}^d} \left((q^d)^{-1}(C_1) + e_1 - \alpha\right)\Big) \cap \mathbb{N}^d. \end{aligned} \tag{8.8}$$

This is another standard set. See the last two pictures in Figure 8.4 for a visualization of the transition from $\Delta_{1\frac{1}{2}}$ to Δ_2: We also create a copy of the set $(q^d)^{-1}(C_1) \cap \Delta_1$ in the hyperplane $\{\beta_1 = 1\}$ of $\mathbb{N}^d$ and subsequently pass to the smallest standard set containing both Δ_1 and that copy. For all heights $a \neq h_0, h_1$, the connected components of $(\Delta_2)^a$ are identical to the connected components of $(\Delta_{1\frac{1}{2}})^a$. For heights $a = h_0, h_1$, the same is true for those connected components of $(\Delta_2)^a$ that do not project to C_0 or C_1. Note that the sets $M_{1\frac{1}{2}}$ and M_2 will in general intersect. The counterpart of C_1 in Δ_2 is the union $C_1 \cup M_{1\frac{1}{2}}$; and the counterpart of C_0 in Δ_2 is $(C_0 \cup M_2) \setminus M_{1\frac{1}{2}}$. The graph $G'(\Delta_2)$ contains all the edges that appear in $G'(\Delta_{1\frac{1}{2}})$, plus an edge from node $(C_0 \cup M_2) \setminus M_{1\frac{1}{2}}$ to node $C_1 \cup M_{1\frac{1}{2}}$: the extra edge exists since $(1, 0, \dots, 0, h_1) \in C_1 \cup M_{1\frac{1}{2}}$ and $z_0 + e_1 \in (C_0 \cup M_2) \setminus M_{1\frac{1}{2}}$ for $z_0 \in (q^d)^{-1}(C_0)$. This establishes the inner induction step. □

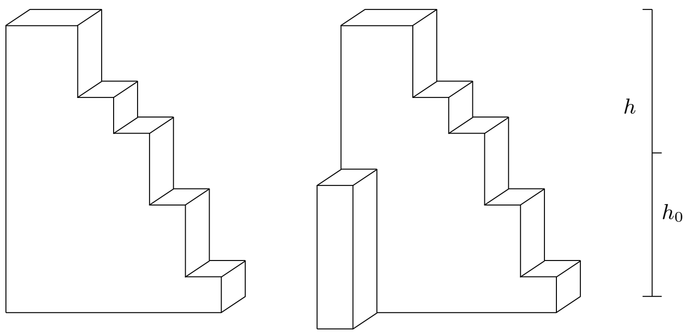

Figure 8.3: From Δ_0 to Δ_1

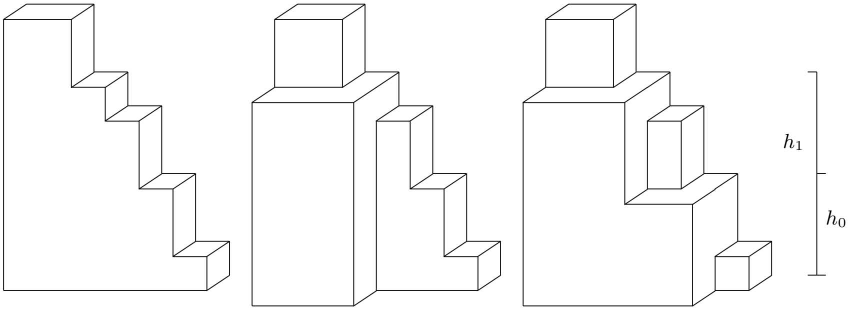

Figure 8.4: From Δ_1 to $\Delta_{1\frac{1}{2}}$ and Δ_2

Readers might wonder how the polynomial dependence from Theorem 1.1 is preserved in Proposition 8.5. Indeed, in the inductive construction of the standard set Δ from the proof of the proposition, the dimension of Δ and the number of elements in it grow rapidly. However, we don't specify Δ as list of its elements, but rather as a list of the minimal generators of the $\mathbb{N}^d$-module $\mathbb{N}^d \setminus \Delta$. This set is also known as the set of *outer corners* of Δ. Doing so, we avoid large data sets when handling large standard sets. We will use this representation of Δ in the proof of Theorem 1.2 below.

