Enumeration of Coherent Matching Fields Modulo Symmetry

Xavier Allamigeon and Yiyuan Chen

INRIA and CMAP, École polytechnique, IP Paris, CNRS
firstname.lastname@inria.fr

Abstract. Coherent matching fields, introduced by Sturmfels and Zelevinsky [21], play an important role in the study of the Grassmannian and its tropical counterpart. They can be identified with the vertices of a Minkowski sum of Birkhoff polytopes associated with the maximal minors of an $n \times r$ matrix.

We present an algorithm for enumerating coherent matching fields up to row and column permutations. Our approach adapts the reverse search method of Avis and Fukuda [7] to enumerate the orbits of the vertices of a polytope under a symmetry group. We further show how the shadow-vertex rule can be used to perform inexpensive pivots between coherent matching fields.

We implemented this algorithm in Rust with distributed parallelism via MPI. This implementation enabled us to enumerate the $n \times r$ coherent matching fields for $n \leq 10$ when $r = 3$ and for $n \leq 8$ when $r = 4$, up to row and column permutations. As an application, we show that there exist configurations of nine points in $\mathbb{R}^2$ in general position whose order type cannot be realized by configurations of tropical points in general position, highlighting a new discrepancy between classical and tropical realizability.

1 Introduction

Given two positive integers $n \geq r$, a *matching field* is a collection $\Lambda = (\Lambda_I)_{I \in \binom{[n]}{r}}$ assigning to every r-subset $I \subset [n]$ a bijection between I and $[r]$ (we denote $[k] \coloneqq \{1, \dots, k\}$, and by $\binom{[n]}{r}$ the set of r-subsets of $[n]$). A matching field Λ is *coherent* if there exists a weight matrix $W \in \mathbb{R}^{n \times r}$ such that, for all $I \in \binom{[n]}{r}$, Λ_I is the unique bijection $\lambda \colon I \to [r]$ maximizing $\sum_{i \in I} W_{i\lambda(i)}$. (We say that Λ is *realized* by the matrix W.)

Coherent matching fields were introduced by Sturmfels and Zelevinsky for the study of the Newton polytope $P_{n,r}$ of the product of maximal minors of an $n \times r$ matrix $(X_{ij})_{ij}$ of indeterminates [21]. Indeed, coherent matching fields are in one-to-one correspondence with the vertices of the polytope $P_{n,r} = \sum_{I \in \binom{[n]}{r}} B_I$, where, for each minor $I \in \binom{[n]}{r}$, $B_I \subset \mathbb{R}^{n \times r}$ is the polytope defined as the convex hull of the $n \times r$ 0/1-matrices associated with bijections from I to $[r]$. Every B_I is isomorphic to the Birkhoff polytope B_r.

Coherent matching fields have been used in the study of the Grassmannian and its toric degenerations [17,9]. They also appear in tropical geometry in the study of the tropical Stiefel map [11]. Coherent matching fields have a strong connection with the cell complex induced by arrangements of tropical hyperplanes, or, equivalently, regular triangulations of the product of simplices $\triangle_{n-r-1} \times \triangle_{r-1}$ [10,20]. The combinatorial types of the latter are in one-to-one correspondence with the *pointed* coherent matching fields, which constitute a $(n-r-1)(r-1)$-dimensional face of the $(n-1)(r-1)$-dimensional polytope $P_{n,r}$; see [11, Example 4.8].

In this paper, we are interested in understanding what arrangements of hyperplanes in general position can be realized by arrangements of tropical hyperplanes. This is motivated by the connection between the complexity of linear programming and that of mean payoff games via tropical linear programming, as studied in a series of works [4,5,6], which opens the way to the application of game algorithms to linear programming. The dual (and equivalent) question is whether a point configuration in general position can be combinatorially realized by a configuration of tropical points. The combinatorial type of a labelled point configuration is called an *order type*. Order types were introduced by Goodman and Pollack [13]. In the planar case, an order type records the orientation of every triple of points. The number of order types (up to relabeling) has been computed by Aichholzer et al. up to ten points in [1]. The following statement is the main outcome of the experimental part of the paper:

Theorem 1. *All order types of n points in $\mathbb{R}^2$ in general position are tropically realizable for $n \leq 8$, while 80 order types (up to relabeling) of nine points in $\mathbb{R}^2$ out of* 158,817 *are not.*

The proof relies on the enumeration of coherent matching fields. More precisely, the order type of every configuration of tropical points can be determined from a coherent matching field and the sign of the coordinates of these points; see Section 4. As the number of coherent matching fields grows extremely fast, we exploit the property that $n \times r$ coherent matching fields are closed under the group action of $\mathfrak{S}_n \times \mathfrak{S}_r$, i.e., permutations of indices of $[n]$ and $[r]$. One contribution of our paper is an algorithm based on the reverse search method [7] that enumerates the orbits of coherent matching fields. To this end, we adapt the approach of Fukuda [12] for the enumeration of vertices of Minkowski sums of polytopes to coherent matching fields modulo symmetry. We present a complete enumeration of coherent matching fields for $r = 3$ and $n \leq 10$, and $r = 4$ and $n \leq 8$. This is done by a parallel implementation of our approach in Rust based on MPI. We report in Table 1 the number of orbits for each size.

