Reducibility of convex cones

Michael Orlitzky

Towson University
February 25th, 2025

Background: convex cones

The setting for everything is either

$\mathbb{R}^n$ – ordered lists of n real numbers, or
$\mathbb{Q}^n$ – so we can do linear algebra on a computer.

In either case we treat them as vector spaces where we can add, scale, or take the dot product of vectors.

(The pictures will usually be of $\mathbb{R}^2$.)

Definition (convex cones).

If K is some subset of $\mathbb{R}^n$, then

- K is convex if the line segment between any two points of K is also in K.
- K is a cone if $\alpha K \subseteq K$ for all $\alpha \geq 0 \in \mathbb{R}$.

If K is already a cone, then K is a convex cone if and only if $x + y \in K$ whenever $x, y \in K$.

Example. A hexagon in $\mathbb{R}^2$ is convex.

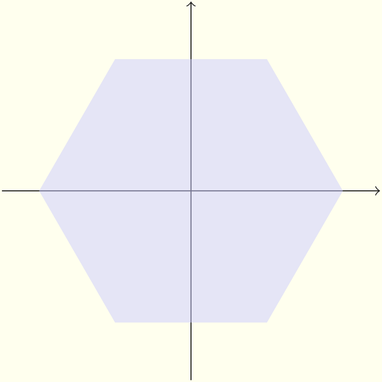

Example. A “ball” in $\mathbb{R}^2$ is convex.

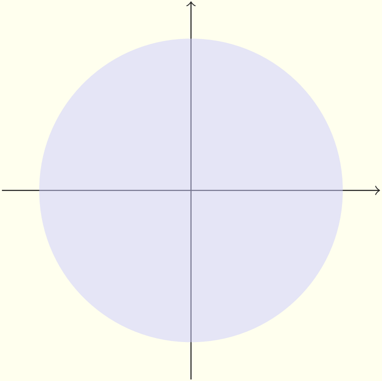

Example. The $(+,+)$ quadrant $\mathbb{R}^2_+$ is a convex cone.

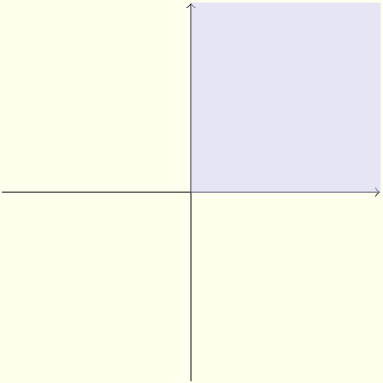

Example. $\mathbb{R}^2_+ \cup \left(-\mathbb{R}^2_+\right)$ is a non-convex cone.

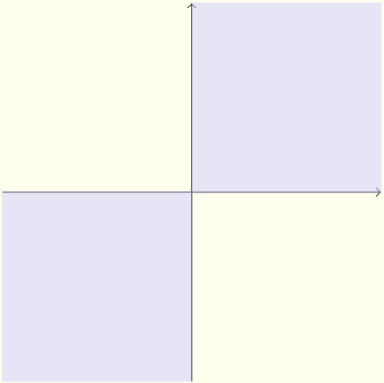

Example. The “ice cream” cone is non-convex.

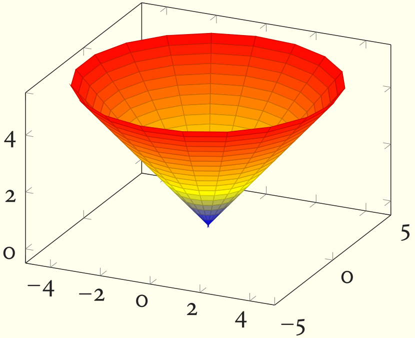

(It would be convex if it contained ice cream.)

Example. $\mathbb{R}^2$ itself is a convex cone.

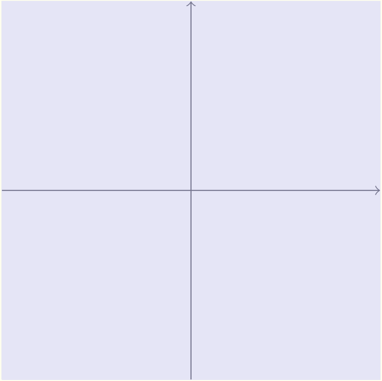

Why care?

As the last example shows, convex cones generalize vector spaces:

- The elements of a vector space can be scaled in any direction (positively or negatively).
- The elements of a convex cone can only be scaled positively.

