When are permutation invariants
Cohen–Macaulay over all fields?

Ben Blum-Smith and Sophie Marques

When are permutation invariants Cohen–Macaulay over all fields?

Ben Blum-Smith and Sophie Marques

We prove that the polynomial invariants of a permutation group are Cohen–Macaulay for any choice of coefficient field if and only if the group is generated by transpositions, double transpositions, and 3-cycles. This unites and generalizes several previously known results. The "if" direction of the argument uses Stanley–Reisner theory and a recent result of Christian Lange in orbifold theory. The "only if" direction uses a local-global result based on a theorem of Raynaud to reduce the problem to an analysis of inertia groups, and a combinatorial argument to identify inertia groups that obstruct Cohen–Macaulayness.

1. Introduction

The invariant ring of a graded action by a finite group G on a polynomial ring

$$k[\boldsymbol{x}] = k[x_1, \ldots, x_n]$$

over a field k is well behaved when the field characteristic is prime to the group order. For example, it is generated in degree $\leq |G|$ (Noether's bound), and it is a Cohen–Macaulay ring (the Hochster–Eagon theorem).

When the characteristic divides the group order (the *modular case*), the situation is much more mysterious. Both of these statements (and many others) can, but do not always, fail. The question of when such pathologies arise has attracted research attention over the last few decades.

In this article we focus on Cohen–Macaulayness. Let $k[\boldsymbol{x}]^G$ be the invariant ring and let

$$p = \operatorname{char} k$$

be the field characteristic. We interpret $k[\boldsymbol{x}]$ as the coordinate ring of $\mathbb{A}_k^n$, so that the action of G on $k[\boldsymbol{x}]$ is induced from an action on $\mathbb{A}_k^n$ by automorphisms. Because the action on $k[\boldsymbol{x}]$ is graded, the corresponding action on $\mathbb{A}_k^n$ is linear, i.e., it arises from a linear representation of G on a k-vector space. Here is a sampling of known results:

- Ellingsrud and Skjelbred [1980] showed that if G is cyclic of order p^m, then $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay unless G fixes a subspace of $\mathbb{A}_k^n$ of codimension ≤ 2.

- Larry Smith [1996] showed that if $n = 3$, then $k[\boldsymbol{x}]^G$ is Cohen–Macaulay. (This was known to hold for $n \leq 2$.)

- Campbell et al. [1999] showed that if G is a p-group, and if the action of G on $\mathbb{A}_k^n$ is the sum of three copies of the same linear representation, then $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay.

- Gregor Kemper [1999] showed that if G is a p-group and $k[\boldsymbol{x}]^G$ is Cohen–Macaulay, then G is necessarily generated by elements g whose fixed-point sets in $\mathbb{A}_k^n$ have codimension ≤ 2, generalizing [Ellingsrud and Skjelbred 1980] beyond cyclic groups and [Campbell et al. 1999] beyond three-copies representations.

See [Kemper 2012] for a more detailed overview.

A theme uniting these results is that generation of G by elements fixing codimension ≤ 2 subspaces is related to good behavior of $k[\boldsymbol{x}]^G$. Further variations on this theme are found in [Dufresne et al. 2009; Gordeev and Kemper 2003; Kac and Watanabe 1982; Lorenz and Pathak 2001]. The main goal of this paper is a result of this kind for permutation groups $G \subset S_n$, acting on $k[\boldsymbol{x}]$ by permuting the x_i. The result characterizes permutation groups generated in this way, and is not restricted to p-groups.

Permutation groups have the feature that the definition of the action is insensitive to the choice of a ground field k. Thus it is natural to ask:

Question 1.1. For which $G \subset S_n$ is $k[\boldsymbol{x}]^G$ Cohen–Macaulay regardless of k?

An additional motivation for this question is that $k[\boldsymbol{x}]^G$ is Cohen–Macaulay for every choice of k if and only if $\mathbb{Z}[\boldsymbol{x}]^G$ is free as a module over the subring $\mathbb{Z}[\boldsymbol{x}]^{S_n}$ of symmetric polynomials, and also if and only if $A[\boldsymbol{x}]^G$ is Cohen–Macaulay for every Cohen–Macaulay ring A. (We will not develop these equivalences here, but see [Blum-Smith 2017, §2.4.1] where the first is worked out in detail, and [Bruns and Herzog 1993, Exercise 5.1.25] for a sketch of the second in a slightly different setting.)