Related Work. Our contribution lies at the crossroads between the enumeration of vertices of Minkowski sums of polytopes, pioneered by [12] and implemented by Weibel [22], and that of symmetric polytopes; see e.g. [8] for a general discussion on dealing with symmetries in polytope representation problems. An important application of enumeration modulo symmetries is the computation of regular triangulations of (symmetric) polytopes, where the work of Jordan, Joswig and Kastner [15] and Rambau [19] constitute the current state-of-the-art; see also [18] for a seminal work.

The works [15,19] are also related to coherent matching fields. They provide the list of regular triangulations of $\triangle_2 \times \triangle_6$, $\triangle_3 \times \triangle_4$ and $\triangle_3 \times \triangle_5$, which correspond to pointed coherent matching fields of size 3×10, 4×9 and 4×10. At the end of Section 3, we explain that our algorithm can also be restricted to enumerate these pointed coherent matching fields.

Organization of the Paper. In Section 2, we provide a self-contained presentation on the reverse-search method for the enumeration of orbits of the vertices of a polytope under a group action. In Section 3, we define the different primitive functions needed by the algorithm in the case of coherent matching fields. Section 4 deals with the application to the enumeration of tropical order types which allows us to establish Theorem 1. Proofs are omitted from the proceedings version of the paper for reasons of space. (An extended version, including proofs, is available on arXiv.) The implementation is available at `https://gitlab.inria.fr/allamige/cmf-enumeration`. The lists of coherent matching fields and tropical order types can be found in the `data` folder.

2 Reverse Search Method for Symmetric Polytopes

Let $P \subset \mathbb{R}^d$ be a polytope with vertex-edge graph $\Gamma(P) = (V(P), E(P))$, where $V(P)$ is the vertex set of P and $E(P)$ the set of edges of P. Let G be a finite group acting on $\mathbb{R}^d$ by affine automorphisms with $G(P) = P$. Then every $g \in G$ induces a graph automorphism of $\Gamma(P)$, so G acts on $V(P)$ preserving adjacency.

We assume that P is equipped with the two essential ingredients required by the reverse search method, namely a local search function and an adjacency oracle. The *local search function*

Table 1. Number of orbits of coherent matching fields under the action of $\mathfrak{S}_n \times \mathfrak{S}_r$.

n	3	4	5	6	7	8	9	10
$r = 3$	1	2	7	66	891	22,582	888,614	51,863,755
$r = 4$		1	7	116	21,758	16,323,593	—	—

1: **procedure** SymReverseSearch
2: $S := \Delta$, $C := \varnothing$, $v := v^\star$
3: **output** v
4: **loop**
5: **if** NextChild(v, S, C) is $(w, \cdot)$
6: $v := w$, $S := \Delta$, $C := \varnothing$
7: **output** v
8: **else if** $v \neq v^\star$
9: $w := v$, $v := f_\phi(v)$
10: $(S, C) :=$ Visited(v, w)
11: **else return**

1: **procedure** NextChild(v, S, C)
2: **while** $S \neq \varnothing$ **do**
3: $\delta := \mathrm{pop}(S)$, $w := \mathrm{Adj}_\phi(v, \delta)$
4: **if** $w \notin C \cup \{\perp\}$ and $f_\phi(w) = v$
5: **return** (w, S)
6: **return** $\perp$

1: **procedure** Visited(v, w)
2: $S := \Delta$, $C := \varnothing$
3: **repeat**
4: $(w', S) :=$ NextChild(v, S, C)
5: $C := C \cup \{w'\}$
6: **until** $w' = w$
7: **return** (S, C)

Fig. 1. The reverse search method for enumerating the orbits of $V(P)$ under the action of G.

is a map $f \colon V(P) \setminus \{v^\star\} \to V(P)$ where $v^\star \in V(P)$ is a distinguished vertex called the *root*, and, for every $v \in V(P) \setminus \{v^\star\}$, $f(v)$ is adjacent to v and there is k such that $f^k(v) = v^\star$. Usually, this function is associated to a ranking function $\rho \colon V(P) \to \mathbb{R}$ such that $\rho(f(v)) > \rho(v)$ for all $v \in V(P) \setminus \{v^\star\}$. The ranking function is commonly chosen as a linear function $\langle c^\star, \cdot \rangle$ where $c^\star$ lies in the interior of the normal cone of the root vertex $v^\star$. The adjacency oracle provides, for every vertex $v \in V(P)$, the complete list of its neighbors in $\Gamma(P)$. It usually takes the form of a map $\mathrm{Adj} \colon V(P) \times \Delta \to V(P) \cup \{\perp\}$ where Δ is a finite superset of indices of the adjacent vertices, and $\perp$ is a special value returned when the index does not correspond to an adjacent vertex. For every vertex w adjacent to v, there is $\delta \in \Delta$ such that $w = \mathrm{Adj}(v, \delta)$.

