

The Graphicahedron

Gabriela Araujo-Pardo*, Maria Del Río-Francos†,
Mariana López-Dudet‡, Deborah Oliveros§

Instituto de Matemáticas
Universidad Nacional Autónoma de México, México

Egon Schulte¶
Department of Mathematics
Northeastern University, Boston, USA

October 13, 2009

Abstract

The paper describes a construction of abstract polytopes from Cayley graphs of symmetric groups. Given any connected graph G with p vertices and q edges, we associate with G a Cayley graph $\mathcal{G}(G)$ of the symmetric group S_p and then construct a vertex-transitive simple polytope of rank q , the *graphicahedron*, whose 1-skeleton (edge graph) is $\mathcal{G}(G)$. The graphicahedron of a graph G is a generalization of the well-known permutahedron; the latter is obtained when the graph is a path. We also discuss symmetry properties of the graphicahedron and determine its structure when G is small.

Key words. Permutahedron, abstract polytopes, Cayley Graphs.

MSC 2000. Primary: 51M20. Secondary: 05C25, 52B15.

*Supported by CONACYT 5737 garaujo@matem.unam.mx

†mad210fcos@gmail.com

‡mdudet17@gmail.com

§Supported by CONACYT CCDG 50151 dolivero@matem.unam.mx

¶schulte@neu.edu

||Supported by NSA-grant H98230-07-1-0005

1 Introduction

In the present paper we describe an interesting abstract polytope, called the *graphicalahedron*, which generalizes the well-known permutahedron. The permutahedron Π_n is a simple convex n -polytope in \mathbb{R}^{n+1} that was apparently first investigated by Schoute in 1911 (see [7, 18, 21]); it was discovered in Guilbaud & Rosenstiehl [8] in 1963 and given the name “permutohedron” (or rather “permutoèdre”, in French). It is the convex hull of all points in \mathbb{R}^{n+1} obtained from $(1, 2, \dots, n+1)$ by permuting the coordinates in all possible ways. In particular, its vertices can be identified with the permutations in the symmetric group S_{n+1} in such a way that two vertices of Π_n are connected by an edge if and only if the corresponding permutations differ by an adjacent transposition.

Our construction of the graphicalahedron builds on Cayley graphs of symmetric groups. Given any connected graph G with p vertices and q edges, we associate with G a Cayley graph $\mathcal{G}(G)$ of the symmetric group S_p and then construct a vertex-transitive simple polytope of rank q , the graphicalahedron \mathcal{P}_G , whose 1-skeleton (edge graph) is $\mathcal{G}(G)$. The generating set of S_p defining $\mathcal{G}(G)$ consists of the transpositions associated with the edges of G . If G is a path of length n , then \mathcal{P}_G is isomorphic to Π_n .

The paper is organized as follows. In Sections 2 and 3, we review basic notation and in particular describe the Cayley graph $\mathcal{G}(G)$ associated with a given graph G . Then, in Section 4, we define the graphicalahedron \mathcal{P}_G , establish that it is an abstract polytope, and prove that its 1-skeleton is the Cayley graph $\mathcal{G}(G)$. In Section 5, we find the structure of the automorphism group of \mathcal{P}_G and then enumerate the graphicalahedra that are regular polytopes. Finally, in Section 6, we describe in more detail the geometric structure of the graphicalahedra when G is a path, a cycle, a star graph, or a small graph with up to four edges. We recover the classical permutahedra as the graphicalahedra associated with paths, and find other interesting highly-symmetric polytopes, including various locally toroidal polytopes.

2 Graphs and polytopes

We briefly review some important concepts for graphs and polytopes, beginning with graphs. For further basic definitions and terminology on graphs and polytopes the reader is referred to Chartrand & Lesniak [3] and McMullen & Schulte [12], respectively.

Throughout, $G = (V(G), E(G))$ denotes a simple graph, without loops or multiple edges and with vertex-set $V(G)$ and edge-set $E(G)$. We say that G is a (p, q) -graph if $|V(G)| = p$ and $|E(G)| = q$. We always assume that p and q are finite. A subgraph of G is called a *spanning* subgraph of G if it contains all the vertices of G . The spanning subgraphs of G are in one-to-one correspondence with the subsets of $E(G)$. In fact, any subset of $E(G)$ uniquely defines a spanning subgraph of G , of which it is the full edge set, and vice versa. For any two spanning graphs H and H' of G we write $H \subseteq H'$ if $E(H) \subseteq E(H')$.

An *abstract polytope* of rank n , or simply an n -polytope, is a partially ordered set (or *poset* for short) \mathcal{P} with a strictly monotone rank function having range $\{-1, 0, \dots, n\}$. The elements of \mathcal{P} are called *faces*, or j -*faces* if their rank is j . The faces of ranks 0, 1 or $n - 1$ are also called the *vertices*, *edges* or *facets* of \mathcal{P} , respectively. Moreover, \mathcal{P} has a smallest face (of rank -1) and largest face (of rank n), denoted by F_{-1} and F_n , respectively; they are the *improper faces* of \mathcal{P} . Each *flag* (maximal chain) of \mathcal{P} contains exactly $n + 2$ faces. Two flags are said to be *adjacent* if they differ in exactly one face; they are j -*adjacent* if this face has rank j . In \mathcal{P} , any two flags Φ and Ψ can be joined by a sequence of flags $\Phi = \Phi_0, \Phi_1, \dots, \Phi_k = \Psi$, all containing $\Phi \cap \Psi$, such that any two successive flags Φ_{i-1} and Φ_i are adjacent; this property is known as the *strong flag-connectedness* of \mathcal{P} . Finally, \mathcal{P} has the following *homogeneity property*, often called the *diamond condition*: whenever $F \leq G$, with $\text{rank}(F) = j - 1$ and $\text{rank}(G) = j + 1$, there are exactly two faces H of rank j such that $F \leq H \leq G$.

Let \mathcal{P} be an n -polytope, and let $0 \leq k \leq n - 1$. The k -*skeleton* of \mathcal{P} is the poset consisting of all proper faces of \mathcal{P} of rank at most k (together with the induced partial order).

3 Cayley Graphs

Given a finite group Γ and any subset \mathcal{T} of Γ consisting of involutions, the *Cayley graph of Γ with respect to \mathcal{T}* , denoted by $\mathcal{G}(\Gamma, \mathcal{T})$, is the graph with vertex-set Γ such that two vertices γ_1 and γ_2 are adjacent (connected by an edge) if and only if $\gamma_2 = \tau\gamma_1$ for some $\tau \in \mathcal{T}$. Here we slightly abuse standard notation and do not require \mathcal{T} to be a generating set of Γ , although this will usually be the case. Note that $\mathcal{G}(\Gamma, \mathcal{T})$ is connected (that is, $\mathcal{G}(\Gamma, \mathcal{T})$ is a Cayley graph in standard terminology) if and only if \mathcal{T} is a generating set of Γ . We are primarily interested in Cayley graphs of symmetric groups S_p .