Kemper [2001] gave an if-and-only if criterion that determines Cohen–Macaulayness of a permutation invariant ring when p divides $|G|$ exactly once. This criterion allows one to determine Cohen–Macaulayness for many specific groups and primes, but does not in general answer Question 1.1 because few permutation groups have squarefree order. Some special cases of Question 1.1 are known:

- If G is a Young subgroup (i.e., a product of symmetric groups acting on disjoint sets), then $k[\boldsymbol{x}]^G$ is a polynomial algebra over k, so it is Cohen–Macaulay regardless of k.

- It follows from the result of Kemper [1999] quoted above that if G is a p-group, then $k[\boldsymbol{x}]^G$ cannot be Cohen–Macaulay over all fields unless G is generated by transpositions and double transpositions, or 3-cycles (and $p = 2$ or 3).

- Kemper [1999] also showed that if $G \subset S_n$ is regular (i.e., its action on

$$[n] = \{1, \ldots, n\}$$

is free and transitive), then $k[\boldsymbol{x}]^G$ is Cohen–Macaulay over every k if it is isomorphic to C_2, C_3, or $C_2 \times C_2$, but not otherwise. (In fact, in other cases, it is not Cohen–Macaulay for any k with char k dividing $|G|$.)

- Victor Reiner [1992; 2003] has shown that A_n, and the diagonally embedded $S_n \hookrightarrow S_n \times S_n \subset S_{2n}$, have invariant rings that are Cohen–Macaulay regardless of the field. (These are the S_n-cases of results he found for all finite Coxeter groups.) Patricia Hersh [2003a; 2003b] has shown the same for the wreath product $S_2 \wr S_n \subset S_{2n}$.

Our main objective in this article is to answer Question 1.1 completely. We will prove the following theorem, which unites all of these cases and ties them into the theme mentioned above.

Theorem 1.2. *Let* $G \subset S_n$. *The ring* $k[\boldsymbol{x}]^G$ *is Cohen–Macaulay for all choices of k if and only if G is generated by transpositions*, *double transpositions*, *and 3-cycles.*

Let N be the subgroup of G generated by transpositions, double transpositions, and 3-cycles. The "if" direction of Theorem 1.2, together with the Hochster–Eagon theorem [Hochster and Eagon 1971, Proposition 13], imply that the characteristics p in which $k[\boldsymbol{x}]^G$ fails to be Cohen–Macaulay must be among those that divide $[G : N]$. This implication will be discussed in more detail in the conclusion (Section 5). The "only if" direction implies that if $[G : N] > 1$, then there is at least one such characteristic p. This p is explicitly constructed in the course of the proof.

The proof of this theorem is methodologically eclectic. The "if" direction uses Stanley–Reisner theory, which relates Cohen–Macaulayness of $k[\boldsymbol{x}]^G$ to the topology of the quotient of a ball by G, and a recent result in orbifold theory by Christian Lange [2016] that characterizes the groups G such that this quotient is a piecewise-linear ball. The "only if" direction is much more algebraic. It is based on a local-global result (Theorem 3.1) reducing the Cohen–Macaulayness of a noetherian invariant ring to that of the invariant rings of its inertia groups acting on strict localizations.

Though Theorem 1.2 is specific to the situation of a polynomial ring $k[\boldsymbol{x}]$ and a permutation group G, a substantial portion of our method for the "only if" direction applies in considerably more generality. Section 2C concerns arbitrary commutative, unital rings, and the local-global result just mentioned only assumes that the invariant ring is noetherian. (Other work on Cohen–Macaulayness of invariants at the generality of noetherian rings includes [Gordeev and Kemper 2003; Lorenz and Pathak 2001].) A secondary goal of this paper is to develop these general tools, which we expect have broader applicability. The fact that Cohen–Macaulayness depends fully on the local action of the inertia groups yields information about Cohen–Macaulayness whenever inertia groups can be accessed directly and are simpler than the whole group, as in the present case.

The method of the "if" direction is similar to the methods used by Reiner [1992; 2003] and Hersh [2003a; 2003b] to prove the results mentioned above. The novelty is the application of Lange's orbifold result [2016] in place of an explicit shelling of a cell complex. The main novelties in the "only if" direction are: the local-global Theorem 3.1, its application to show that certain kinds of inertia p-groups obstruct

Cohen–Macaulayness (Proposition 3.11), and a combinatorial argument that exhibits such an inertia p-group explicitly in the case at hand (Lemma 4.5).

The organization of the paper is as follows. Section 2 collects together the needed background from commutative algebra, Stanley–Reisner theory, and piecewise-linear topology, and introduces notation that is used throughout the article. Section 3 contains the general results on Cohen–Macaulayness and inertia groups that are needed for the "only if" direction of Theorem 1.2, including the local-global Theorem 3.1 and the p-group obstruction Proposition 3.11. Section 4 proves the "if" direction of Theorem 1.2, and then using this, proves the "only if" direction. Finally, Section 5 draws out some implications and poses questions for further inquiry.