9. Reduction to standard graphs with unique maximal nodes

The following proposition provides the fourth and last step in proving that C4 decomposition and standard decomposition of labeled graphs are equivalent. Here is a small example illustrating its assertion.

EXAMPLE 9.1. Let G be the graph with nodes x and y, both of label 1, and no edges. Let G' be the graph with nodes x and y of label 1 and z of label 2, with edges from x and from y to z. Figure 9.1 shows that there is a bijection between the standard decompositions of G and the standard decompositions of G'. $\Diamond$

PROPOSITION 9.2. *Let G be a labeled graph. Then there exists a graph G' such that*

(i) G is a subgraph of G',

(ii) G' has a unique node of maximal label that is reachable from all nodes of G' and,

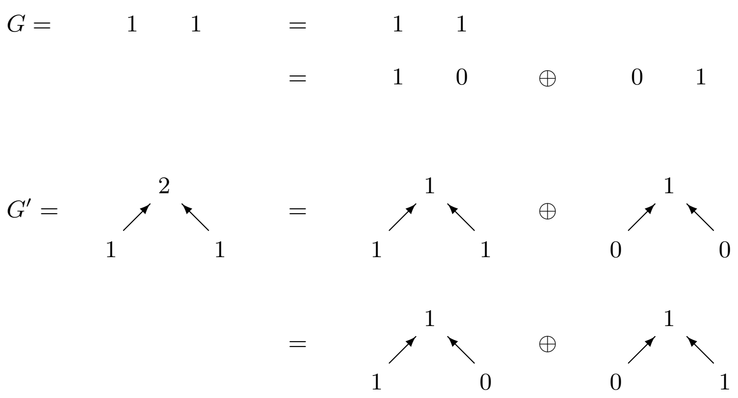

Figure 9.1: The decompositions of the graphs from Example 9.1

(iii) the standard decompositions of G and G' are in canonical bijection.

Proof. Let l be the maximal label among all nodes of G. Let G' be equal to G, except that G' has an extra node v with label $l+1$, and v has an edge to it from all other nodes. The first two conditions are immediate, so it remains to show that the standard decompositions of G and G' are in canonical bijection.

It is not hard to see that the function

$$\begin{aligned} f: \{\text{stdd comps of } G'\} \setminus \{\{v\}\} &\to \{\text{stdd comps of } G\} \\ H &\mapsto H \setminus \{v\} \text{ without edges in } H \text{ with target } v \end{aligned}$$

is a bijection. We extend f to a map of multisets of standard components by applying it to each component individually, so for example $f(\{A, B\}) := \{f(A), f(B)\}$.

Let D' be a standard decomposition of G'. Write D' as a union of D and V where V is a multiset that contains only copies of $\{v\}$ while D does not contain $\{v\}$ at all. Then obviously $f(D)$ is a decomposition of G.

For the other direction, let D be a standard decomposition of G and let $D' := f^{-1}(D)$. Then $G' - \sum D'$ is a graph in which all nodes but v have label zero, node v having a label $l > 0$. Let V be the multiset that contains l copies of $\{v\}$. Then $D' \cup V$ is a decomposition of G'. Let g be the function $D \mapsto D' \cup V$. It is not hard to see that f and g are mutual inverses. □

We can now prove that C4 decomposition and standard decomposition of labeled graphs are equivalent.

Proof (Proof of Theorem 1.2). (i) A solution of problem (a) implies a solution of problem (b) by Proposition 6.4. Assume we are able to solve problem (b), and are given a labeled graph G. We pass to the canonicalization G', which has the same standard decompositions as G by Proposition 7.3.