A natural approach to enumerate the G-orbits of the vertices of P is to specify a function ϕ mapping every vertex to a canonical representative of its orbit. We call *normalized vertices* the set of such representatives, i.e., the elements of $\phi(V(P))$. We denote by $\Gamma_\phi(P) = (V_\phi(P), E_\phi(P)) := (\phi(V(P)), \phi(E(P)))$ the image of the graph $\Gamma(P)$ under the map ϕ. It is isomorphic to the quotient graph. We define the adjacency oracle for $\Gamma_\phi(P)$ by $\mathrm{Adj}_\phi := \phi \circ \mathrm{Adj}$ (by defining $\phi(\perp) := \perp$). We can prove that for every vertex $v \in V_\phi(P)$ and neighbor w of v in $\Gamma_\phi(P)$, there exists $\delta \in \Delta$ such that $w = \mathrm{Adj}_\phi(v, \delta)$. In order to specify the local search function for $\Gamma_\phi(P)$, we make the following assumption:

Assumption 1. For all $v \in V_\phi(P)$, $\rho(v) = \max\{\rho(v') \colon v' = v \mod G\}$.

Under this assumption, the root $v^\star$ is a normalized vertex (recall that $v^\star$ is the unique vertex of P maximizing the map ρ). We then define the local search function for $\Gamma_\phi(P)$ by $f_\phi := \phi \circ f$. We readily get that $\rho(f_\phi(v)) > \rho(v)$ for all $v \in V_\phi(P) \setminus \{v^\star\}$. Moreover, v and $f_\phi(v)$ are adjacent vertices in $\Gamma_\phi(P)$ as v and $f(v)$ are adjacent in $\Gamma(P)$, $\phi(v) = v$ and $f_\phi(v) = \phi(f(v))$.

The resulting reverse search method is given in Figure 1. It consists in recursively exploring a spanning tree of $\Gamma_\phi(P)$ with root vertex $v^\star$, where a vertex w is defined as a child of a vertex v if the equality $f_\phi(w) = v$ holds. The variable S is a stack used to iterate over the elements of the finite set Δ. The main difference with the presentation of [12] lies in the use of an additional set C to collect the already visited neighbors of v. This is required by the fact that several vertices w adjacent to v in $\Gamma(P)$ may belong to the same orbit. More precisely, the auxiliary function NextChild(v, S, C) pops elements from S until a normalized vertex w that has not been already visited and that satisfies $f_\phi(w) = v$ is found. If there is none, it returns the special value $\perp$. In this case, if in addition $v \neq v^\star$, the algorithm backtracks to the parent vertex $f_\phi(v)$ (Lines 8 to 10 in SymReverseSearch). More precisely, after setting w to v and v to $f_\phi(v)$, it calls the function Visited to rebuild the set C of neighbors of v visited up to the child vertex w, as well as the stack S of indices. In this way, our reverse search method remains a memoryless enumeration algorithm, since the set of already explored children does not need to be kept in memory during the exploration of the next children. In the case where NextChild$(v, S, C) = \perp$

(no more child) and $v = v^\star$, the exploration of the spanning tree is complete and the algorithm returns; this is done at Section 2 in SymReverseSearch.

3 Enumeration of CMF modulo $\mathfrak{S}_n \times \mathfrak{S}_r$

We denote by $\mathrm{CMF}(n, r)$ the set of coherent matching fields. Given $\Lambda \in \mathrm{CMF}(n, r)$, we denote by $v(\Lambda)$ the associated vertex of $P_{n,r}$, i.e., $v(\Lambda) \coloneqq \sum_{I \in \binom{[n]}{r}} E(\Lambda_I)$ where $E(\lambda)_{ij} \coloneqq 1$ if $\lambda(i) = j$, 0 otherwise. The one-to-one correspondence between $\mathrm{CMF}(n, r)$ and the vertices of $P_{n,r}$ relies on the fact that the interior of the normal cone of $v(\Lambda)$ precisely consists of the weight matrices $W \in \mathbb{R}^{n\times r}$ realizing the coherent matching field Λ (see Section 1). Consequently, for enumeration purposes, we use coherent matching fields and the corresponding vertices of $P_{n,r}$ interchangeably.

The set $\mathrm{CMF}(n, r)$ is equipped with the action of the group $\mathfrak{S}_n \times \mathfrak{S}_r$ as follows. Given $(\sigma, \tau) \in \mathfrak{S}_n \times \mathfrak{S}_r$ and $\Lambda \in \mathrm{CMF}(n, r)$, $(\sigma, \tau) \cdot \Lambda$ is the coherent matching field formed by the matchings $\{(\sigma(i), \tau(j)) \colon (i, j) \in \Lambda_I\}$ for all $I \in \binom{[n]}{r}$. In particular, if $W \in \mathbb{R}^{n\times r}$ realizes Λ, then the matrix obtained from W by permuting its rows by σ and its columns by τ realizes $(\sigma, \tau) \cdot \Lambda$. We denote by $(\sigma, \tau) \cdot W$ the permuted weight matrix, noting that $v\big((\sigma, \tau) \cdot \Lambda\big) = (\sigma, \tau) \cdot v(\Lambda)$. We point out that the action of $\mathfrak{S}_n \times \mathfrak{S}_r$ on $\mathrm{CMF}(n, r)$ is not free. For instance, the *all-diagonal* coherent matching field, which maps every minor $I = \{i_1 < \cdots < i_r\} \in \binom{[n]}{r}$ to $1, \dots, r$, is invariant under the action of reversing rows and columns (i.e., the pair $(i \mapsto n + 1 - i$, $j \mapsto r + 1 - j))$. In contrast, $\mathfrak{S}_n$ and $\mathfrak{S}_r$ both act freely on $\mathrm{CMF}(n, r)$. This is because, for every $\Lambda \in \mathrm{CMF}(n, r)$, the rows (resp. columns) of $v(\Lambda)$ are pairwise distinct.