2. Background

Throughout this paper, A denotes an arbitrary commutative, unital ring, k denotes a field, p denotes the characteristic of k, $k[\boldsymbol{x}]$ denotes the polynomial ring $k[x_1, \ldots, x_n]$, $[n]$ denotes the set $\{1, \ldots, n\}$, and G denotes a finite group with a faithful action on $k[\boldsymbol{x}]$ by permutations of the x_i's, or on A by arbitrary automorphisms. In Section 4B, the prime number p will be conceptually prior to k, and k will be chosen to satisfy char $k = p$.

2A. _Cohen–Macaulayness._ Recall that the _depth_ of a local noetherian ring is the length of the longest regular sequence contained in the maximal ideal. The depth is always bounded above by the dimension. When equality is achieved, the ring is said to be _Cohen–Macaulay_. A general noetherian ring is defined to be Cohen–Macaulay if its localization at every maximal, or equivalently at every prime, is Cohen–Macaulay [Bruns and Herzog 1993, Definition 2.1.1 and Theorem 2.1.3(b)].

Although there has been work on extending the theory of Cohen–Macaulayness to the nonnoetherian setting [Hamilton and Marley 2007], in this paper we will follow tradition by regarding noetherianity as a requirement of Cohen–Macaulayness.

Cohen–Macaulayness is automatic for artinian rings, since if the dimension is zero, the depth of a localization cannot be strictly lower than this. For example, fields are Cohen–Macaulay. Noetherian regular rings, for example polynomial rings over fields, are also Cohen–Macaulay [Bruns and Herzog 1993, Corollary 2.2.6].

For our purposes it will be necessary to know how the Cohen–Macaulayness of a ring relates to that of a flat extension. The needed fact [Bruns and Herzog 1993, Theorem 2.1.7] is that if $A \to B$ is a flat extension of noetherian rings, then B is Cohen–Macaulay if and only if, for each prime ideal $\mathfrak{q}$ of B and its contraction $\mathfrak{p}$ in A, both $A_\mathfrak{p}$ and $B_\mathfrak{q}/\mathfrak{p}B_\mathfrak{q}$ are Cohen–Macaulay. It is enough to quantify this statement over maximal ideals $\mathfrak{q}$ of B. We will use this fact repeatedly in Section 3.

When a noetherian ring is finite over a regular subring, Cohen–Macaulayness is related to flatness as a module over the subring. In the traditional situation of invariant theory, this fact has a particularly nice formulation. For if $k[\boldsymbol{x}]$ is a polynomial ring over a field, and G acts by graded automorphisms, then $k[\boldsymbol{x}]^G$ is finitely generated and graded, and the Noether normalization lemma guarantees a graded

polynomial subring (generated by a *homogeneous system of parameters*) over which $k[\boldsymbol{x}]^G$ is finite. In this situation, $k[\boldsymbol{x}]^G$ is Cohen–Macaulay if and only if it is a *free* module over this subring (the *Hironaka criterion*). We will not build on this fact directly, but we mention it both because it motivates interest in Cohen–Macaulayness, and because we do use a result [Reiner 2003, Theorem A.1] that depends on it, whose proof we outline in the next section.

2B. *Combinatorial commutative algebra and* PL *topology.* The proof of the "if" direction of Theorem 1.2 relies on results in combinatorial commutative algebra and some basic facts about PL topology. For motivation, we describe the plan of the proof before recalling these results.

By work of Adriano Garsia and Dennis Stanton [1984], refined by Victor Reiner [2003], Cohen–Macaulayness of the polynomial invariant ring $k[\boldsymbol{x}]^G$ can be deduced from the Cohen–Macaulayness of the *Stanley–Reisner ring* of a certain cell complex (specifically a *boolean complex*) that depends on G. The Cohen–Macaulayness of this Stanley–Reisner ring can in turn be deduced from information about the complex that depends only on the homeomorphism class of its total space. For G generated as in Theorem 1.2, a recent result of Christian Lange [2016] hands us this topological information. This is the structure of the proof, which will be assembled in Section 4A. Here, we recall the needed results and definitions regarding boolean complexes and Stanley–Reisner rings.

Let P be a finite poset and k a field.

Definition 2.1. The *Stanley–Reisner ring* of P over k, written $k[P]$, is the quotient of the polynomial ring $k[\{y_\alpha\}_{\alpha \in P}]$, with indeterminates indexed by the elements of P, by the ideal generated by products $y_\alpha y_\beta$ indexed by incomparable pairs $\alpha, \beta \in P$.