If G' has multiple nodes of locally maximal label l, we pass to the graph G'' with only one node of maximal label $l+1$ from Proposition 9.2. G'' still has the same standard decompositions as G. Then we replace G'' by its transitive closure G'''. By Lemma 8.4, this transition does not harm the decompositions either. Finally, Proposition 8.5 provides a standard set Δ''' whose canonicalized standard graph is G'''. Problem (a) is solved.

(ii) This assertion depends on the representations of G and Δ. In the proof of Theorem 5.1, we explained that we specify a graph as a list of nodes with labels and a list of edges. After the proof of Proposition 8.5, we explained that we specify a standard set by its outer corners.

Let us first show that for any graph G with n nodes and e edges, a staircase Δ whose graph equals G can be computed in polynomial time. We may assume that G is canonical and transitive, and has only one node of maximal label, since the operations

- ∘ passing to the canonicalization,
- ∘ passing to a graph with only one node of maximal label, and
- ∘ passing to the transitive closure

are obviously polynomial in the datum of G. It therefore remains to show that the construction from the proof of Proposition 8.5 is polynomial. That construction builds Δ using two nested inductions over n and e. The respective base cases being trivial, it suffices to show that both induction steps are polynomial in the datum of G. Let us stick to the notation from the proof of Proposition 8.5. In addition to that notation, we define $\mathscr{C}_i \subseteq \mathbb{N}^d$ as the set of corners of Δ_i for $i = 0, 1, 2$. In both the inner and the outer induction, the dimension of the standard sets involved rises by one. Thus the dimension d is polynomial in the datum of G. The outer induction step is the passage from Δ_0 to Δ_1, as defined in ((8.6)). That definition shows that $e_1 \in \mathscr{C}_0$ and

$$\mathscr{C}_1 = (\mathscr{C}_0 \setminus \{e_1\}) \cup \{e_1 + e_i \,|i = 1, \ldots, d-1\} \cup \{h_0 e_d\}\,,$$

cf. Figure 8.3. The inner induction step is thus polynomial.

The inner induction step is the passage from Δ_1 via $\Delta_{1\frac{1}{2}}$ to Δ_2. Remember that for $i = 0, 1$, the node $v_i \in G_1$ corresponds to a connected component C_i of $(\Delta_1)^{h_i}$. Let $\mathscr{C}'$ be the union of the following three sets:

- ∘ all corners $\alpha \in \mathscr{C}_1$ such that $\alpha - e_j \in (q^d)^{-1}(C_1) \cap \Delta_1$ for some $e_j \neq e_1$,
- ∘ the projections to the hyperplane $\{x_d = 0\}$ of all corners $\alpha \in \mathscr{C}_1$ such that $\alpha - e_j \in (q^d)^{-1}(C_1) \cap \Delta_1$ for some $e_j \neq e_1, e_d$, and
- ∘ the elements $2e_1$ and $h_1 e_d$.

Then $\mathscr{C}'$ is the set of corners of $M_{1\frac{1}{2}}$ from ((8.7)). Remember that $\Delta_{1\frac{1}{2}}$ is the union of Δ_1 and $M_{1\frac{1}{2}}$. The set $\mathscr{C}_{1\frac{1}{2}}$ of corners of $\Delta_{1\frac{1}{2}}$ is therefore obtained by

- collecting the exponents of least common multiples of $x^\alpha\ x^\beta$, for all $\alpha \in \mathscr{C}_1$ and all $\beta \in \mathscr{C}'$,
- and subsequently cleaning that set up, that is, detecting pairs α, β such that $\alpha \in \beta + \mathbb{N}^d$ and deleting each such α.

This establishes the passage from Δ_1 to $\Delta_{1\frac{1}{2}}$. As for the passage from $\Delta_{1\frac{1}{2}}$ to Δ_2, we construct a set of corners $\mathscr{C}''$ in an analogous way as we constructed $\mathscr{C}'$ in the three bulleted items above, but using C_0 rather than C_1 and h_0 rather than h_1. Then Δ_2 is the union of $\Delta_{1\frac{1}{2}}$ and the standard set with corners $\mathscr{C}''$. The set $\mathscr{C}_2$ is therefore obtained from sets $\mathscr{C}_{1\frac{1}{2}}$ and $\mathscr{C}''$ by the method of taking least common multiples and cleaning up which we employed above. All operations are polynomial.