We now instantiate the symmetric reverse-search framework of Section 2 for the polytope $P_{n,r}$ under the action of $G = \mathfrak{S}_n \times \mathfrak{S}_r$. Throughout this section, we denote $\langle A, B\rangle \coloneqq \sum_{ij} A_{ij} B_{ij}$ the Frobenius inner product of matrices A, B.

Normalization Function. We choose as a canonical representative of an orbit modulo $\mathfrak{S}_n \times \mathfrak{S}_r$ the coherent matching field Λ such that the matrix $v(\Lambda)$ is maximal in the sense of the lexicographic ordering $\leq_{\mathrm{lex}}$, with the convention that the entries are read column by column. The computation of such a representative is provided by the following statement:

Proposition 2. *Let $\Lambda \in \mathrm{CMF}(n, r)$. The algorithm* NormalizeCMF(Λ) *computes an action $(\sigma, \tau) \in \mathfrak{S}_n \times \mathfrak{S}_r$ which maximizes $v\big((\sigma, \tau) \cdot \Lambda\big)$ w.r.t. $\leq_{\mathrm{lex}}$ in time $O(r^2 n \log n)$.*

The algorithm NormalizeCMF(Λ) can be described as follows. Let $A \coloneqq v(\Lambda)$. After k iterations, the matrix $A^{(k)} \coloneqq \big(A_{\sigma(i)\tau(l)}\big)_{i\in[n],\, l\in[k]}$ is the lexicographically maximal k-column prefix obtainable from A by permuting rows and choosing k columns. The set $J \subset [r]$ contains the selected columns, and the bijection $\tau \colon [k] \to J$ records their order. At iteration k, we compute a partition B of equal rows of $A^{(k)}$ (Section 3). For each unused column index $j \notin J$, we sort the column vector $(A_{\sigma(i)j})_{i\in[n]}$ in every block of the partition B, which provides the sorted vector c_j and a corresponding permutation $\sigma_j \in \mathfrak{S}_n$ of the rows (Section 3). Since all candidate prefixes with $(k + 1)$ columns share the same first k columns $A^{(k)}$, choosing any $j^\star$ such that $c_{j^\star}$ is lexicographically maximal yields $A^{(k+1)}$.

Following this, we set $\phi(\Lambda) \coloneqq (\sigma, \tau) \cdot \Lambda$ where $(\sigma, \tau) = $ NormalizeCMF(Λ).

Adjacency Oracle. The construction of the adjacency oracle closely follows the one of [12]. By the Minkowski sum structure, every vertex of $P_{n,r}$ is a sum of the vertices of its summands B_I. Moreover, every edge of $P_{n,r}$ is parallel to an edge of one of the summands. Recall that each B_I is an embedded copy of the Birkhoff polytope B_r, whose vertices are the permutation matrices of size $r \times r$, and two vertices are adjacent if and only if they differ from a nontrivial cycle. As a consequence, in order to capture all edges to $P_{n,r}$, it suffices to define the index set as $\Delta \coloneqq \binom{[n]}{r} \times \mathfrak{C}_r$, where $\mathfrak{C}_r$ is the set of nontrivial cycles in $\mathfrak{S}_r$. For each $\delta = (I, \sigma) \in \Delta$ and $\Lambda \in \mathrm{CMF}(n, r)$, we set $e(\Lambda, \delta) \coloneqq E(\sigma \circ \Lambda_I) - E(\Lambda_I) \in \mathbb{R}^{n\times r}$, which is the direction of the edge from Λ_I to $\sigma \circ \Lambda_I$ in the summand B_I. The tangent cone at $v(\Lambda)$ is then given by the pointed

procedure NormalizeCMF(Λ)
  $A \coloneqq v(\Lambda)$, $\sigma \coloneqq$ identity map over $[n]$
  $k \coloneqq 0$, $J \coloneqq \varnothing$, $\tau \coloneqq$ empty map
  while $k < r$ **do**
    $B \coloneqq \textsc{RowPartition}((A_{\sigma(i)\tau(\ell)})_{(i,\ell)\in[n]\times[k]})$
    for all $j \notin J$ **do**
      $(c_j, \sigma_j) \coloneqq \textsc{Sort}((A_{\sigma(i)j})_{i\in[n]}, B)$
    let $j^\star$ be any index such that $c_{j^\star}$ is maximal in the lexicographic sense
    $\sigma \coloneqq \sigma_{j^\star}$, $J \coloneqq J \cup \{j^\star\}$, $\tau \coloneqq \tau[k+1 \mapsto j^\star]$, $k \coloneqq k+1$
  return (σ, τ)