Remark 2.2. This is a special case of a more general definition, which we will not use directly: the Stanley–Reisner ring of a simplicial complex. (We will use a further generalization — see Definition 2.5 below.) The Stanley–Reisner ring of a poset is nothing but the Stanley–Reisner ring of the *chain complex* of the poset, i.e., the simplicial complex with vertex set the elements of the poset, whose simplices are the chains in the poset. It is helpful to keep in mind that the Stanley–Reisner ring of a poset has an underlying simplicial complex as well.

Write $[n] = \{1, \ldots, n\}$. Let B_n be the boolean algebra on the set $[n]$, i.e., the set of subsets of $[n]$, ordered by inclusion. Then the Stanley–Reisner ring $k[B_n \setminus \{\varnothing\}]$ is, in a sense that can be made precise, a coarse approximation of the polynomial ring $k[\boldsymbol{x}]$. In particular, it carries a natural action of S_n via the latter's action on the set $[n]$, and if $G \subset S_n$, then $k[\boldsymbol{x}]^G$ is Cohen–Macaulay whenever $k[B_n \setminus \{\varnothing\}]^G$ is Cohen–Macaulay. This is the content of [Reiner 2003, Theorem A.1].

The proof is given in full there, and also in great detail in [Blum-Smith 2017, Section 2.5.3], and in any case is essentially a characteristic-neutral reformulation of an argument of Adriano Garsia and Dennis Stanton [1984], building on Garsia's earlier work [1980]. However, we would like this result to be better-known, so we indicate the line of proof.

As mentioned in Section 2A, a finitely generated graded k-algebra is Cohen–Macaulay if and only if it is free as a module over the subring generated by any homogeneous system of parameters. Thus,

Cohen–Macaulayness can be established by showing the existence of a module basis over such a subring. For any $G \subset S_n$, $k[\boldsymbol{x}]^{S_n}$ and $k[B_n \setminus \{\varnothing\}]^{S_n}$ are such subrings, respectively, of $k[\boldsymbol{x}]^G$ and $k[B_n \setminus \{\varnothing\}]^G$, and they are isomorphic. Thus, Cohen–Macaulayness may be passed from $k[B_n \setminus \{\varnothing\}]^G$ to $k[\boldsymbol{x}]^G$ by showing that the existence of a module basis for the former over the common subring $k[B_n \setminus \{\varnothing\}]^{S_n} \cong k[\boldsymbol{x}]^{S_n}$ implies the existence of a basis for the latter. Garsia [1980] introduced a k-linear, S_n-equivariant map $\mathscr{G} : k[B_n \setminus \{\varnothing\}] \to k[\boldsymbol{x}]$ sending

$$y_U \mapsto \prod_{i \in U} x_i,$$

where $U \in B_n \setminus \{\varnothing\}$ is any nonempty subset of $[n]$. The map $\mathscr{G}$ is first extended multiplicatively to all monomials of $k[B_n \setminus \{\varnothing\}]$, and then k-linearly to the whole ring. This map is an isomorphism of k-vector spaces, and also, in a sense made precise in [Blum-Smith 2017, Proposition 2.5.66], a coarse approximation of a ring homomorphism. In particular, for any $G \subset S_n$, if $k[B_n \setminus \{\varnothing\}]^G$ is Cohen–Macaulay, it maps an appropriately chosen $k[B_n \setminus \{\varnothing\}]^{S_n}$-basis of $k[B_n \setminus \{\varnothing\}]^G$ to a $k[\boldsymbol{x}]^{S_n}$-basis of $k[\boldsymbol{x}]^G$. This statement about bases was proven by Garsia and Stanton [1984] with $k = \mathbb{Q}$, in which case both rings are automatically Cohen–Macaulay — Garsia and Stanton's interest was in the explicit construction of bases — but it was observed in [Reiner 2003, Theorem A.1] that the argument is characteristic-neutral and so allows one to deduce Cohen–Macaulayness of $k[\boldsymbol{x}]^G$ from that of $k[B_n \setminus \{\varnothing\}]^G$ in the modular situation.

Remark 2.3. The map $\mathscr{G}$ is referred to as the *transfer map* in [Garsia 1980; Garsia and Stanton 1984; Reiner 2003]. Other authors in invariant theory [Neusel and Smith 2002; Smith 1995] use the same phrase to denote the A^G-linear map

$$\mathrm{Tr} : A \to A^G$$
$$x \mapsto \sum_{g \in G} g(x).$$

While this latter map is also called the *trace*, there are well-established usages of *transfer* to describe maps analogous to Tr in both topology and group theory, so we prefer to call $\mathscr{G}$ the *Garsia map* to avoid competition for the term and to honor Garsia's introduction of it [1980]. The present paper makes no use of the Garsia map except implicitly in quoting [Reiner 2003, Theorem A.1].