Let us now show that for each standard set Δ, its canonicalized graph $G'(\Delta)$ can be computed in polynomial time. In other words, we have to compute the connected components of the isohypses in polynomial time. We assume Δ to be given by its corner set $\mathscr{C}$. For each $\alpha \in \mathscr{C}$, we define $\Delta_\alpha := q^d(\alpha) + \oplus_{i=1}^{d-1}\mathbb{N}e_i$. For each height a, we define Δ_a as the union of all Δ_α, for all α with $|\alpha| \leq a$. Then the a-th isohypse is

$$\Delta^a = \Delta_a \setminus \Delta_{a-1} = \cup_{|\alpha|=a}(\Delta_\alpha \setminus T_{a-1}).$$

Obviously each $E_\alpha := \Delta_\alpha \setminus T_{a-1}$ is connected. Moreover, it is easy to see that $E_\alpha \cup E_\beta$ is connected if, and only if, the least common multiple of the monomials $x^{q^d(\alpha)}$ and $x^{q^d(\beta)}$ has its exponent outside of T_{a-1}. Upon applying this observation to all α, β of total degree a, we compute the connected components of the a-th isohypse in polynomial time. □

A. A generating function

We will now present a natural generating function for the number of standard decompositions of a standard graph G. The analogue of this generating function in the setting of standard sets is discussed in Lederer (2014, Section 2.3).

It is good to temporarily forget about labelings. So let F be an unlabeled directed graph. Let $\mathcal{E}$ be the set of all standard 0-1 subgraphs of F[4] with node set $\mathcal{V}_F$. We identify each $E \in \mathcal{E}$ with the *characteristic function* of the labeling, that is, with the vector $\chi_E := (\chi_{E,v})_{v\in\mathcal{V}_F}$ indexed by nodes of F, with entries

$$\chi_{E,v} := \begin{cases} 1 & \text{if } v \in \mathcal{V}_E \\ 0 & \text{else.} \end{cases}$$

We define $\chi := (\chi_{E,v})_{E\in\mathcal{E},v\in\mathcal{V}_F}$ to be the matrix whose rows are indexed by $\mathcal{E}$, the row with index E being the vector χ_E. Moreover, we introduce

[4] We defined standard 0-1 subgraphs only for labeled graphs; if F is unlabeled, we give each node the trivial label 1; then the notion of 0-1 subgraphs is well-defined.

a vector $t := (t_v)_{v \in \mathcal{V}_F}$ of indeterminates, also indexed by nodes of F. If $w := (w_v)_{v \in \mathcal{V}_F}$ is any vector of non negative integers, indexed by nodes of F, we write $t^w := \prod_{v \in \mathcal{V}_F} t_v^{w_v}$. Consider the power series

$$g := \prod_{E \in \mathcal{E}} \frac{1}{1 - t^{\chi_E}}.$$

We define integers $\Phi_\chi(w)$, one for each integer-valued vector w as above, by expanding the power series g,

$$g =: \sum_{v \in \mathbb{N}^{\mathcal{V}_F}} \Phi_\chi(w) t^w.$$

Φ_χ is called a *vector partition function*, see Sturmfels (1995). Note that labelings of graphs G with the same nodes and edges as F correspond to vectors w as above via

$$w = (\mathcal{L}_G(v))_{v \in \mathcal{V}_F}.$$

We denote by G_w the labeled graph G with the same nodes and edges as F and labeling given by w.

PROPOSITION A.1. *(i) Given any vector $w \in \mathbb{N}^{\mathcal{V}_F}$, the coefficient $\Phi_\chi(w)$ vanishes unless the labeled graph G_w is standard.*

(ii) If the labeled graph G_w is standard, the coefficient $\Phi_\chi(w)$ equals the number of standard decompositions of F.