Every vector space is a convex cone, but not every convex cone is a vector space.

Convex cones arise naturally in optimization where you are trying to minimize some function within a set, and there is a *good direction* you want to go in.

- OK to move positively in the good direction, bad to move away from it.
- The directions you're allowed to move in will also form a convex cone (if your set is nice).
- Doing both at the same time leads to a third, smaller cone of directions that are both good and allowable.

Example (search directions).

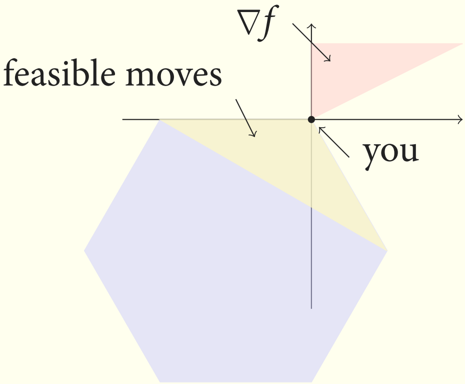

Example (conic programming).

If K_1, K_2 are convex cones, if A is linear, and if b, c are two points, then

$$\begin{array}{rr} \text{minimize} & \langle b, x \rangle \\ \text{subject to} & A\left(x\right) \in K_2 + c \\ & x \in K_1 \end{array}$$

can solve NP-hard problems:

- The quadratic assignment problem
- Maximum stable set of a graph
- Chromatic number of a graph

Definition (conic hull, generators).

Convex cones are usually generated by a subset.

The *conic hull* is the span, but with all coefficients ≥ 0:

$$\operatorname{cone}(G) := \left\{ \sum_{i=1}^{m} \alpha_i g_i \;\middle|\; g_i \in G,\ \alpha_i \geq 0, m \in \mathbb{N} \right\}.$$

If $\operatorname{cone}(G) = K$ we say that K is generated by G and the elements of G are generators of K.

Definition (cone species).

A convex cone is…

- *Closed* if it is topologically closed as a set. For simplicity, assume all cones are closed.
- *Pointed* if it contains no lines.
- *Solid* if it spans the ambient space.
- *Proper*, if all of the above.

Lemma (characterizations of pointedness).

TFAE for a closed convex cone K:

- K is pointed
- $K \cap (-K) = \{0\}$
- $K \setminus \{0\}$ lives in a strict half-space of $\mathbb{R}^n$

The last item means that K is generated by a cross-section of itself.

Example. Ice cream cones are pointed,

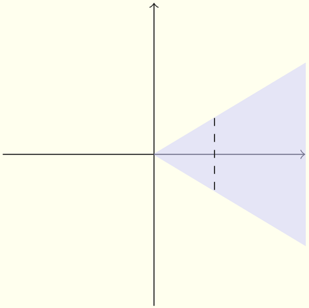

and are generated by any cross section (where the ice cream is one inch deep, for example).

Example. A cone that contains the *y*-axis is not.

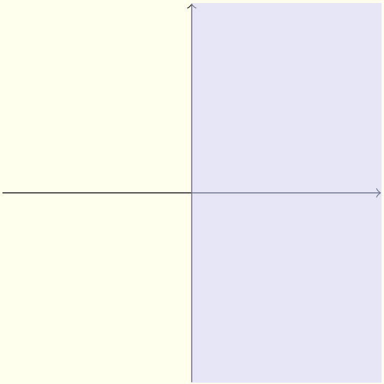

Example (nonnegative orthant).

If $\mathbf{e}$ is the standard basis in $\mathbb{R}^n$, then

$$\begin{aligned}\mathbb{R}^n_+ &:= \left\{ \left(x_1, x_2, \ldots, x_n \right)^T \;\middle|\; x_i \geq 0 \text{ for all } i \right\} \\ &= \operatorname{cone}\left(\mathbf{e} \right)\end{aligned}$$

is a proper cone.

Example (vector subspace).

If **e** is a basis for a vector subspace V of $\mathbb{R}^n$, then $V = \operatorname{cone}\left(\left\{\pm e_i \mid e_i \in \mathbf{e}\right\}\right)$ is pointed/solid only in trivial cases:

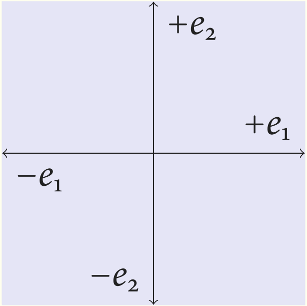

Example (PSD cone).