The work cited above reduces proving Cohen–Macaulayness of $k[\boldsymbol{x}]^G$ to the analogous statement for $k[B_n \setminus \{\varnothing\}]^G$. The Cohen–Macaulayness of this latter ring can be assessed using a topological criterion, following a general philosophy in Stanley–Reisner theory that the Cohen–Macaulayness of a Stanley–Reisner ring is equivalent to a condition on the homology of the underlying simplicial complex. In the present situation, $k[B_n \setminus \{\varnothing\}]^G$ is not the Stanley–Reisner ring of a poset or simplicial complex, but it turns out to be the Stanley–Reisner ring of a *boolean complex*. We recall the needed definitions:

Definition 2.4. A *boolean complex* is a regular CW complex in which every face has the combinatorial type of a simplex.
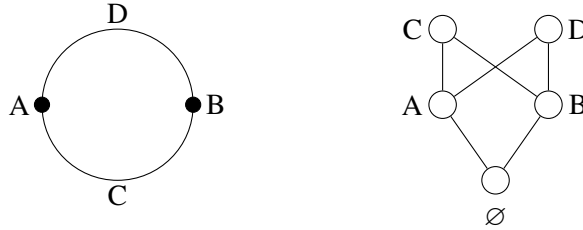
Figure 1. Left: a boolean complex with total space homeomorphic to a circle. Right: its face poset.

This is a mild generalization of a simplicial complex, in which it is possible for two faces to intersect in an arbitrary subcomplex rather than a single subface. (For example, two faces can have all the same vertices.) See Figure 1. The terminology is due to Garsia and Stanton [1984].

The *face poset* of a cell complex is the poset whose elements are the cells (*faces*), and the relation $\alpha \leq \beta$ means that α's closure is contained in β's closure. For our purposes it is convenient to modify this definition to include an additional *empty face* $\varnothing$, with $\varnothing \leq \alpha$ for all faces α. With this convention, a boolean complex can be characterized as a regular CW complex whose face poset has the property that every lower interval is a finite boolean algebra; this is the etymology of the name *boolean complex*. Face posets of boolean complexes are referred to as *simplicial posets*, a term introduced by Richard Stanley [1986].

Stanley [1991] generalized the notion of a Stanley–Reisner ring to a boolean complex Ω, as follows. Let k be a field and let Q be the face poset of Ω, including the minimal element $\varnothing$. Let $k[\{z_\alpha\}_{\alpha \in Q}]$ be a polynomial ring with indeterminates indexed by the elements of Q. Let I be the ideal of this ring generated by:

(1) The element $z_\varnothing - 1$.

(2) All products $z_\alpha z_\beta$ where $\alpha, \beta \in Q$ have no common upper bound.

(3) All elements of the form

$$z_\alpha z_\beta - z_{\alpha \wedge \beta} \sum_{\gamma \in \mathrm{lub}(\alpha, \beta)} z_\gamma$$

where α and β have at least one common upper bound and $\mathrm{lub}(\alpha, \beta)$ denotes the (consequently nonempty) set of least upper bounds of α and β.

The greatest (common) lower bound $\alpha \wedge \beta$ of α and β exists and is unique in the above formula because, as remarked above, every lower interval, and in particular the lower interval below any common upper bound for α and β, is a boolean algebra and therefore a lattice. Thus whenever α and β have any common upper bound, they have a unique greatest common lower bound in some lower interval containing them both, and thus in the whole poset.

Definition 2.5. The quotient ring $k[\{z_\alpha\}_{\alpha \in P}]/I$ is called the *Stanley–Reisner ring* of Ω and denoted $k[\Omega]$.
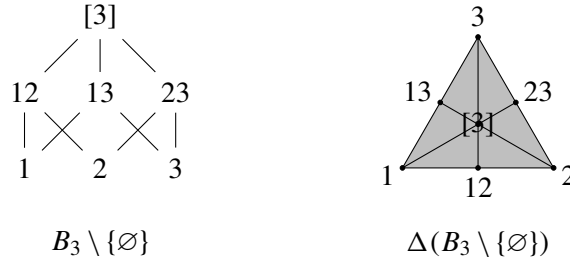
Figure 2. The poset $B_3 \setminus \{\varnothing\}$, and its order complex, which is a 2-ball.