PROOF. We expand each term $\frac{1}{1-t^{\chi_E}}$ in the product expression of g as a geometric series,

$$g = \prod_{E \in \mathcal{E}} (1 + t^{\chi_E} + t^{2 \cdot \chi_E} + t^{3 \cdot \chi_E} + t^{4 \cdot \chi_E} + \ldots).$$

Upon expanding the product, we see that each monomial appearing in the series takes the shape $m = \prod_{E \in \mathcal{F}} t^{n_E \cdot \chi_E}$ for some finite $\mathcal{F} \subseteq \mathcal{E}$ and some $n_E \in \mathbb{N}$. We replace the set $\mathcal{F}$ by the multiset $\mathcal{H}$ in which each $E \in \mathcal{F}$ appears n_E times. Since each member of $\mathcal{H}$ is a standard 0-1 subgraph of F, the graph $G := \sum \mathcal{H}$ standard graph and has the same nodes and edges as F. The above monomial m equals $\prod_{v \in \mathcal{V}_F} t^{\mathcal{L}_G(v)}$. This establishes (1).

As for (2), let G be a standard graph with the same nodes and edges as $\mathcal{V}_F$. The above discussion shows that the coefficient of the monomial $m := \prod_{v \in \mathcal{V}_F} t^{\mathcal{L}_G(v)}$ shows up in the expansion of g, and its coefficient counts the number of ways of writing G as a sum $G = \sum \mathcal{H}$ of elements of $\mathcal{E}$. This is just the number of standard decompositions of G. □

B. Partitions of partitions

Appendix A suggests a connection between standard decompositions and partitions. Let us further investigate this.

Example B.1.
- The set of partitions of an integer n is in natural bijection with the set of standard sets of cardinality n by identifying a partition and its Young diagram (in the French notation).
- If $p = \{n_1, \ldots, n_h\}$ (a multiset) is a partition of n and for each i, p_i is a partition of n_i, we call $\{p_1, \ldots, p_h\}$ a *partition of partition of* n. The set of partitions of partitions of n is in natural bijection with the set of standard sets $\Delta \subseteq \mathbb{N}^3$ of cardinality n, *together with all their C4 decompositions.*

◊

Both bijections are visualized in Figure B.1. For generalizing the statements, we introduce the notion of *C4 games.*

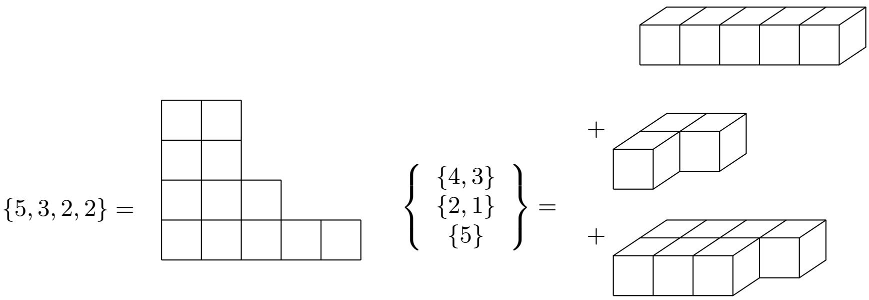

Figure B.1: Partitions (of partitions, resp.) and C4 games in $\mathbb{N}^2$ (in $\mathbb{N}^3$, resp.) correspond to each other

Definition B.2 (Iterated partition). *Let n be a positive integer. We recursively define a q-*fold iterated partition of n *as follows:*

- *for $q = 1$, it is a partition of n, that is, a multiset $p = \{n_1, \ldots, n_h\}$ of positive integers such that $\sum n_i = n$;*
- *for $q > 1$, it is a multiset $p = \{p_1, \ldots, p_h\}$ of $(q-1)$-fold iterated partitions of integers $n_1, \ldots, n_h$ such that $\sum n_i = n$.*

In other words, we look at all partitions of n into n_i, together with all partitions of all parts n_i into $n_{i,j}$, together with all partitions of all parts $n_{i,j}$ into $n_{i,j,k}$, etc.