Let G be a (p, q) -graph with vertex set $V(G) := \{1, \dots, p\}$ and edge set

$E(G) = \{e_1, \dots, e_q\}$, where $p \geq 1$ and $q \geq 0$ (if $q = 0$, then $E(G) = \emptyset$). In most applications, G will be connected. We associate with G a Cayley graph on S_p as follows. If $e = \{i, j\}$ is an edge of G (with vertices i and j), define $\tau_e := (i j)$; this is the transposition in S_p that interchanges i and j . Let $\mathcal{T}_G := \{\tau_{e_1}, \dots, \tau_{e_q}\}$ denote the set of transpositions determined by the edges of G , and let $T_G := \langle \tau_{e_1}, \dots, \tau_{e_q} \rangle$ denote the subgroup of S_p generated by \mathcal{T}_G . If G is connected, then $T_G = S_p$ (see Lemma 3.1 below). More generally, if $K \subseteq E(G)$, then we define $\mathcal{T}_K := \{\tau_e \mid e \in K\}$ and $T_K := \langle \tau_e \mid e \in K \rangle$. If K happens to be the full edge set of a subgraph H of G , then we also write \mathcal{T}_H or T_H in place of \mathcal{T}_K and T_K ; that is, $\mathcal{T}_H := \mathcal{T}_{E(H)}$ and $T_H := T_{E(H)}$. If $K = \emptyset$, then T_K is the trivial group.

Now the *Cayley graph of G* , denoted by $\mathcal{G}(G)$, is the Cayley graph of S_p with respect to \mathcal{T}_G ; that is, $\mathcal{G}(G) := \mathcal{G}(S_p, \mathcal{T}_G)$. Thus $V(\mathcal{G}(G)) = S_p$, and $\{\gamma_1, \gamma_2\} \in E(\mathcal{G}(G))$ if and only if $\tau_e \gamma_1 = \gamma_2$ for some $e \in E(G)$. For instance, if G has no edges, then $\mathcal{G}(G)$ also has no edges; and if G has only one edge, then $\mathcal{G}(G)$ is a matching on S_p . On the other hand, connected graphs are characterized by the following lemma.

Lemma 3.1 *The graph G is connected if and only if \mathcal{T}_G generates S_p . Equivalently, G is connected if and only if $\mathcal{G}(G)$ is connected.*

Proof. Suppose G is connected. We only need to verify that every transposition of S_p is generated by \mathcal{T}_G . Let $(i j)$ be any transposition in S_p . Since G is connected, there exist a path along edges of G connecting the vertices i and j of G . (If $\{i, j\} \in E(G)$, then the path has just one edge.) Suppose the path transverses the edges f_1, \dots, f_r of G , in this order, so that i is the first vertex of f_1 and j is the last vertex of f_r . Then $(i j) = \tau_{f_r} \cdots \tau_{f_2} \tau_{f_1} \tau_{f_2} \cdots \tau_{f_r}$.

Conversely, suppose G is not connected. If H is a connected component of G and \mathcal{T}_H is the corresponding set of transpositions, then the subgroup T_H generated by \mathcal{T}_H is the symmetric group $S_{V(H)}$ on the vertex-set $V(H)$ of H and is a proper subgroup of $S_p = S_{V(G)}$. It follows that the subgroup T_G generated by \mathcal{T}_G is isomorphic to the direct product of symmetric groups $S_{V(H)}$, where H runs over all connected components of G . In any case, T_G is a proper subgroup of S_p . Thus \mathcal{T}_G fails to generate S_p . ■

Clearly, if G and G' are isomorphic (p, q) -graphs then the corresponding Cayley graphs $\mathcal{G}(G)$ and $\mathcal{G}(G')$ are also isomorphic. The following lemma summarizes certain basic properties of the Cayley graphs $\mathcal{G}(H)$ associated with spanning subgraphs H of G ; they all are subgraphs of $\mathcal{G}(G)$.

Lemma 3.2 (a) A subgraph H of G is a spanning subgraph of G if and only if $\mathcal{G}(H)$ is a spanning subgraph of $\mathcal{G}(G)$.
 (b) If H and H' are spanning subgraphs of G and $H \subseteq H'$, then also $\mathcal{G}(H) \subseteq \mathcal{G}(H')$.
 (c) For any k with $0 \leq k \leq q$ we have $\mathcal{G}(G) = \cup \mathcal{G}(G_k)$, where G_k runs over all spanning subgraphs of G with k edges.

Figures 1 and 2 show the Cayley graphs of the path $G = P_2$ of length 2 and the cycle $G = C_3$ of length 3, respectively; now $p = 3$ and $q = 2$ or 3 , respectively. Here the basic transpositions in \mathcal{T}_G can be represented by different colors; in fact, color representations of Cayley graphs are quite common and some authors use the term *Cayley color graph* to emphasize this idea (see [3]).




Figure 1: The basic graph P_2 and its Cayley color graph $\mathcal{G}(P_2) = \mathcal{G}(S_3, \{\tau_{\{1,2\}}, \tau_{\{2,3\}}\})$

The following lemma relates the Cayley graph of the path P_n on $n + 1$ vertices to the permutahedron Π_n in \mathbb{R}^{n+1} . For $n = 2$, the permutahedron is a hexagon (see Figure 1).

Lemma 3.3 Let P_n denote the simple path with n edges and $n + 1$ vertices. Then the Cayley graph $\mathcal{G}(P_n)$ of P_n and the 1-skeleton of the permutahedron Π_n are isomorphic.




Figure 2: The basic graph C_3 and its Cayley color graph $\mathcal{G}(C_3) = \mathcal{G}(S_3, \{\tau_{\{1,2\}}, \tau_{\{1,3\}}, \tau_{\{2,3\}}\})$

4 The Graphicahedron

Let G be a connected (p, q) -graph. In this section, we describe a partially ordered set \mathcal{P}_G , called the *graphicahedron*, and show in particular that \mathcal{P}_G is an abstract polytope of rank q . The i -faces of \mathcal{P}_G are given by the connected components of the Cayley graphs of the form $\mathcal{G}(H)$, where H is a spanning subgraph of G with i edges.

In defining the graphicahedron it is convenient to initially suppress the (smallest) face of rank -1 and concentrate on faces of ranks $0, \dots, q$. The missing face of rank -1 will be appended at the end.

For $i \in I := \{0, \dots, q\}$ define

$$C_i := \{(K, \alpha) \mid K \subseteq E(G), |K| = i, \alpha \in S_p\}. \quad (1)$$

Then we have the following equivalence relation on $\bigcup_{i \in I} C_i$.

Definition 4.1 *Two elements $(K, \alpha), (L, \beta) \in \bigcup_{i \in I} C_i$ are said to be equivalent, or $(K, \alpha) \sim (L, \beta)$ for short, if and only if $K = L$ and $T_K \alpha = T_L \beta$ (equality as right cosets in S_p).*

By a slight abuse of notation we use the same symbol for both, an element of $\bigcup_{i \in I} C_i$ and its equivalence class under \sim .