Fig. 2. Normalization of coherent matching fields

convex cone generated by the elements $e(\Lambda,\delta)$ for $\delta \in \Delta$. Moreover, the index δ determines an edge of $P_{n,r}$ if and only if $e(\Lambda,\delta)$ is an extreme ray of this cone. The latter condition can be verified by determining if there exists a separating element $Z \in \mathbb{R}^{n\times r}$, i.e.,

$$\langle Z, e(\Lambda,\delta)\rangle \geq 1 \text{ and } \langle Z, e(\Lambda,\delta')\rangle \leq 0 \text{ for all } \delta' \text{ s.t. } e(\Lambda,\delta') \neq e(\Lambda,\delta)\,. \tag{$\mathsf{F}_{(\Lambda,\delta)}$}$$

We point out that the condition $e(\Lambda,\delta') \neq e(\Lambda,\delta)$ amounts to the fact that the rays generated by $e(\Lambda,\delta')$ and $e(\Lambda,\delta)$ are distinct, since $e(\Lambda,\delta')$ and $e(\Lambda,\delta)$ are both matrices with entries in $\{0,\pm1\}$. As in any Minkowski sum, an edge of $P_{n,r}$ in a given direction is the sum of all summand edges with that direction, so moving to the adjacent vertex requires pivoting simultaneously along all of them. For $\delta \in \Delta$, let $\mathcal{R}(\Lambda,\delta) \coloneqq \{\delta' \in \Delta \colon e(\Lambda,\delta') = e(\Lambda,\delta)\}$, and let $\operatorname{Pivot}(\Lambda, R)$ be the matching field obtained by replacing Λ_I with $\sigma \circ \Lambda_I$ for all $(I,\sigma) \in R$, leaving the other minors unchanged. We define $\operatorname{Adj}(v(\Lambda),\delta) \coloneqq v(\operatorname{Pivot}(\Lambda,\mathcal{R}(\Lambda,\delta)))$ if $(\mathsf{F}_{(\Lambda,\delta)})$ is feasible, and $\operatorname{Adj}(v(\Lambda),\delta) \coloneqq \perp$ otherwise.

Local Search Function. We fix the root vertex to the matrix $v^\star \coloneqq v(\Lambda^\star)$ where $\Lambda^\star$ is the all-diagonal coherent matching field. The ranking function is then chosen as $\rho(v) \coloneqq \langle C^\star, v\rangle$ where we define

$$C^\star \coloneqq \begin{pmatrix} M^{nr-1} & M^{n(r-1)-1} & \cdots & M^{n-1} \\ M^{nr-2} & M^{n(r-1)-2} & \cdots & M^{n-2} \\ \vdots & \vdots & & \vdots \\ M^{n(r-1)} & M^{n(r-2)} & \cdots & 1 \end{pmatrix}$$

and M is a real number manipulated as a symbolic value satisfying $M \gg 1$. The vertex $v^\star$ is the unique maximizer of ρ over $P_{n,r}$, since the diagonal permutation term strictly dominates all other permutations in every $r\times r$ minor of $C^\star$. Moreover, the normalization function ϕ defined above satisfies Assumption 1, as $\langle C^\star, A\rangle \leq \langle C^\star, B\rangle$ is equivalent to $A \leq_{\text{lex}} B$.

As in [12], the local search function consists in applying one step of the shadow-vertex simplex pivot rule induced by the objective vector $C^\star$ and a vector of the interior of the normal cone of the current coherent matching field. For brevity, write $N(\Lambda)$ for the normal cone of $v(\Lambda)$ in $P_{n,r}$, and $\operatorname{int} N(\Lambda)$ for its interior (recall $N(\Lambda)$ is full-dimensional). We elaborate on a method to find such a vector that is more direct than the general approach described in [12] for the case of coherent matching fields. This relies on the following statement:

Proposition 3. *Let $\Lambda \in \operatorname{CMF}(n,r)$. The interior of the normal cone of $v(\Lambda)$ consists of the rays generated by matrices $W \in \mathbb{R}^{n\times r}$ satisfying*

$$z^I_{\Lambda_I(i)} + W_{i\Lambda_I(i)} \geq z^I_j + W_{ij} + 1\,, \qquad \forall I \in \binom{[n]}{r}\,,\ \forall i \in I\,,\ \forall j \in [r]\setminus\{\Lambda_I(i)\} \tag{H_Λ}$$

for some vectors $z^I \in \mathbb{R}^r$ ($I \in \binom{[n]}{r}$).