Remark 2.6. Definition 2.5 generalizes Definition 2.1, but in a somewhat subtle way. Given a poset P, one can form its chain complex Ω, regarded as a boolean complex, and then the $k[P]$ of Definition 2.1 will be isomorphic to the $k[\Omega]$ of Definition 2.5; however, the poset Q of the latter definition will not be P. Instead, its elements will be *chains* in P, ordered by inclusion. For example, let $P = B_2 \setminus \{\varnothing\}$. Then the elements of P may be abbreviated 1, 2, and 12, and the only incomparable pair consists of 1 and 2. Thus

$$k[P] = k[y_1, y_2, y_{12}]/(y_1 y_2)$$

according to Definition 2.1. However, Q consists of the six chains in P: the empty chain $\varnothing$, three chains of length 1 (1, 2, and 12), and two chains of length 2 ($1 \subset 12$ and $2 \subset 12$). Thus

$$k[\Omega] = k[z_\varnothing, z_1, z_2, z_{12}, z_{1 \subset 12}, z_{2 \subset 12}]/I$$

where I is as described above. The isomorphism is given by mapping the z of a given chain to the product of y's corresponding to elements of the chain, for example $z_{1 \subset 12} \mapsto y_1 y_{12}$. Indeed, the definition of I becomes much more transparent after considering why this map is an isomorphism.

The ring of interest to us is the invariant ring $k[B_n \setminus \{\varnothing\}]^G$ inside the Stanley–Reisner ring of the poset $B_n \setminus \{\varnothing\}$. This ring can be identified with the Stanley–Reisner ring of a boolean complex using a result of Victor Reiner, as follows. Let Δ be the order complex of $B_n \setminus \{\varnothing\}$, i.e., the simplicial complex whose vertices are the elements of $B_n \setminus \{\varnothing\}$, and whose faces are the chains in $B_n \setminus \{\varnothing\}$. As a simplicial complex, Δ is the barycentric subdivision of an $(n-1)$-simplex, thus it is topologically an $(n-1)$-ball. See Figure 2.

The simplicial complex Δ carries a natural simplicial action of S_n, via the latter's action on $[n]$. The quotient cell complex Δ/G is usually not simplicial, but it is a boolean complex. This is because Δ is a *balanced* complex, and the action of G is a balanced action.

Definition 2.7. A boolean complex of dimension d is *balanced* if there is a labeling of its vertices by $d+1$ labels such that the vertices of any one face have distinct labels. Given such a labeling, a cellular action by a group is a *balanced action* if it preserves the labeling.
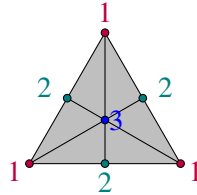
Figure 3. The labeling of the order complex of $B_3 \setminus \{\varnothing\}$, showing it is balanced.

In the present case, the vertices of Δ are the nonempty subsets of $[n]$, and thus Δ is balanced by associating a subset to its cardinality. (Here, $d = n - 1$, so the n possible cardinalities give the right number of labels.) The action of S_n is clearly balanced with respect to this labeling. See Figure 3.

It is straightforward to check that the quotient of a balanced boolean complex by a balanced action is again a balanced boolean complex. (Details are given in [Blum-Smith 2017, Lemma 2.5.86].) Thus Δ/G is a balanced boolean complex.

Victor Reiner [1992, Theorem 2.3.1] showed that if a group G acts cellularly and balancedly on a balanced boolean complex Ω, then the invariant ring $k[\Omega]^G$ inside the Stanley–Reisner ring of Ω is isomorphic to $k[\Omega/G]$, the Stanley–Reisner ring of the quotient boolean complex Ω/G. In the present situation, this gives us

$$k[\Delta/G] \cong k[B_n \setminus \{\varnothing\}]^G. \tag{1}$$

Thus the problem is reduced to showing that $k[\Delta/G]$ is Cohen–Macaulay.

Finally, the Cohen–Macaulayness of $k[\Delta/G]$ can be assessed topologically. In general, the Cohen–Macaulayness of the Stanley–Reisner ring of a boolean complex Ω is equivalent (just as for a simplicial complex) to a condition on $|\Omega|$, the underlying topological space of Ω, that depends only on its homeomorphism class. Namely, $k[\Omega]$ is Cohen–Macaulay if and only if

$$\tilde{H}_i(|\Omega|; k) = 0 \quad \text{and} \quad H_i(|\Omega|, |\Omega| - q; k) = 0 \tag{2}$$

for all points $q \in |\Omega|$ and all $i < \dim \Omega$. (Here, $H_i(|\Omega|, |\Omega|-q; k)$ is relative singular homology and $\tilde{H}_i(|\Omega|; k)$ is reduced singular homology.) This theorem is the product of work of Gerald Reisner (building on work of Melvin Hochster), James Munkres, Richard Stanley, and Art Duval. Reisner [1976] proved that for a *simplicial* complex Ω, Cohen–Macaulayness of $k[\Omega]$ is equivalent to a homological vanishing condition that a priori depends on the simplicial structure and not just the underlying topological space. Munkres [1984] showed that Reisner's condition is equivalent to the purely topological condition stated above. Richard Stanley [1991] showed that the direction