Definition 4.2 *Let $\mathcal{P}_G := (\bigcup_{i \in I} C_i) / \sim$. When $(K, \alpha), (L, \beta) \in \mathcal{P}_G$, define $(K, \alpha) \leq (L, \beta)$ if and only if $K \subseteq L$ and $T_K \alpha \subseteq T_L \beta$.*

It is straightforward to check that \leq is well-defined and indeed makes \mathcal{P}_G a partially ordered set. As before, its *faces* are its elements, (K, α) . In particular, \mathcal{P}_G has a strictly monotone rank function given by

$$\text{rank}(K, \alpha) := |K|. \quad (2)$$

Hence (K, α) is an i -face of \mathcal{P}_G if and only if $|K| = i$. In particular, each vertex (0-face) of \mathcal{P}_G is of the form (\emptyset, α) with $\alpha \in S_p$ (and \emptyset as first component). Thus \mathcal{P}_G has $p!$ vertices. Moreover, there is just one q -face, $(E(G), \alpha)$ for any α , which is incident with any other face of \mathcal{P}_G . (Recall here that we are still excluding $i = -1$.)

It is helpful to relate the faces of \mathcal{P}_G to the connected components of the Cayley graphs of spanning subgraphs of G . First recall that the subsets of $E(G)$ are in one-to-one correspondence with the spanning subgraphs of G . Thus we may think of a subset K of $E(G)$ as a spanning subgraph, \tilde{K} (say), of G . Then the right coset $T_K\alpha$ of the subgroup $T_K = T_{\tilde{K}}$ in S_p involved in Definition 4.1 is just the connected component of α in the Cayley graph $\mathcal{G}(\tilde{K})$, and equivalence of faces in \mathcal{P}_G is just equality of connected components. In summary, we have

Lemma 4.1 *The faces of \mathcal{P}_G are in one-to-one correspondence with the connected components of the Cayley graphs of spanning subgraphs of G . Under this correspondence, if K denotes the spanning subgraph of G with edge set K , then the face (K, α) corresponds to the connected component of α in $\mathcal{G}(\tilde{K})$.*

It is important to point out that the partial order in \mathcal{P}_G is not in general equivalent to inclusion of connected components of Cayley graphs, or, equivalently, inclusion of right cosets in S_p . Clearly, incidence in \mathcal{P}_G does imply inclusion of cosets; in fact, by definition, if $(K, \alpha) \leq (L, \beta)$ in \mathcal{P}_G , then $T_K\alpha \subseteq T_L\beta$ in S_p . However, the converse is not true in general, since subgroups generated by transpositions often have quite different generating sets of transpositions; there may exist subsets K and L , not related by inclusion, but nevertheless satisfying $T_K\alpha \subseteq T_L\beta$ for some α, β . On the other hand, the following Remark describes an interesting special case when the converse does hold.

Remark 4.1 *If G is a tree, then the partial order in \mathcal{P}_G is equivalent to inclusion of right cosets in S_p . In other words, in this case we have $(K, \alpha) \leq (L, \beta)$ in \mathcal{P}_G if and only if $T_K\alpha \subseteq T_L\beta$ in S_p .*

Here the additional condition that $K \subseteq L$ is already implied by the condition that $T_K\alpha \subseteq T_L\beta$; that is, the former is not required in the definition of the partial

order of \mathcal{P}_G . This follows from the fact that, for a tree G , any subgroup of S_p of the form T_K is isomorphic to a direct product of symmetric groups, one for each connected component of the spanning subgraph of G with edge set K . In fact, if $T_K\alpha \subseteq T_L\beta$ in S_p , then $T_L\beta = T_L\alpha$ and hence also $T_K \subseteq T_L$; but then, in turn, the direct product structure of the subgroups T_K and T_L implies that $K \subseteq L$.

Resuming our general discussion, let again G be any connected (p, q) -graph. Next we investigate the flags of \mathcal{P}_G . First note that, if (K, α) and (L, β) are two faces of \mathcal{P}_G with $(K, \alpha) \leq (L, \beta)$, then $(L, \beta) = (L, \alpha)$ in \mathcal{P}_G , so in representing the latter face we may replace β by α . In fact, since $(K, \alpha) \leq (L, \beta)$, we have $\alpha \in T_K\alpha \subseteq T_L\beta$ and hence $T_L\alpha = T_L\beta$. It follows that any chain of mutually incident faces in \mathcal{P}_G can be represented in such a way that all second components are the same; in fact, we can always use the second component of the smallest face in the chain.

In particular, a flag Φ of \mathcal{P}_G can be described by only two parameters, namely a maximal nested family of subsets of $E(G)$, denoted by $\mathcal{K}_\Phi := \mathcal{K} := \{K_0, K_1, \dots, K_q\}$, and a single element $\alpha \in S_p$; that is,

$$\Phi = (\mathcal{K}, \alpha) := \{(K_0, \alpha), (K_1, \alpha), \dots, (K_q, \alpha)\}.$$

By a *maximal nested* family of subsets of a given finite set we mean a flag in the Boolean lattice (power set) associated with this set. Note here that K_i contains exactly i edges, so, in particular, $K_0 = \emptyset$ and $K_q = E(G)$. (Bear in mind that we are still ignoring the (-1) -face.) Thus, as will become clear, a flag is parametrized by its vertex (in a sense, α) and a flag of the vertex-figure at this vertex (in a sense, \mathcal{K}).

Our next lemma is aimed at establishing the strong flag-connectedness of \mathcal{P}_G .

Lemma 4.2 *Let $\mathcal{K} = \{K_0, K_1, \dots, K_q\}$ and $\mathcal{F} = \{F_0, F_1, \dots, F_q\}$ be two maximal nested families of subsets of $E(G)$. Then there exists a sequence $\mathcal{K} = \mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_{s-1}, \mathcal{K}_s = \mathcal{F}$ of maximal nested families of subsets of $E(G)$ such that*

- (a) \mathcal{K}_j differs from \mathcal{K}_{j+1} in exactly one element, for all $j = 0, \dots, s-1$;
- (b) $\mathcal{K} \cap \mathcal{F} \subseteq \mathcal{K}_j$, for all $j = 0, \dots, s$.

Proof. One possible way to establish the lemma is to appeal to the strong flag-connectedness of the face-lattice of the $(q-1)$ -simplex; the latter is isomorphic to the Boolean lattice on $E(G) = \{e_1, \dots, e_q\}$. Viewed in this setting, the two families \mathcal{K} and \mathcal{F} are flags of the $(q-1)$ -simplex and can be joined by a sequence of successively adjacent flags of the $(q-1)$ -simplex all containing $\mathcal{K} \cap \mathcal{F}$. ■

Next we settle the strong flag-connectedness of \mathcal{P}_G . We will see that, in essence, the previous lemma corresponds to the strong flag-connectedness of the vertex-figures of \mathcal{P}_G .

Lemma 4.3 *The poset \mathcal{P}_G is strongly flag-connected. In other words, if Φ and Ψ are two flags of \mathcal{P}_G , then there exists a sequence of successively adjacent flags $\Phi = \Phi_0, \Phi_1, \dots, \Phi_{s-1}, \Phi_s = \Psi$ of \mathcal{P}_G such that $\Phi \cap \Psi \subseteq \Phi_i$ for all $i = 0, \dots, s$.*

Proof. Let $\Phi = \{(K_0, \alpha), \dots, (K_q, \alpha)\}$ and $\Psi = \{(F_0, \beta), \dots, (F_q, \beta)\}$ be two flags of \mathcal{P}_G , where $\alpha, \beta \in S_p$ and $\mathcal{K} := \{K_0, \dots, K_{q-1}\}$ and $\mathcal{F} := \{F_0, \dots, F_{q-1}\}$ are maximal nested families of subsets of $E(G)$, respectively; that is, $\Phi = (\mathcal{K}, \alpha)$ and $\Psi = (\mathcal{F}, \beta)$. Let J denote the set of suffices j with $(K_j, \alpha) = (F_j, \beta)$; then J consists of the positions j where Φ and Ψ agree, and $J \neq \emptyset$ since $q \in J$. Let m denote the smallest suffix in J .