The characterization is based on the dual formulation of the perfect matching problem applied to every minor. More precisely, given a fixed minor $I \in \binom{[n]}{r}$, the existence of a vector $z^I \in \mathbb{R}^r$ satisfying the condition (H_Λ) is equivalent to the fact that the submatrix $(W_{ij})_{i\in I,\, j\in[r]}$ is in the interior of the normal cone of the vertex $E(\Lambda_I)$ of the Birkhoff polytope. As a consequence of Proposition 3, we can use a LP solver to compute an element $W(\Lambda) \in \operatorname{int} N(\Lambda)$.

We define $f(v(\Lambda))$ for each $\Lambda \neq \Lambda^\star$ by one step of the shadow-vertex rule. More precisely, we consider the (affine) ray $\gamma(\mu) := \mu W(\Lambda) + C^\star$ for μ from $+\infty$ to 0. Then $\gamma(\mu) \in \operatorname{int} N(\Lambda)$ for all sufficiently large μ, whereas $\gamma(0) \notin N(\Lambda)$. Let $\mu^\star := \inf\{\mu > 0 \colon \gamma(\mu) \in N(\Lambda)\}$. Since $N(\Lambda) = \{C \in \mathbb{R}^{n\times r} \colon \langle e(\Lambda,\delta), C\rangle \leq 0 \text{ for all } \delta \in \Delta\}$, $\mu^\star$ coincides with the largest value μ_δ for which the corresponding constraint becomes active, namely, $\langle e(\Lambda,\delta), \gamma(\mu_\delta)\rangle = 0$. These values are given by $\mu_\delta = -\frac{\beta_\delta}{\alpha_\delta}$ where $\alpha_\delta := \langle e(\Lambda,\delta), W(\Lambda)\rangle < 0$ and $\beta_\delta := \langle e(\Lambda,\delta), C^\star\rangle$. Note that α_δ is independent of M, whereas β_δ is a polynomial in M. Hence the values μ_δ can be compared symbolically by their highest powers of M, which is equivalent to comparing the matrices $-\frac{1}{\alpha_\delta} e(\Lambda,\delta)$ according to $\leq_{\text{lex}}$. In this way, we can easily compute the set $\mathcal{R}^\star(\Lambda) := \{\delta \in \Delta \colon \mu_\delta = \mu^\star\}$. We then define the local search map by $f(v(\Lambda)) := v(\operatorname{Pivot}(\Lambda, \mathcal{R}^\star(\Lambda)))$. In the next statement, we show that the local search function returns a vertex of $P_{n,r}$, and that the latter satisfies the requirements of the reverse search method described in Section 2.

Proposition 4. *For all $\Lambda \in \operatorname{CMF}(n,r)$ such that $\Lambda \neq \Lambda^\star$, $f(v(\Lambda))$ is a vertex of $P_{n,r}$ adjacent to $v(\Lambda)$, and $\langle C^\star, f(v(\Lambda))\rangle > \langle C^\star, v(\Lambda)\rangle$. Moreover, $f(v(\Lambda))$ can be computed in time $O(nr\#\Delta)$ from a normal vector $W(\Lambda)$.*

Parallel Implementation in Rust. We have implemented the enumeration of the orbits of coherent matching fields in Rust. Linear feasibility problems are solved using `HiGHS` [14]. Our implementation is a distributed parallel implementation based on MPI, and relies on a boss/worker task distribution inspired by [22]. Suppose that the boss process explores the subtree rooted at a vertex v. Every time it finds a new child w, it assigns the task of enumerating the subtree rooted at w to some idle worker process. If there is none, the boss process continues the exploration itself. In case the boss finishes its job and gets idle before some workers, it unassigns the oldest busy worker which returns to the boss the current state of its job. The boss process resumes the job and can continue distributing subtree exploration jobs to worker processes. As a result, processes do not share any memory and only exchange messages through the MPI library: (i) for worker processes, idleness notifications or returning the current state of a job; (ii) for the boss process, job assignments or unassignment notification. Thanks to the memoryless nature of the reverse search method, messages carrying job assignments or the current state of a stolen job only consist of a pair formed by an integer d representing a depth and an integer vector encoding the current coherent matching field Λ. This information fully characterizes the current state of the exploration of a subtree: if $d > 0$, once the subtree rooted at Λ has fully been explored, the depth d is decremented, and the exploration of the parent subtree continues with the exploration of the children located after Λ in the ordering induced by the adjacency oracle.

The experiments were conducted on a cluster of nine machines, each with 56 Intel Core processors (Haswell architecture, 2.3 GHz). The 9×3, 10×3 and 8×4 coherent matching fields were enumerated in about 3 hours, 3 days 16 hours and 7 days 7 hours, respectively (the other entries were all obtained in less than 10 minutes). As a consequence of the memoryless nature of the reverse search method, memory usage remains negligible. While we did not have the opportunity to evaluate the parallel scaling, we observed that, as n and r grow, the job distribution becomes less efficient and many workers may be left waiting to be assigned a job. This is because discovering new children becomes computationally more demanding, making the work of the boss the bottleneck. We plan to investigate this aspect further and evaluate whether alternative job-splitting techniques improve efficiency.