(2) is satisfied for all $q \in |\Omega|$ and $i < \dim \Omega \Rightarrow k[\Omega]$ is Cohen–Macaulay

generalizes to boolean complexes, and Art Duval [1997] showed that this generalization is bidirectional. See [Blum-Smith 2017, §2.5.2] for more details.
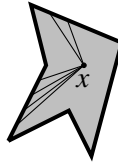
Figure 4. A compact cone neighborhood of a point in $\mathbb{R}^2$. The link K is drawn in bold, and the star S is the entire set, the union of segments from x to the points of K. Some of these segments are also drawn. Note each point of $S \setminus \{x\}$ is on exactly one such segment.

Remark 2.8. Since we only use Stanley–Reisner theory to show the "if" direction of Theorem 1.2 and thus we only need it to deduce Cohen–Macaulayness, and not the *failure* of Cohen–Macaulayness, the proof of Theorem 1.2 only uses Stanley's and not Duval's part of the generalization of (2) to boolean complexes.

Combining the results quoted above, we see that to demonstrate the Cohen–Macaulayness of the ring $k[\boldsymbol{x}]^G$, it is sufficient to prove that the boolean complex $\Omega = \Delta/G$ satisfies the homological vanishing condition (2) for all $x \in |\Delta/G|$ and all $i < n - 1$. The proof of the "if" direction of Theorem 1.2 will consist in showing that this condition holds when G is generated by transpositions, double transpositions, and 3-cycles.

This will be accomplished by quoting a recent result of Christian Lange (see Section 4A) that is stated in the language of piecewise-linear (PL) topology, so we also need to recall a few definitions and a basic fact from this field. We follow [Lange 2016, §3.1] and [Rourke and Sanderson 1972, Chapters 1 and 2] for these details. A *polyhedron* is a subset X of $\mathbb{R}^m$ in which each point has a compact cone neighborhood, i.e., given $x \in X$, there is a compact set $K \subset X$ such that (i) the union S of line segments from x to points of K is contained in X, (ii) each point of $S \setminus \{x\}$ is on a unique such line segment from x, and (iii) S is a neighborhood of x in X, i.e., it contains an open subset of X containing x. The set S is called a *star* of x in X, and K is called a *link* of x. See Figure 4.

Remark 2.9. This definition of *polyhedron* is a technical device, used here to define the concepts *piecewise-linear* and *polyhedral star*. It includes the more conventional meaning of a three-dimensional polytope as a special case, but is much, much broader. For example, any open subset of $\mathbb{R}^n$, or of any polytope, is a polyhedron.

More broadly, our use of PL topology in this paper is only to serve a technical need linking Lange's result to our setting.

If $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$ are polyhedra, a continuous map $f : X \to Y$ is a *piecewise-linear* (or PL) *map* if its graph $\{(x, f(x)) : x \in X\} \subset \mathbb{R}^{m+n}$ is a polyhedron. A *piecewise-linear* (or PL) *space* is a second-countable, Hausdorff topological space equipped with a covering by open sets U_i, each with a homeomorphism $\varphi_i : X_i \to U_i$ from a polyhedron X_i in some $\mathbb{R}^{m_i}$, such that the transition maps

$$\varphi_j^{-1} \circ \varphi_i|_{\varphi_i^{-1}(U_i \cap U_j)}$$

are PL. A PL space is a PL *manifold* (with or without boundary) if the charts X_i can be taken to be open subsets of $\mathbb{R}^n$ or the half-space $\mathbb{R}^{n-1} \times \mathbb{R}^{\geq 0}$.

A subset P of a PL space Y is called a polyhedron if for each of the charts $\varphi_i : X_i \to U_i \subset Y$, the preimage $\varphi_i^{-1}(P) \subset X_i \subset \mathbb{R}^{m_i}$ is a polyhedron.

If $X \subset \mathbb{R}^n$ is a polyhedron and $x \in X$, one may always find a link and star for x that are polyhedra [Rourke and Sanderson 1972, p.5]. It then follows from the definitions that if Y is a PL space, any point y of Y has a neighborhood S contained in some $U_i \ni y$, such that the preimage $\varphi_i^{-1}(S) \subset X_i$ is both a polyhedron and a star of $\varphi_i^{-1}(y)$ in X_i. We will refer to such an S as a *polyhedral star* of y.