First we settle the case that Φ and Ψ share a 0-face (vertex), or equivalently, that $\alpha = \beta$, or that $m = 0$. If $\alpha = \beta$, we can appeal to Lemma 4.2 to obtain a sequence $\mathcal{K} = \mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_{t-1}, \mathcal{K}_t = \mathcal{F}$ of maximal nested families of subsets of $E(G)$, such that any two consecutive families differ in exactly one element and all families contain $\mathcal{K} \cap \mathcal{F}$. Now define $\Phi_i := (\mathcal{K}_i, \alpha)$ for $i = 0, \dots, t$. Then the sequence of flags $\Phi_0, \Phi_1, \dots, \Phi_{s-1}, \Phi_s$ of \mathcal{P}_G has all the required properties. Thus the vertex-figure of \mathcal{P}_G at its vertex $(K_0, \alpha) = (F_0, \beta)$ (in fact, at any of its vertices) is strongly flag-connected.

The proof of the general case rests on the connectedness of G (recall that the latter had been assumed from the start). Thus the underlying Cayley graph $\mathcal{G}(G)$ is connected and the transpositions in \mathcal{T}_G generate the full symmetric group S_p (that is, $T_G = S_p$). Then, since $(K_m, \alpha) = (F_m, \beta)$ in \mathcal{P}_G (recall the definition of m), there exist edges f_1, \dots, f_s in $K_m (= F_m)$ and corresponding transpositions $\tau_{f_1}, \dots, \tau_{f_s}$ in \mathcal{T}_{K_m} such that $\beta = \tau_{f_s} \cdots \tau_{f_2} \tau_{f_1} \alpha$, provided $m \geq 1$. Here it is convenient to also allow $m = 0$, meaning that $\beta = \alpha$ and $s = 0$ (that is, no transpositions occur). If $m = 1$, then Φ and Ψ agree in their 1-face but not their 0-face; since K_1 consists of a single edge, we must have $s = 1$, $K_1 = \{f_1\}$, and $\beta = \tau_{f_1} \alpha$. Recall from Lemma 3.1 that $T_{K_m} = S_p$ if $m = q$.

We now proceed by induction on s . The case $s = 0$ (that is, $\beta = \alpha$) has already been settled. Now let $s \geq 1$ and hence also $m \geq 1$ (possibly $m = q$). The key inductive step is to properly join the initial flag $\Phi = (\mathcal{K}, \alpha)$ to a suitable new flag $\Lambda = (\mathcal{L}, \gamma)$ associated with a maximal nested family \mathcal{L} of subsets of $E(G)$ and the element $\gamma := \tau_{f_1} \alpha$ of S_p . At the end, since $\beta = \tau_{f_s} \cdots \tau_{f_2} \gamma$, the inductive hypothesis for $s - 1$, applied to Λ and Ψ , will yield a sequence of flags joining Λ to Ψ . Finally, then, the two sequences of flags connecting Φ to Λ and Λ to Ψ ,

respectively, can be concatenated to obtain a sequence of flags connecting Φ to Ψ .

It remains to construct Λ . To begin with, observe that in \mathcal{P}_G we have

$$(K_0, \alpha), (F_0, \tau_{f_1} \alpha) \leq (\{f_1\}, \alpha) \leq (K_m, \alpha) = (K_m, \beta),$$

since both $\tau_{f_1} \alpha \in \langle \tau_{f_1} \rangle \alpha$ and $\tau_{f_1} \alpha = (\tau_{f_s} \cdots \tau_{f_2})^{-1} \beta \in T_{K_m} \beta = T_{K_m} \alpha$, respectively. Now we can proceed in three steps as follows. First, by our observation, we may choose a flag Λ' (say) which contains (K_0, α) as its 0-face and $(\{f_1\}, \alpha)$ as its 1-face and also includes the entire subset $\Phi \cap \Psi$; this is possible even if $m = 1$. Then, since Φ and Λ' share a 0-face, we can join them by a sequence of successively adjacent flags, all containing $\Phi \cap \Lambda'$ and therefore also $\Phi \cap \Psi$. Second, substitute the 0-face of Λ' by $(F_0, \tau_{f_1} \alpha)$ to obtain a new flag Λ'' 0-adjacent to Λ' , and then append Λ'' to the existing flag sequence. Third, choose any maximal nested family $\mathcal{L} = \{L_0, \dots, L_q\}$ of sets of edges of G such that $L_1 = \{f_1\}$ and $L_j = K_j = F_j$ for all $j \in J$; this is possible, since $f_1 \in K_m$. Define

$$\Lambda := (\mathcal{L}, \tau_{f_1} \alpha) = \{(L_0, \tau_{f_1} \alpha), \dots, (L_q, \tau_{f_1} \alpha)\}.$$

Then Λ shares a vertex with Λ'' and contains $\Phi \cap \Psi$; note for the latter that $f_1 \in K_j = L_j$ for all $j \in J$. Finally, since Λ'' and Λ share a common 0-face, we can further extend the already existing flag sequence by a sequence of successively adjacent flags, which all contain $\Lambda'' \cap \Lambda$ and hence also $\Phi \cap \Psi$, and which joins Λ'' to Λ . We have now joined Φ to Λ , and we are done. ■

We also need to prove that \mathcal{P}_G satisfies the diamond condition.

Lemma 4.4 *For each $i = 1, \dots, n-1$, given an $(i-1)$ -face (K, α) and an $(i+1)$ -face (L, β) of \mathcal{P}_G such that $(K, \alpha) \leq (L, \beta)$, there exist exactly two i -faces (J, γ) of \mathcal{P}_G such that $(K, \alpha) \leq (J, \gamma) \leq (L, \beta)$. Similarly, given a 1-face (L, β) of \mathcal{P}_G , there exist exactly two 0-faces (J, γ) of \mathcal{P}_G such that $(J, \gamma) \leq (L, \beta)$.*

Proof. For the first part, suppose $(K, \alpha) \leq (J, \gamma) \leq (L, \beta)$ in \mathcal{P}_G . Then, as explained earlier, we may assume that $\gamma = \alpha = \beta$. Moreover, $K \subset J \subset L$ and K, J, L contain exactly $i-1, i$ or $i+1$ elements, respectively. This leaves precisely two possible choices for J . Conversely, any such choice determines an i -face (J, α) such that $(K, \alpha) \leq (J, \alpha) \leq (L, \alpha)$.

For the second part, suppose $(J, \gamma) \leq (L, \beta)$ in \mathcal{P}_G . Now $J = \emptyset$ and L has only one element. If $L = \{e\}$ (say), then necessarily $\gamma \in T_L \beta = \langle \tau_e \rangle \beta = \{\beta, \tau_e \beta\}$.

This leaves only two choices for γ , namely β or $\tau_e\beta$. Conversely, either of these choices yields a 0-face (J, γ) with $(J, \gamma) \leq (L, \beta)$. ■

Finally, then, we need to append a face of rank -1 as smallest face of \mathcal{P}_G . In particular, from now on, \mathcal{P}_G will denote the extended partially ordered set.