In a future work, we plan to implement the enumeration of pointed coherent matching fields. We expect this to be done with minor modifications: setting the root of the search tree to the lexicographically maximal pointed coherent matching field; keeping only the edges of the adjacency oracle between pointed coherent matching fields; restricting the normalization to the permutations that preserve pointedness. By the correspondence recalled in Section 1, the pointed coherent matching fields of size $n \times r$ are exactly the regular triangulations of $\triangle_{n-r-1} \times \triangle_{r-1}$.

This will allow us to compare our approach with those of [15,19], in which all triangulations of a point configuration are enumerated and the non-regular ones are filtered out afterwards.

4 Tropical Realizability of Order Types

Let $p_1, \dots, p_n \in \mathbb{R}^d$ be a configuration of n labelled points. The *order type* is a signed vector $\chi \in \{0, \pm 1\}^{\binom{n}{d+1}}$ which, for each $(d+1)$-subset $I = \{i_1 < \dots < i_{d+1}\}$, provides the orientation $\chi_I \coloneqq \operatorname{sgn}\det\begin{pmatrix} 1 & \cdots & 1 \\ p_{i_1} & \cdots & p_{i_{d+1}} \end{pmatrix}$ of the point set $(p_{i_1}, \dots, p_{i_{d+1}})$. In this section, we study configurations of points in *general position*, i.e., no $d+1$ points lie on a common affine hyperplane. Equivalently, $\chi_I \in \{\pm 1\}$ for all $I \in \binom{[n]}{d+1}$. We refer to [13] for an introduction to order types.

The tropical analogue of order types is defined by means of points with tropical signed numbers as entries. A *tropical signed number* is a pair $(\varepsilon, a) \in (\{\pm 1\} \times \mathbb{R}) \cup \{(0, -\infty)\}$ consisting of a sign ε and a modulus $a \in \mathbb{R} \cup \{-\infty\}$. We denote by $\mathbb{T}_\pm$ the set of tropical signed numbers, and by $\operatorname{tsgn}(\cdot)$ and $|\cdot|$ the sign and modulus functions. Signed tropical numbers arise as the image under the signed valuation $z \mapsto (\operatorname{sgn} z, \operatorname{val} z)$ of elements of an ordered valued field, e.g. real Puiseux series with real exponents [16]; see [6] for a discussion on other such fields. The sum $(\varepsilon_1, a_1) \oplus \cdots \oplus (\varepsilon_p, a_p)$ of finitely many tropical signed numbers (ε_i, a_i) is defined as any of the term $(\varepsilon_{i^\star}, a_{i^\star})$ of maximal modulus $a_{i^\star}$ whenever all terms attaining this maximum have the same sign, leaving the sum undefined otherwise.[1] The product of two tropical signed numbers $(\varepsilon, a), (\varepsilon', a')$ is defined as $(\varepsilon, a) \odot (\varepsilon', a') \coloneqq (\varepsilon\varepsilon', a + a')$. The unit element is $(+1, 0)$. The *tropical determinant* of a tropical signed matrix $A \in \mathbb{T}_\pm^{k \times k}$ is defined as $\operatorname{tdet} A \coloneqq \bigoplus_{\sigma \in \mathfrak{S}_k} (\operatorname{sgn}\sigma, 0) \odot \bigodot_{i=1}^k A_{i\sigma(i)}$. The matrix A is *generic* if $\operatorname{tdet} A$ is a well-defined tropical number distinct from $(0, -\infty)$; equivalently, the signs of all the terms $(\operatorname{sgn}\sigma, 0) \odot \bigodot_{i=1}^k A_{i\sigma(i)}$ with maximal modulus are equal and nonzero.

A configuration of tropical points $p_1, \dots, p_n \in \mathbb{T}_\pm^d$ is in *general position* if, for each $(d+1)$-subset $I = \{i_1 < \dots < i_{d+1}\}$, the tropical signed matrix $P_I \coloneqq \begin{pmatrix} (+1,0) & \cdots & (+1,0) \\ p_{i_1} & \cdots & p_{i_{d+1}} \end{pmatrix}$ is generic. As a consequence of [2, Theorem 6.5], this holds if and only if no $d+1$ points lie in the same tropical affine hyperplane; see [3] for more details. The order type $\chi \in \{\pm 1\}^{\binom{n}{(d+1)}}$ of such a point configuration is then defined by $\chi_I \coloneqq \operatorname{tsgn}(\operatorname{tdet} P_I)$. Every order type of a configuration of tropical points in general position can be realized by a configuration of points in general position over the reals.

In order to enumerate tropical order types, we use the fact that any configuration $p_1, \dots, p_n$ of tropical points in general position is associated to a face of $P_{n,d+1}$ induced by the weights $(|P_{ji}|)_{ij}$, where $P \coloneqq \begin{pmatrix} (+1,0) & \cdots & (+1,0) \\ p_1 & \cdots & p_n \end{pmatrix}$. More precisely, the tropical order type can be reconstructed from any coherent matching field Λ such that $v(\Lambda)$ is a vertex of the face, and the signs of the P_{ji} for (i,j) in the support of $v(\Lambda)$. Hence, it suffices to enumerate the tropical order types arising from every pair of $n \times (d+1)$ coherent matching field Λ and sign pattern $\varepsilon_{ij} \in \{\pm 1\}$ for (i,j) in the support of $v(\Lambda)$, with the constraint that $\varepsilon_{i1} = +1$. By [21, Corollary 3.4], there are $2^{d \times (n-d)}$ such sign patterns for every Λ.