In the space $\mathcal{S}^n$ of real, symmetric $n \times n$ matrices, the PSD subset,

$$\mathcal{S}^n_+ \coloneqq \left\{ A \in \mathcal{S}^n \mid \langle Ax, x \rangle \geq 0 \text{ for all } x \in \mathbb{R}^n \right\}$$

is a proper cone generated by the matrices xx^T such that $x \in \mathbb{R}^n$.

(This is the spectral decomposition.)

Example (Cartesian products).

If K_1 and K_2 are both convex cones, then $K_1 \times K_2$ is too.

If G_1 and G_2 generate K_1 and K_2, then

$$\left(G_1 \times \left\{0\right\}\right) \cup \left(\left\{0\right\} \times G_2\right)$$

generates $K_1 \times K_2$.

Example (linear images).

If K is a convex cone and if A is invertible, then $A(K)$ is a convex cone.

If G generates K, then $A(G)$ generates $A(K)$.

We can combine these two tricks to create many new convex cones $A(K_1 \times K_2)$ from existing K_1, K_2.

Question.

But what if we are *given* $A\left(K_1 \times K_2\right)$, is there a way to recover K_1 and K_2?

Definition (informal).

A convex cone is *reducible* if it is the nontrivial direct sum of two convex cones, and *irreducible* if not.

Reducibility

Definition (formal).

A convex cone $K \subseteq \mathbb{R}^n$ is *reducible* if and only if…

- There exist subspaces $V_1, V_2 \neq \{0\}$ such that $\mathbb{R}^n = V_1 + V_2$ and $V_1 \cap V_2 = \{0\}$.
- There exist convex cones $K_1 \subseteq V_1$ and $K_2 \subseteq V_2$ such that $K = K_1 + K_2$.

The first part just says that $\mathbb{R}^n$ is a direct sum of V_1 and V_2, and the second part says that K splits along that direct sum.

Applications of reducibility:

- The book by Berman and Plemmons[†].
- The automorphism group of an EJA acts transitively on the set of Jordan frames if and only if the cone of squares is irreducible or $\mathbb{R}^n_+$.
- If K is a proper polyhedral cone, its Lyapunov rank is one if and only if it is irreducible.
- If K is a proper polyhedral cone in $\mathbb{R}^n$, then K is irreducible if and only if $\dim\left(\mathbf{Z}\left(K\right)^*\right) \geq n^2 - 1$.

[†] Nonnegative Matrices in the Mathematical Sciences, 1994.

And in general:

It’s better to have two small simple things than one big complicated thing.

(Most algorithms have superlinear complexity, so it’s faster to do half the work twice.)

If you start with a reducible cone and reduce it repeatedly, you end up with a direct sum of irreducible cones.

The factors are unique (up to order) if the original cone was proper. This result has been proved many times in many different settings:

- E.B. Vinberg, 1963.
- Bleicher and Schneider, 1964.
- Rothaus, 1966.

$\vdots$

- Kaneyuki, 1967.
- Asano, 1968.
- Gruber, 1970.
- Gale and Klee, 1975.
- Meschiari (Gentili and O'Connor), 1982.
- Faraut and Korányi, 1994.
- Hauser and Güler, 2002.
- Rohrer, 2011.
- Bremner, Sikirić, Pasechnik, Rehn and Schürmann, 2014.

In the best case, when K is proper, we obtain a unique decomposition into irreducible factors $K_1, K_2, \ldots, K_r$:

$$\begin{array}{c} K = \bigoplus_{i=1}^{r} K_i \\ \subseteq \\ \mathbb{R}^n = \bigoplus_{i=1}^{r} V_i. \end{array}$$

The order doesn't matter, but as a set the factors are unique.

Example (proper, unique).

If $\mathbf{e} = \{e_1, e_2, \ldots, e_n\}$ is the usual basis, then $\mathbb{R}^n_+ = \operatorname{cone}(\mathbf{e})$ has a unique decomposition:

$$\begin{aligned} \mathbb{R}^n_+ &= \bigoplus_{i=1}^{n} \operatorname{cone}(\{e_i\}) \\ &\subseteq \\ \mathbb{R}^n &= \bigoplus_{i=1}^{n} \operatorname{span}(\{e_i\}). \end{aligned}$$

Example (non-pointed, non-unique).

As a subset of itself, $\mathbb{R}^2$ is solid but not pointed.