The symmetric group S_p acts faithfully on \mathcal{P}_G as a group of polytope automorphisms. In fact, every $\gamma \in S_p$ determines an element of $\Gamma(\mathcal{P}_G)$, again denoted by γ , given by

$$(K, \alpha) \rightarrow (K, \alpha\gamma) \quad (K \subseteq E(G), \alpha \in S_p). \quad (3)$$

(Note here that γ is realized by right multiplication on the second component, not by left multiplication. Left multiplication would not in general lead to an incidence preserving mapping of \mathcal{P}_G .) It is straightforward that the group S_p of all such polytope automorphisms acts simply transitively on the vertices of \mathcal{P}_G . In general, however, S_p is only a proper subgroup of the full automorphism group of \mathcal{P}_G . The precise relationship is clarified in Theorem 5.1.

Call an abstract q -polytope *simple* if all its vertex-figures are isomorphic to the $(q - 1)$ -simplex. In \mathcal{P}_G , the flags which contain a given vertex (\emptyset, α) are all of the form (\mathcal{K}, α) , where \mathcal{K} is a maximal nested family of subsets of $E(G)$. It follows that \mathcal{P}_G is simple. Moreover, if (K, α) is a facet (that is, $(q - 1)$ -face) of \mathcal{P}_G containing (\emptyset, α) , then $|K| = q - 1$ and K is obtained from $E(G)$ by removing a single element. This gives exactly q choices for K and hence q choices of facets containing (\emptyset, α) . Therefore, since S_p acts vertex-transitively on \mathcal{P}_G , there are exactly q types of facets in \mathcal{P}_G , each type occurring at every vertex of \mathcal{P}_G . Note here that it can happen that two distinct facets (K, α) and (L, α) containing a given vertex (\emptyset, α) determine the same coset in S_p , that is, $T_K\alpha = T_L\alpha$ (the facets still are distinct because K and L are distinct).

In summary, we have established the following theorem.

Theorem 4.2 *Let G be a connected (p, q) -graph. Then \mathcal{P}_G is a simple abstract polytope of rank q with $p!$ vertices and $p!q!$ flags. Moreover, the symmetric group S_p acts simply transitively on the vertices of \mathcal{P}_G .*

Note that \mathcal{P}_G is a 0- or 1-polytope, respectively, if G is the trivial graph (with a single vertex and no edge) or a graph with a single edge (and two vertices). If G has exactly two edges, then \mathcal{P}_G is a hexagon (see also Theorem 6.1).

The polytope \mathcal{P}_G of Theorem 4.2 is called the *graphicahedron associated with G* , or simply the *G -graphicahedron*. The following theorem lies at the heart of our construction and is largely responsible for the use of the term “graphicahedron”.

Theorem 4.3 *Let G be a connected (p, q) -graph. Then the 1-skeleton of \mathcal{P}_G is isomorphic to the Cayley graph $\mathcal{G}(G)$ associated with G .*

Proof. Clearly, the vertices of \mathcal{P}_G can be identified with the elements of S_p , via $(\emptyset, \alpha) \rightarrow \alpha$. Moreover, in \mathcal{P}_G , two vertices (\emptyset, α) and (\emptyset, β) are incident with a common 1-face $(\{e\}, \gamma)$ (say) if and only if $\langle \tau_e \rangle \gamma = \{\alpha, \beta\}$; that is, if and only if $\beta = \tau_e \alpha$. Thus, adjacency in the 1-skeleton of \mathcal{P}_G corresponds precisely to adjacency in the Cayley graph $\mathcal{G}(G)$. ■

Cayley graphs are a fruitful source for the construction of polytopes. For two quite different applications of Cayley graph techniques to polytopes see also Monson-Weiss [15] and Pellicer [16].

5 The group of the graphicahedron

The automorphism group $\Gamma(\mathcal{P}_G)$ of the graphicahedron \mathcal{P}_G will usually be larger than the underlying symmetric group S_p . In fact, this happens precisely when G has non-trivial graph automorphisms (symmetries). Let $\Gamma(G)$ denote the group of graph automorphisms of G . Note that, if G is connected, then $\Gamma(G)$ is faithfully represented by its actions on the vertices of G , and on the edges of G provided G has at least two edges or no edge.

Theorem 5.1 *Let G be a connected (p, q) -graph, and let $q \neq 1$. Then $\Gamma(\mathcal{P}_G) = S_p \times \Gamma(G)$.*

Proof. First we show that $\Gamma(\mathcal{P}_G)$ contains a subgroup isomorphic to a semi-direct product $S_p \times \Gamma(G)$. We already know that S_p can be viewed as a subgroup of $\Gamma(\mathcal{P}_G)$ acting by right multiplication on the second component of the faces of \mathcal{P}_G .

Now consider $\Gamma(G)$. Recall that $q \geq 2$. Every graph automorphism κ of G is an incidence preserving permutation of the vertices and edges of G and hence determines two mappings. First, κ clearly acts on the subsets of $E(G)$ via its action on the q edges of G . Second, κ induces a group automorphism on S_p

via its action on the p vertices of G , namely through conjugation in S_p by κ ; in particular, if e is an edge of G , then $\kappa(e)$ is also an edge of G and $\kappa\tau_e\kappa^{-1} = \tau_{\kappa(e)}$ in S_p . Hence, κ also determines the following mapping on \mathcal{P}_G , again denoted by κ :

$$(K, \alpha) \rightarrow (\kappa(K), \alpha^\kappa) \quad (K \subseteq E(G), \alpha \in S_p), \quad (4)$$

with $\alpha^\kappa := \kappa\alpha\kappa^{-1}$. We need to show that this is indeed a polytope automorphism of \mathcal{P}_G .

First note that in S_p we have

$$T_{\kappa(K)} = \langle \tau_{\kappa(e)} \mid e \in K \rangle = \kappa \langle \tau_e \mid e \in K \rangle \kappa^{-1} = \kappa T_K \kappa^{-1}$$

and hence also

$$T_{\kappa(K)}\alpha^\kappa = \kappa T_K \kappa^{-1} \kappa \alpha \kappa^{-1} = \kappa T_K \alpha \kappa^{-1}.$$

Therefore, if $K \subseteq L$ and $\alpha, \beta \in S_p$, then $T_K\alpha \subseteq T_L\beta$ if and only if $T_{\kappa(K)}\alpha^\kappa \subseteq T_{\kappa(L)}\beta^\kappa$. In \mathcal{P}_G , then this says that $(K, \alpha) \leq (L, \beta)$ if and only if $(\kappa(K), \alpha^\kappa) \leq (\kappa(L), \beta^\kappa)$. Thus κ is an incidence preserving bijection, that is, a polytope automorphism of \mathcal{P}_G .

By slight abuse of notation, we let $\Gamma(G)$ also denote the subgroup of $\Gamma(\mathcal{P}_G)$ consisting of all polytope automorphism κ of \mathcal{P}_G derived in this way from graph automorphisms of G . Note here that this is a faithful copy of the group of graph automorphisms of G ; in fact, if every face (K, α) in (4) is fixed under κ , then the corresponding graph automorphism must map every subset K of $E(G)$ to itself and so must be the identity mapping since G is connected.