In this way, we enumerate the order types up to the action of $\mathfrak{S}_n$ in the case $d = 2$, and compare the resulting counts with those given in [1, Table 1]. Orbits modulo $\mathfrak{S}_n$ are computed using the so-called λ-matrices; see [13]. For $n \leq 8$, we recover the same numbers of order types, while for $n = 9$ we obtain 158,737 tropical order types up to relabeling. This proves Theorem 1.

Acknowledgement. The two authors warmly thank Stéphane Gaubert and Xavier Goaoc for numerous discussions on coherent matching fields and order types. They are also grateful to Michael Joswig for discussions on enumeration algorithms handling symmetries, and to Guillaume Dindart for his support in using the cluster. Both authors were partially supported by ANR projects HilbertXfield (ANR-23-CE40-0012) and ZADyG (ANR-25-CE48-7058).

[1] We refer to [2] for the construction of a semiring of tropical signed numbers where the tropical sum is always defined. The nonzero tropical signed numbers defined here correspond to the “unbalanced” elements of [2], and the addition and multiplication remain consistent with those of the semiring.

References

1. Aichholzer, O., Aurenhammer, F., Krasser, H.: Enumerating order types for small point sets with applications. Order **19**(3) (2002)
2. Akian, M., Cohen, G., Gaubert, S., Nikoukhah, R., Quadrat, J.: Linear systems in (max, +) algebra. In: 29th IEEE Conference on Decision and Control (1990)
3. Akian, M., Gaubert, S., Guterman, A.: Tropical cramer determinants revisited. In: Tropical and Idempotent Mathematics and Applications, Contemporary Mathematics, vol. 616. American Mathematical Society (2014)
4. Akian, M., Gaubert, S., Guterman, A.: Tropical polyhedra are equivalent to mean payoff games. International Journal of Algebra and Computation **22**(01) (2012)
5. Allamigeon, X., Benchimol, P., Gaubert, S., Joswig, M.: Combinatorial simplex algorithms can solve mean payoff games. SIAM J. on Optimization **24**(4) (2014)
6. Allamigeon, X., Benchimol, P., Gaubert, S., Joswig, M.: What tropical geometry tells us about the complexity of linear programming. SIAM Review **63**(1) (2021)
7. Avis, D., Fukuda, K.: A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. Discrete & Computational Geometry **8**(3) (1992)
8. Bremner, D., Sikirić, M.D., Schürmann, A.: Polyhedral representation conversion up to symmetries. In: Polyhedral Computation. CRM Proceedings & Lecture Notes, vol. 48. American Mathematical Society (2009)
9. Clarke, O., Mohammadi, F.: Toric degenerations of Grassmannians and Schubert varieties from matching field tableaux. Journal of Algebra **559** (2020)
10. Develin, M., Sturmfels, B.: Tropical convexity. Documenta Mathematica **9** (2004)
11. Fink, A., Rincón, F.: Stiefel tropical linear spaces. Journal of Combinatorial Theory, Series A **135** (2015)
12. Fukuda, K.: From the zonotope construction to the Minkowski addition of convex polytopes. Journal of Symbolic Computation **38**(4) (2004)
13. Goodman, J.E., Pollack, R.: Multidimensional sorting. SIAM Journal on Computing **12**(3) (1983)
14. Huangfu, Q., Hall, J.A.J.: Parallelizing the dual revised simplex method. Mathematical Programming Computation **10**(1) (2018)
15. Jordan, C., Joswig, M., Kastner, L.: Parallel enumeration of triangulations. The Electronic Journal of Combinatorics **25**(3) (Jul 2018)
16. Markwig, T.: A field of generalised Puiseux series for tropical geometry. Rendiconti del Seminario Matematico della Università di Torino **68**(1) (2010)
17. Mohammadi, F., Shaw, K.: Toric degenerations of Grassmannians from matching fields. Algebraic Combinatorics **2**(6) (2019)
18. Rambau, J.: Topcom: Triangulations of point configurations and oriented matroids. In: Mathematical Software—ICMS 2002. World Scientific (2002)
19. Rambau, J.: Symmetric lexicographic subset reverse search for the enumeration of circuits, cocircuits, and triangulations up to symmetry (Jan 2026)
20. Santos, F.: The Cayley trick and triangulations of products of simplices. In: Integer Points in Polyhedra—Geometry, Number Theory, Algebra, Optimization, Contemporary Mathematics, vol. 374. American Mathematical Society (2005)
21. Sturmfels, B., Zelevinsky, A.: Maximal minors and their leading terms. Advances in Mathematics **98**(1) (1993)
22. Weibel, C.: Implementation and parallelization of a reverse-search algorithm for Minkowski sums. In: Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX'10) (2010)