Definition B.3 (C4 game). *Let n be a positive integer. We recursively define a* C4 game of size n in $\mathbb{N}^d$ *as follows:*

- *for $d = 1, 2$, it is standard set $\Delta \subseteq \mathbb{N}^d$ of cardinality n;*
- *for $d > 2$, it is a multiset $\{g_1, \ldots, g_h\}$ of C4 games of respective sizes n_i in $\mathbb{N}^{d-1}$ such that $\sum n_i = n$.*

In other words, we look at all standard sets $\Delta \subseteq \mathbb{N}^d$ of a cardinality n, together with all C4 decompositions of Δ into $\Delta_i \subseteq \mathbb{N}^{d-1}$, together with all C4 decompositions of all Δ_i into $\Delta_{i,j} \subseteq \mathbb{N}^{d-2}$, together with all C4 decompositions of all $\Delta_{i,j}$ into $\Delta_{i,j,k} \subseteq \mathbb{N}^{d-3}$, etc.

Proposition B.4. *For all $d, n \in \mathbb{N}$, there is a natural bijection*

$$f_d : \{(d-1)\text{-fold iterated partitions of } n\} \to \left\{\text{C4 games of size } n \text{ in } \mathbb{N}^d\right\}.$$

Proof. The assertion is obvious for $d = 1, 2$. For $d > 2$, the bijection f_d sends each multiset $\{H_1, \ldots, H_l\}$ of $(d-2)$-fold iterated partitions of integers $n_1, \ldots, n_l$ to the multiset $\{f_{d-1}(H_1), \ldots, f_{d-1}(H_l)\}$. □

Note that the bijection is only natural up to the choice of coordinate axes in $\mathbb{N}^d$. In other words, replacing the tuple $(e_1, \ldots, e_d)$ of standard basis elements by $(e_{\sigma(1)}, \ldots, e_{\sigma(d)})$ for some permutation σ induces an automorphism of the source of bijection f_d. For $d = 2$, this corresponds to the ambiguity between a partition and its transpose.

References

David Eisenbud (1995). *Commutative algebra*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York. ISBN 0-387-94268-8; 0-387-94269-6, xvi+785 . With a view toward algebraic geometry.

Laurent Evain (2005). On the postulation of s^d fat points in $\mathbb{P}^d$. *J. Algebra* **285**(2), 516–530. ISSN 0021-8693. URL http://dx.doi.org/10.1016/j.jalgebra.2004.09.034.

André Hirschowitz (1985). La méthode d'Horace pour l'interpolation à plusieurs variables. *Manuscripta Math.* **50**, 337–388. ISSN 0025-2611. URL http://dx.doi.org/10.1007/BF01168836.

M. Lederer (2008). The vanishing ideal of a finite set of closed points in affine space. *J. Pure Appl. Algebra* **212**, 1116–1133.

Mathias Lederer (2014). Components of Gröbner strata in the Hilbert scheme of points. *Proc. Lond. Math. Soc. (3)* **108**(1), 187–224. ISSN 0024-6115. URL http://dx.doi.org/10.1112/plms/pdt018.

Hiraku Nakajima (1999). *Lectures on Hilbert schemes of points on surfaces*, volume 18 of *University Lecture Series*. American Mathematical Society, Providence, RI. ISBN 0-8218-1956-9, xii+132 .

Bernd Sturmfels (1995). On vector partition functions. *J. Combin. Theory Ser. A* **72**(2), 302–309. ISSN 0097-3165. URL http://dx.doi.org/10.1016/0097-3165(95)90067-5.

Manuscript received 05 April 2013

Laurent Evain
Université d'Angers
Faculté des Sciences
Département de mathématiques
2, Boulevard Lavoisier
49045 Angers Cedex 01
FRANCE
laurent.evain@univ-angers.fr

Mathias Lederer
Department of Mathematics
Cornell University
Malott Hall
Ithaca
New York 14853
USA
mlederer@math.cornell.edu

Bjarke Hammersholt Roune
Department of Mathematics
Cornell University
Malott Hall
Ithaca
New York 14853
USA
bjarke.roune@gmail.com