At this point we know that both S_p and $\Gamma(G)$ can be viewed naturally as subgroups of $\Gamma(\mathcal{P}_G)$. These subgroups intersect only trivially, since every polytope automorphism in S_p leaves the first component K of every face (K, α) unchanged, while only the trivial polytope automorphism in $\Gamma(G)$ has this property. Moreover, the subgroup S_p is invariant under conjugation in $\Gamma(\mathcal{P}_G)$ by a polytope automorphisms κ in $\Gamma(G)$. In fact, if $\gamma \in S_p$, then the polytope automorphism $\kappa\gamma\kappa^{-1}$ takes the face (K, α) of \mathcal{P}_G to

$$(K, (\alpha^{\kappa^{-1}\gamma}^\kappa)) = (K, \alpha\gamma^\kappa),$$

and hence must coincide with the polytope automorphism determined by the element γ^κ of the underlying symmetric group S_p ; as a graph automorphism, κ is a permutation of the vertices, so γ^κ really is an element of S_p . Thus S_p is normalized by $\Gamma(G)$, and we have a semi-direct product $S_p \rtimes \Gamma(G)$ realized as a subgroup of $\Gamma(\mathcal{P}_G)$.

It remains to show that every polytope automorphism of \mathcal{P}_G lies in this semi-direct product. Suppose ρ is a polytope automorphism of \mathcal{P}_G . Then ρ takes

vertices of \mathcal{P}_G to vertices of \mathcal{P}_G , and in particular takes (\emptyset, ϵ) to a vertex (\emptyset, γ) (say) with $\gamma \in S_p$. (As before, ϵ denotes the identity element in S_p .) Under γ^{-1} , viewed as a polytope automorphism of \mathcal{P}_G , this new vertex is mapped back to (\emptyset, ϵ) . It follows that the polytope automorphism $\kappa := \gamma^{-1}\rho$ maps (\emptyset, ϵ) to itself and hence yields a polytope automorphism of the vertex-figure of \mathcal{P}_G at this vertex. Our goal is to show that κ is contained in the subgroup $\Gamma(G)$ of $\Gamma(\mathcal{P}_G)$. Then this would imply

$$\rho = \gamma\kappa \in S_p \cdot \Gamma(G),$$

as required.

We know that a polytope automorphism is completely determined by its effect on a single flag; this is implied by the strong flag-connectedness. Now pick any flag of \mathcal{P}_G of the form $\Phi := (\mathcal{K}, \epsilon)$, where, as usual, $\mathcal{K} = \{\emptyset, K_1, \dots, K_q\}$ is a maximal nested family of subsets of $E(G)$. Note that (\emptyset, ϵ) is the vertex contained in Φ . Since κ preserves the vertex-figure of \mathcal{P}_G at (\emptyset, ϵ) , the image of Φ under κ is a flag $\Lambda := (\mathcal{L}, \epsilon)$, where $\mathcal{L} = \{\emptyset, L_1, \dots, L_q\}$ is a maximal nested family of subsets of $E(G)$. In particular, κ maps (K_j, ϵ) to (L_j, ϵ) for each $j = 1, \dots, q$.

Suppose the edges of G have been labeled f_1, \dots, f_q and g_1, \dots, g_q , so that $K_j = \{f_1, \dots, f_j\}$ and $L_j = \{g_1, \dots, g_j\}$ for each j . Define the permutation κ' of $E(G)$ by $\kappa'(f_j) = g_j$ for each j . Then $\kappa'(K_j) = L_j$ for each j , so in particular, κ maps (K_j, ϵ) to $(\kappa'(K_j), \epsilon)$ for each j . Suppose we already know that κ' is the edge permutation of G determined by a graph automorphism of G , again denoted by κ' . Then (4) shows that κ' , viewed as a polytope automorphism, maps the face (K_j, ϵ) of \mathcal{P}_G to

$$(\kappa'(K_j), \epsilon^{\kappa'}) = (\kappa'(K_j), \epsilon) = (L_j, \epsilon),$$

for each j , and therefore maps Φ to Λ . Since polytope automorphisms are uniquely determined by their effect on a single flag, this would imply that $\kappa = \kappa' \in \Gamma(G)$, as desired.

Therefore we must prove that κ' is the edge permutation of G given by a graph automorphism of G . Here it remains to find the action of the latter on the vertices of G . We show that κ' determines a vertex permutation κ'' (say) of G , such that κ' and κ'' together yield an incidence preserving bijection of G , that is, a graph automorphism of G .

First note that a graph automorphism of a connected graph is completely determined by its effect on the vertices of valency at least 2 and their neighbors of valency 1. Thus we may ignore vertices of G of valency 1, since we know the effect on the edges that contain them. Let j be a vertex of G of valency at least

2, and let $E(j)$ be the set of edges of G that contain vertex j . It suffices to show that, given any vertex j of G of valency at least 2, the image of $E(j)$ under κ' is contained in a set $E(j'')$ for some vertex j'' of G (which then trivially must have valency at least 2). In fact, once this has been proved, we can simply define the desired vertex permutation κ'' by $\kappa''(j) = j''$ for any vertex j of G of valency at least 2; this determines κ'' completely, and we are done.

The key for the final step is to remember that κ' was derived from the above polytope automorphism κ of the underlying graph \mathcal{P}_G . We know that the face structure of \mathcal{P}_G is completely determined by the structure of G . For example, the 2-face $(\{e, f\}, \epsilon)$ of \mathcal{P}_G is a hexagon or square, depending on whether or not e, f are edges of G with a common vertex; note here that $T_{\{e, f\}} = \langle \tau_e, \tau_f \rangle = S_3$ or $C_2 \times C_2$, according as e and f share a vertex or not. On the other hand, polytope automorphisms preserve isomorphism types of faces, so κ must necessarily map square 2-faces to square 2-faces and hexagonal 2-faces to hexagonal 2-faces; in particular, $T_{\{\kappa'(e), \kappa'(f)\}} \cong T_{\{e, f\}}$. Hence, the edges e and f of G share a common vertex in G if and only if their images $\kappa'(e)$ and $\kappa'(f)$ share a common vertex in G .

Finally, then, let j be any vertex of G of valency $s \geq 2$, and let $E(j)$ be as above. When $s = 2$ we are done by the previous argument; in fact, if e and f are the two edges of $E(j)$ (sharing vertex j), then $\kappa'(e)$ and $\kappa'(f)$ also share a vertex, j'' , and $\kappa'(E(j)) \subseteq E(j'')$. Suppose $s \geq 3$. We now consider 3-faces of \mathcal{P}_G . If e, f, g are any edges in $E(j)$ (all sharing the vertex j), then we must have $T_{\{e, f, g\}} = \langle \tau_e, \tau_f, \tau_g \rangle = S_4$; in fact, if k, l, m (say) are the end vertices of e, f, g distinct from j , respectively, then

$$\tau_e = (j k), \quad \tau_f = (j l), \quad \tau_g = (j m),$$

and these transpositions generate all permutations of j, k, l, m . In particular, the 3-face $(\{e, f, g\}, \epsilon)$ of \mathcal{P}_G must have 24 vertices, so the same must be true for its image $(\{\kappa'(e), \kappa'(f), \kappa'(g)\}, \epsilon)$ under κ . Hence $T_{\{\kappa'(e), \kappa'(f), \kappa'(g)\}} = \langle \tau_{\kappa'(e)}, \tau_{\kappa'(f)}, \tau_{\kappa'(g)} \rangle$ must also have order 24. Since we already know from our previous considerations that any two of $\kappa'(e), \kappa'(f), \kappa'(g)$ share a common vertex in G , this now forces all three to have a vertex in common (otherwise we would only obtain a subgroup S_3). It follows that κ' maps any three edges e, f, g in $E(j)$ to three concurrent edges. However, if any three edges in $E(j)$ are mapped to three concurrent edges, then the entire set $E(j)$ is mapped to a set of edges of G all sharing a common vertex, j'' . Thus $\kappa'(E(j)) \subseteq E(j'')$, and the proof is complete. ■