It can be decomposed, but not uniquely, because any basis for $\mathbb{R}^2$ forms a new decomposition:

$$\mathbb{R}^2 = \operatorname{span}\left(\left\{(1,0)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(0,1)^T\right\}\right),$$
$$\mathbb{R}^2 = \operatorname{span}\left(\left\{(1,1)^T\right\}\right) \oplus \operatorname{span}\left(\left\{(1,-1)^T\right\}\right),$$

and so on.

Example (non-solid, non-unique).

$K = \{0\}$ is a pointed convex cone in $\mathbb{R}^2$ that can be decomposed, but not uniquely, because it is not solid.

We can always write

$$K = \{0\} \oplus \{0\} \subseteq V_1 \oplus V_2$$

but as in the previous example we have many choices for the subspaces V_1 and V_2.

Definition (lineality space).

The *lineality space* of a convex cone K is defined to be $\operatorname{linspace}(K) \coloneqq K \cap (-K)$, and is the largest subspace contained in K.

The dimension of $\operatorname{linspace}(K)$ measures "how not pointed" the cone K is, and the dimension of $\operatorname{span}(K)^{\perp}$ measures "how not solid" it is:

- K is pointed if and only if $\operatorname{linspace}(K) = \{0\}$
- K is solid if and only if $\operatorname{span}(K)^{\perp} = \{0\}$

Theorem (well known).

Every closed convex cone K has an orthogonal direct sum decomposition,

$$K = \operatorname{linspace}(K) \oplus \left[K \cap \operatorname{linspace}(K)^{\perp}\right],$$

where $K \cap \operatorname{linspace}(K)^{\perp}$ is pointed.

Proof.

Wets & Witgall showed in 1967 that every generating set G of K contains a basis for $\operatorname{linspace}(K)$.

Proof (cont'd).

Take one, and orthogonalize it:

$$\begin{aligned}
b_1 &\leftarrow b_1 \\
b_2 &\leftarrow \|b_1\|^2 \, b_2 - \langle b_2, b_1 \rangle \, b_1 \\
b_3 &\leftarrow \|b_1\|^2 \, \|b_2\|^2 \, b_3 - \|b_1\|^2 \, \langle b_3, b_2 \rangle \, b_2 - \|b_2\|^2 \, \langle b_3, b_1 \rangle \, b_1 \\
&\vdots
\end{aligned}$$

This is essentially Gram-Schmidt without normalization that works over the *integers*.

Proof (cont'd).

Afterwards, do the same thing with the remaining $g \in G$ to make them orthogonal to $\operatorname{linspace}(K)$. If L is the orthogonal basis from the previous slide, then

$$g \leftarrow \left(\prod_{\ell \in L} \|\ell\|^2 \right) g - \sum_{\ell \in L} \left(\prod_{\substack{m \in L \\ m \neq \ell}} \|m\|^2 \right) \langle g, \ell \rangle \, \ell.$$

The cone generated by all such g is still K, because we can recover the original G using $\operatorname{cone}(\cdots)$. □

Theorem (still more or less known).

Every convex cone K has a unique decomposition $K = L \oplus Z \oplus P$ where

$$\begin{aligned} L &= \operatorname{linspace}(K) \text{ in itself} \\ Z &= \{0\} \text{ in } \operatorname{span}(K)^{\perp} \\ P &= \bigoplus_{i=1}^{r} P_i \text{ in } \operatorname{span}(K) \cap \operatorname{linspace}(K)^{\perp} \end{aligned}$$

and each P_i is proper and irreducible in $\operatorname{span}(P_i)$.

Proposition.

If $K = K_1 \oplus K_2$ is pointed in $V_1 \oplus V_2$, then the generators of K_i are those of K that live in V_i.

Proof.

Suppose $K = \operatorname{cone}(G)$. One direction is easy: if $x \in \operatorname{cone}(G \cap V_1)$, then $x \in K_1$.

Proof (cont'd).

Conversely, let $x \in K_1$. Write each $g \in G \subseteq K$ as $p + q$ with $p \in K_1$ and $q \in K_2$ so that $x = \sum_{i=1}^{m} \alpha_i \left(p_i + q_i\right)$. Group these,

$$x = \underbrace{\left(\sum_{i=1}^{m} \alpha_i p_i\right)}_{\in V_1} + \underbrace{\left(\sum_{i=1}^{m} \alpha_i q_i\right)}_{\in V_2}.$$

Since $x \in V_1$, the second sum must be zero. But K_2 is pointed (K is pointed), so all q_i must be zero. Thus $g_i = p_i \in V_1$ for all g_i used to express x. $\square$

This result will be the key to decomposing a given convex cone in practice.

To reiterate:

If $K = \operatorname{cone}(G)$ is pointed, the elements of G can be grouped into distinct generating sets G_i of the irreducible factors of $K = \bigoplus K_i$.

Theorem.

If $K = \operatorname{cone}(G)$ and if $G \subseteq \mathbb{Q}^n$ is finite, then there is an algorithm to compute the decomposition

$$K = L \oplus Z \oplus \left(\bigoplus_{i=1}^{r} P_i \right).$$

Proof.

G contains a basis for $L = \operatorname{linspace}(K)$. Use the Gram-Schmidt trick to obtain a set of generators H for the pointed component $P = \bigoplus_{i=1}^{r} P_i$.

Proof (cont'd).

H contains a basis for $\operatorname{span}(P)$. Place the $h \in H$ in a matrix as the columns, and use Gaussian elimination to reduce it to row-echelon form. The pivot columns will then identify a linearly-independent subset of H, i.e. a basis $\left\{h_{i_1}, h_{i_2}, \ldots, h_{i_m}\right\}$ for $\operatorname{span}(P)$.

At this point, stop and compute $\operatorname{span}(K)^{\perp}$, the orthogonal complement of $L \oplus \operatorname{span}(P)$. This is the subspace that $Z = \{0\}$ lives in.

Proof (cont'd).

It remains to find the components P_i. Express each $h_i \in H$ in terms of our new basis:

$$h_i = \alpha_1(h_i)\,h_{i_1} + \underbrace{\alpha_2(h_i)}_{\in\mathbb{R}}\,h_{i_2} + \cdots + \alpha_m(h_i)\,h_{i_m}.$$

Build a graph (V, E) with vertices $V \coloneqq H$ and edges

$$E \coloneqq \left\{ \left(h_i, h_{i_k}\right) \mid \alpha_k(h_i) \neq 0 \right\}.$$

(This trick is due to Bremner, Sikirić, Pasechnik, Rehn and Schürmann.)

Proof (cont'd).

Basis and direct-sum representations are unique, so this is well-defined.

The edges of the graph indicate whether or not each generator h_i of P has a nonzero h_{i_k} component.

From the preceding proposition, there is a path from h_i to h_{i_k} if and only if h_i and h_{i_k} are generators of the same irreducible factor of P.

Proof (cont'd).

Every graph can be partitioned into connected components. For finite graphs, this is a standard depth-first search algorithm.

Do it: find the vertices H_i in each connected component of the graph (V, E).

Set $P_i \coloneqq \operatorname{cone}(H_i)$, and we are done. $\square$

This algorithm has been implemented in SageMath. The simplest example is of decomposing $\mathbb{R}^n_+$, the nonnegative orthant, into n rays.

```
sage: K = cones.nonnegative_orthant(3)
sage: K
3-d cone in 3-d lattice N
sage: K.irreducible_factors()
{1-d cone in 3-d lattice N,
 1-d cone in 3-d lattice N,
 1-d cone in 3-d lattice N}
```

If decomposing it does nothing, your cone is irreducible. For example, the ℓ_1-norm cone.

```
sage: K = Cone([ (x,y,1) for x in [-1,1]
....:                    for y in [-1,1] ])
sage: K
3-d cone in 3-d lattice N
sage: K.irreducible_factors()
{3-d cone in 3-d lattice N}
sage: K.is_reducible()
False
```

Automorphisms

Definition (automorphisms).

An automorphism of a convex cone K is any invertible linear map A such that $A(K) = K$.

The set of all automorphisms of K form a group, $\operatorname{Aut}(K)$.

Automorphisms are not usually easy to find, but in some cases it can be done and it happens to be an application of reducibility.

Definition (polyhedral cone).

A convex cone K is *polyhedral* if $K = \operatorname{cone}(G)$ for a finite set G, or (equivalently) if K is the intersection of a finite number of half-spaces.

These equivalent conditions are known as the $\mathcal{V}$- and $\mathcal{H}$-representations, for vertex ($\mathcal{V}$) and half-space ($\mathcal{H}$). Often one is more useful than the other. Using both at the same time is the *double-description* method.

Proposition.

If K is a pointed polyhedral cone and if G is a **minimal** generating set of K, then for all $g \in G$, there exists a $\tilde{g} \in G$ such that $A\left(\operatorname{cone}\left(\{g\}\right)\right) = \operatorname{cone}\left(\{\tilde{g}\}\right)$.