Note that Theorem 5.1 fails if G is a connected graph with only one edge. In fact, in this case \mathcal{P}_G is a 1-polytope with automorphism group S_2 , while $\Gamma(G)$

also has order 2 (but has a trivial action on the edge set of G).

The following theorem characterizes the graphicahedra which are regular polytopes. By P_n and C_n , respectively, we denote the path or cycle of length n (with n edges), and by $K_{1,n}$ the star graph with n edges emanating from a central vertex. Note that $K_{1,n} = P_n$ if $n = 0, 1, 2$, so in particular $K_{1,0}$ is the trivial graph (with a single vertex and no edge).

Corollary 5.1 *Let G be a connected (p, q) -graph. Then \mathcal{P}_G is regular if and only if $G = C_3$ or $G = K_{1,q}$ with $q \geq 0$.*

Proof. If $q = 0$ or 1 , then \mathcal{P}_G is a 0- or 1-polytope and hence is regular; in this case $G = K_{1,q}$. Hence we may assume that $q \geq 2$.

A polytope is regular if and only if the order of its automorphism group is the same as the number of its flags (see [12, 2A5]). By Theorem 5.1, $\Gamma(\mathcal{P}_G) = S_p \times \Gamma(G)$; and by Theorem 4.2, \mathcal{P}_G has $p!q!$ flags. Hence \mathcal{P}_G is regular if and only if $\Gamma(G)$ has order $q!$. Since graph automorphisms are completely determined by their effect on the edges of the graph, the latter is equivalent to saying that $\Gamma(G) = S_q$, or, that every permutation of edges of G arises from a graph automorphism of G .

It remains to enumerate the graphs G with q edges and with $\Gamma(G) = S_q$. Now suppose G is a graph with q edges and with $\Gamma(G) = S_q$.

First we show that G is a tree unless $G = C_3$. Suppose C is a cycle of smallest length, k (say), in G . Let f_1, \dots, f_k denote the edges of C , in cyclic order. Now, since $\Gamma(G) = S_q$, if f is any edge of G , then the edges f, f_2, \dots, f_k must also form a k -cycle, C_f (say), of G ; in fact, any edge permutation that maps f_1 to f and fixes each f_j for $j \geq 2$ must necessarily come from a graph automorphism that maps C to a k -cycle of G . But then $C_f = C$ and hence $f = f_1$, since the two cycles share all but one edge (double edges are not permitted). Hence the edges of C comprise all the edges of G , and hence G itself is a cycle of length $k = q$. The automorphism group of a cycle of length q is a dihedral group of order q , and is isomorphic to S_q only if $q = 3$. This leaves C_3 as the only possibility. Thus G is acyclic unless $G = C_3$.

If G has only vertices of valency at most 2, then G is the path of length q . Since the automorphism group of a path of any length has order 2, this forces $q = 2$, that is, $G = P_2 = K_{1,2}$.

Now suppose G has a vertex j (say) of valency at least 3. Let e, f, g be edges of G containing j . If e' is any edge of G , then e' must also contain j ; in fact, any edge permutation that maps e to e' and fixes f and g must necessarily come from

a graph automorphism that fixes the vertex j and hence takes e to an edge, e' , that also contains j . Thus every edge of G contains the vertex j of G . It follows that $G = K_{1,q}$.

In summary, we have shown that G must be a graph C_3 or $K_{1,q}$ with $q \geq 2$. Conversely, each such graph G has S_3 or S_q , respectively, as its automorphism group, so that \mathcal{P}_G must be regular. ■

6 Examples

In this section we discuss in more detail the geometric structure of the graphicahedron \mathcal{P}_G for some specific choices of (p, q) -graph G . Recall that G must be connected.

We begin with the path $G = P_n$ of length n , for which $p = n + 1$ and $q = n$. As mentioned earlier, the graphicahedron associated with a given graph G generalizes the permutahedron Π_n . The following theorem shows that \mathcal{P}_G is in fact isomorphic to Π_n if $G = P_n$.

Theorem 6.1 *The graphicahedron \mathcal{P}_{P_n} associated with a path P_n of length n is isomorphic to the face-lattice of the permutahedron Π_n of rank n .*

Proof. If we knew that \mathcal{P}_{P_n} could be realized as a convex n -polytope in some euclidean space, then, bearing in mind Lemma 3.3 and Theorem 4.3, we could simply appeal to the well-known fact that a simple convex polytope of a given dimension is uniquely determined up to isomorphism by its 1-skeleton (see [2, 10]). However, we do not a priori have such a realization and must proceed in a different manner. We outline the proof.

We use some basic Coxeter group theory for the symmetric group S_{n+1} , or, equivalently, the symmetry group of the regular n -simplex. The permutahedron Π_n was defined earlier as the convex hull of the $(n+1)!$ points (in a hyperplane) in \mathbb{R}^{n+1} obtained from $(1, 2, \dots, n+1)$ by permuting the coordinates in all possible ways. However, Π_n has other equivalent realizations. In fact, Π_n can also be obtained from Wythoff's construction applied to an interior point x of the standard fundamental simplex for the symmetry group S_{n+1} of the regular n -simplex (see Coxeter [4]); in other words, Π_n is the convex hull of all images of x under this group. Since the (interior) walls of the fundamental simplex are in one-to-one correspondence with the standard generating system of adjacent transpositions

(1 2), (2 3), \dots , (n $n+1$) of S_{n+1} , it is immediately clear that the 1-skeleton of Π_n is the Cayley graph of S_{n+1} . In the notation of [4, §11.6-7], Π_n may be represented by a Coxeter diagram A_n in which all nodes are ringed, and the structure of its faces can be read off this diagram. Combinatorially, the boundary complex of Π_n is the dual of the barycentric subdivision of the boundary complex of the regular n -simplex. Comparison with the graphicahedron \mathcal{P}_{P_n} then shows that the faces of Π_n may be viewed as geometric realizations of those of \mathcal{P}_{P_n} . In particular, a 0-face (\emptyset, α) of \mathcal{P}_G corresponds to the vertex $\alpha(x)$ of Π_n , and, more generally, a face (K, α) of \mathcal{P}_G corresponds (via the coset $T_K\alpha$) to the convex hull of the images of $\alpha(x)$ under the subgroup T_K of the symmetry group S_{n+1} . Thus the graphicahedron \mathcal{P}_{P_n} is isomorphic to the face-lattice of Π_n . ■

Next we analyse the graphicahedra associated with more general kinds of connected graphs G . We are particularly interested in graphs with few vertices and edges.

For the n -cycle C_n with $n = 3$ or 4 it is not hard to verify directly that \mathcal{P}_{C_n} is a vertex-transitive tessellation of the 2-torus or 3-torus, respectively. In particular, by Corollary 5.1, \mathcal{P}_{C_3} is a regular 3-polytope isomorphic to the toroidal map $\{6, 3\}_{(1,1)}$ (see [5]). The structure of \mathcal{P}_{C_n} for general n will be analyzed in the forthcoming paper [6]; in fact, \mathcal{P}_{C_n} is always a tessellation on the $(n - 1)$ -torus.

Next let $K_{1,n}$ be a star graph with n edges emanating from a central vertex, so that $p = n + 1$ and $q = n$. Then $\mathcal{P}_{K_{1,n}}$ is a regular n -polytope for each $n \geq 0$, with automorphism group $S_{n+1} \times S_n$ (see Theorem 5.1). Except for the C_3 -graphicahedron, these are the only graphicahedra which are regular polytopes.

When $n = 3$ we obtain the toroidal regular map $\mathcal{P}_{K_{1,3}} = \{6, 3\}_{(2,2)}$. This occurs as the facet type of the regular 4-polytope $\mathcal{P}_{K_{1,4}}$. In particular, $\mathcal{P}_{K_{1,4}}$ is the universal regular 4-polytope $\{\{6, 3\}_{(2,2)}, \{3, 3\}\}$ with automorphism group $S_5 \times S_4$ and with 20 toroidal facets each isomorphic to $\{6, 3\}_{(2,2)}$ (see [12, 11C8] and also [9, 11, 14]). Isomorphism of $\mathcal{P}_{K_{1,4}}$ with the universal polytope follows from that fact that both polytopes have groups of the same order; here, the semi-direct $S_5 \times S_4$ actually also is a direct product $S_5 \times S_4$ (however, with different factors). More generally, for each $n \geq 0$, the regular n -polytope $\mathcal{P}_{K_{1,n}}$ is the facet type of the regular $(n + 1)$ -polytope $\mathcal{P}_{K_{1,n+1}}$, and the latter has $(n + 1)(n + 2)$ such facets; see also [19, p.320] for a related polytope, possibly isomorphic to the graphicahedron of a star graph.

Clearly, a connected $(p, 3)$ -graph must necessarily be isomorphic to P_3 , C_3 or $K_{1,3}$, so there are no graphicahedra of rank 3 other than those already mentioned. This completely settles the case $q = 3$. On the other hand, when $q = 4$ there

are two connected $(p, 4)$ -graphs in addition to P_4 , C_4 or $K_{1,4}$ (see Figure 3). The corresponding graphicahedra of rank 4 are clearly not regular; their automorphism group is $S_5 \times C_2$ or $S_4 \times C_2$, respectively, with the factor C_2 arising from the graph symmetry. For the graph on the left, the graphicahedron has twenty-five facets, ten permutahedra, five regular toroids $\{6, 3\}_{(2,2)}$, and ten hexagonal prisms. For the graph on the right, the graphicahedron has seven facets, of which two are permutahedra, four are regular toroids $\{6, 3\}_{(1,1)}$, and one is a regular toroid $\{6, 3\}_{(2,2)}$.



Figure 3: Graphs with four edges distinct from P_4 , C_4 or $K_{1,4}$.

Acknowledgment: Some of the results described here were presented and discussed at the International Workshop on Polytopes and Symmetries held in March 2008 during the Coloquio Victor Neumann-Lara de Teoria de las Gráficas, Combinatoria y sus Aplicaciones at the Universidad Autónoma de Zacatecas in Zacatecas, Mexico. We wish to extend our thanks to the participants, including in particular Javier Bracho, Dimitri Leemans, Barry Monson, Alen Orbanic, Daniel Pellicer, Tomaž Pisanski and Asia Weiss.

References

- [1] A. Björner, M. Las Vergnas, B. Sturmfels, N. White and G. Ziegler: *Oriented Matroids, Second Edition*, Encyclopedia of Mathematics and its Applications **46**, Cambridge University Press, 1999.
- [2] R. Blind and P. Mani: *On puzzles and polytope isomorphism*, Aequationes Math. **34** (1987), 287–297.
- [3] G. Chartrand and L. Lesniak: *Graphs and Digraphs, Third Edition*, Chapman and Hall, London, UK, 1996.
- [4] H.S.M. Coxeter: *Regular Polytopes, Third Edition*, Dover, New York, 1973.

- [5] H.S.M. Coxeter and W.O.J. Moser, *Generators and Relations for Discrete Groups, Fourth Edition*, Springer, New York, 1980.
- [6] M. Del Río-Francos, I. Hubard, D. Oliveros and E. Schulte, *Face-transitive graphicahedra*, in preparation.
- [7] S. Fomin and N. Reading, *Root systems and generalized associahedra*, In: Geometric Combinatorics (E. Miller, V. Reiner and B. Sturmfels, eds.), IAS/Park City Math. Ser., 13, Amer. Math. Soc., Providence, RI, 2007, pp. 63–131.
- [8] G.-Th. Guilbaud, P. Rosenstiehl, *Analyse algébrique d'un scrutin*, Mathématiques et Sciences Humaines **4** (1963), 9–33.
- [9] B. Grünbaum, *Regularity of graphs, complexes and designs*, In: Problèmes combinatoires et théorie des graphes, Coll. Int. C.N.R.S. **260**, Orsey (1977), 191–197.
- [10] G. Kalai: *A simple way to tell a simple polytope from its graph*, J. Combinatorial Theory, Ser. A, **49** (1988), 381–383.
- [11] P. McMullen and E. Schulte, *Hermitian forms and locally toroidal regular polytopes*, Advances in Math. **82** (1990), 88–125.
- [12] P. McMullen and E. Schulte: *Abstract Regular Polytopes*, Cambridge University Press, 2002.
- [13] P. McMullen: *Space-tiling zonotopes*, Mathematika **22** (1975), 202–211.
- [14] B.R. Monson and A.I. Weiss, *Eisenstein integers and related C-groups*, Geom. Dedicata **66** (1997), 99–117.
- [15] B.R. Monson and A.I. Weiss, *Cayley graphs and symmetric 4-polytopes*, Preprint, 2008.
- [16] D. Pellicer, *CPR graphs and regular polytopes*, European Journal of Combinatorics **29** (2008), 59–71.
- [17] G.C. Shephard: *Combinatorial properties of the associated zonotope*, Canadian J. Math. **26** (1974), 302–321.
- [18] P.H. Schoute: *Analytic treatment of the polytopes regularly derived from the regular polytopes*, Verhandelingen der Koninklijke Akademie van Wetenschappen te Amsterdam **11**, No. 3, Johannes Müller, Amsterdam, 1911, 87 pages.

- [19] E. Schulte, *Regular incidence-polytopes with euclidean or toroidal faces and vertex-figures*, J. Combin. Theory, Ser. A, **40** (1985), 305–330.
- [20] N. White: *Combinatorial Geometries*, Cambridge University Press, 1987.
- [21] G. Ziegler: *Lectures on Polytopes*, Springer-Verlag, New York, 1994.