The “minimal” here is important, because nothing is stopping you from adding redundant elements to a generating set.

Example.

Let $G_1 \coloneqq \{e_1, e_2, e_3\}$ be the standard basis in $\mathbb{R}^3$ so that $\mathbb{R}^3_+ = \operatorname{cone}(G_1)$. Suppose we add $e_1 + e_2$ to this set, and call the result G_2. Then $\mathbb{R}^3_+ = \operatorname{cone}(G_2)$ as well, but G_2 contains a redundant entry.

Apropos of automorphisms, define $A : \mathbb{R}^3 \to \mathbb{R}^3$ by

$$A(e_1) = e_2, \quad A(e_2) = e_3, \quad A(e_3) = e_1.$$

This A belongs to $\operatorname{Aut}(\mathbb{R}^3_+)$, and it permutes G_1, but not G_2 because $A(e_1 + e_2) \notin G_2$.

The moral of the story is:

1. Ignoring scaling, automorphisms of a pointed convex cone permute a minimal generating set.
2. For polyhedral cones, those sets are finite.
3. The permutations of a finite set are again finite.

If something is finite, you can enumerate it.

(The possible *scalings* are not finite, but with a little extra work[†] they too can be found.)

[†] Computing symmetry groups of polyhedra by Bremner et al. (2014)

If K is not pointed or not solid, it's not a big deal in this context.

The automorphisms of $L = \operatorname{linspace}(K)$ and $Z = \{0\}$ are just the invertible linear maps on $\operatorname{linspace}(K)$ and $\operatorname{span}(K)^{\perp}$, respectively.

So without much loss of generality, we now compute the automorphisms of a *proper* polyhedral cone K ignoring scaling.

Theorem.

If $K = \operatorname{cone}(G)$ is proper and if $G \subseteq \mathbb{Q}^n$ is finite, then there is an algorithm to find the subset of $\operatorname{Aut}(K)$ that permutes G.

Proof.

Remove the redundant generators in G: an $\mathcal{H}$-representation makes it trivial, and in any case checking if one point is in the conic hull of the others is an LP feasibility problem. G is now minimal.

Proof (cont'd).

Every n-element subset of G potentially determines a linear transformation A on $\mathbb{Q}^n$. We therefore consider every **pair** of n-element subsets $S_1, S_2 \subseteq G$.

If there is a solution A to the linear system $AS_1 = S_2$ that permutes G, we keep it; otherwise, we try the next (S_1, S_2). Any A that we keep is automatically invertible because G contains a basis for $\mathbb{Q}^n$. $\square$

This is not a deep algorithm, but it is relevant for two reasons:

1. It clearly demonstrates the benefit of reducing the cone if possible.
2. Extending it to find *all* automorphisms requires you to reduce the cone anyway.

To back up the first claim, we give an example.

Most of the time spent in the algorithm comes from enumerating the pairs of subsets (S_1, S_2). If G contains N elements, then there are $2\binom{N}{n}$ ways to choose (S_1, S_2).

Example.

Let $K = \operatorname{cone}(G)$ in $\mathbb{Q}^3$ be formed by raising the regular octagon in $\mathbb{Q}^2$ up to height one[†]. For this K we have $N = 8$ and $n = 3$, so there are $2\binom{8}{3} = 112$ ways to choose S_1 and S_2.

[†] An octagon won't have rational coordinates, but imagine something close.

Example (cont'd).

Now glue together two copies of K to form

$$J \coloneqq K \times K = \operatorname{cone}(H)$$

where

$$H \coloneqq (G \times \{0\}) \cup (\{0\} \times G).$$

If we naively run the algorithm on J, we get $N = 16$ and $n = 6$ for $2\binom{16}{6} = 16,016$ pairs of subsets.

Example (conclusion).

It is better to enumerate the 112 pairs for K twice than it is to enumerate the $16,016$ pairs for J.

(This is true even if we include the $N\left(n^2 + n\right) = 672$ operations needed to do the decomposition.)

(And even if we don't overlook the automorphisms that swap the two copies of K.)

Final observations:

- A futher benefit of reduction is that a change of basis can be used to make your vectors shorter.
- Everything we did for automorphisms (one cone) can be done for *isomorphisms* (two cones).
- Isomorphic cones have isomorphic face lattices. This is another opportunity to punt to a graph algorithm rather than get our hands dirty.

The end