The key fact we need is that if X is a polyhedron and $x \in X$, then any two polyhedral stars of x in X are PL-homeomorphic, in other words the star is a PL-homeomorphism invariant of x [Rourke and Sanderson 1972, pp.20–21]. It follows from the above discussion that the same is true in any PL space.

If Y is a PL manifold, one may take each chart X_i to be an open subset in $\mathbb{R}^n$ or $\mathbb{R}^{n-1} \times \mathbb{R}^{\geq 0}$. In any open subset of $\mathbb{R}^n$, the star of a point $(x_1, \ldots, x_n)$ may be taken to be the cube $[x_1 - \varepsilon, x_1 + \varepsilon] \times \cdots \times [x_n - \varepsilon, x_n + \varepsilon]$ for sufficiently small $\varepsilon > 0$; and in $\mathbb{R}^{n-1} \times \mathbb{R}^{\geq 0}$ it can be taken to be the intersection of this cube with the closed half-space $\{x_n \geq 0\}$. In all cases, this is topologically a closed ball. It then follows from the fact quoted in the previous paragraph that *every* polyhedral star in a PL manifold is topologically a ball.

The "if" direction of Theorem 1.2 will be proven by quoting the result of Lange mentioned above to show that if G is generated by transpositions, double transpositions, and 3-cycles, then Δ / G is a polyhedral star of a point in a PL manifold, and therefore a ball. Thus it meets the homological vanishing criterion described above, regardless of the field k.

2C. *Generalities about group actions on a ring.*
The purpose of this section is to develop the commutative algebra needed to prove the general results in Section 3, which are then used in section Section 4B to prove the "only if" direction of Theorem 1.2.

Let 1 denote the group identity. (In commutative diagrams, let it also denote a trivial group.) Let A^G denote the ring of invariants, and similarly for any subgroup of G. It is well known that A is always integral over A^G [Bourbaki 1964, Chapitre V §1.9, Proposition 22].

Let $\mathfrak{P} \lhd A$ be a prime ideal.

Recall that the *decomposition group* $D_G(\mathfrak{P})$ of $\mathfrak{P}$ is the stabilizer of $\mathfrak{P}$ in G:

$$D_G(\mathfrak{P}) = \{g \in G : g\mathfrak{P} = \mathfrak{P}\}.$$

The decomposition group acts on the integral domain $A/\mathfrak{P}$. The *inertia group* $I_G(\mathfrak{P})$ of $\mathfrak{P}$ is the kernel of this action:

$$I_G(\mathfrak{P}) = \{g \in G : (g-1)A \subset \mathfrak{P}\},$$

where

$$(g-1)A = \{ga - a : a \in A\}.$$

The notations $I_G(\mathfrak{P})$ and $D_G(\mathfrak{P})$ implicitly specify the ring A being acted on by G, since $\mathfrak{P}$ belongs to A.

We recall some basic facts in this setup [Bourbaki 1964, Chapitre V §2.2, Théorème 2], which we use freely in what follows: (i) G acts transitively on the prime ideals of A lying over $\mathfrak{P}^\star = \mathfrak{P} \cap A^G$ and (ii) the extension of residue fields $\kappa(\mathfrak{P})/\kappa(\mathfrak{P}^\star)$ is a normal field extension, and the canonical map from $D_G(\mathfrak{P})$ to the group of $\kappa(\mathfrak{P}^\star)$-automorphisms of $\kappa(\mathfrak{P})$ is a surjection with kernel $I_G(\mathfrak{P})$, i.e., the sequence

$$1 \to I_G(\mathfrak{P}) \to D_G(\mathfrak{P}) \to \mathrm{Aut}_{\kappa(\mathfrak{P}^\star)}(\kappa(\mathfrak{P})) \to 1$$

is exact.

If $N \lhd G$ is a normal subgroup, then the quotient group G/N acts on the invariant ring A^N, and the decomposition and inertia groups in G and G/N relate straightforwardly. Note that, by their definitions,

$$I_N(\mathfrak{P}) = I_G(\mathfrak{P}) \cap N \quad \text{and} \quad D_N(\mathfrak{P}) = D_G(\mathfrak{P}) \cap N.$$

Lemma 2.10. *We have*

$$D_{G/N}(\mathfrak{P} \cap A^N) \cong D_G(\mathfrak{P})/D_N(\mathfrak{P})$$

and

$$I_{G/N}(\mathfrak{P} \cap A^N) \cong I_G(\mathfrak{P})/I_N(\mathfrak{P}).$$

We believe this and the next lemma may be well known; however, as we were unable to locate references, we include full proofs.

Proof. The sequences

$$1 \to D_N(\mathfrak{P}) \to D_G(\mathfrak{P}) \xrightarrow{\varphi} D_{G/N}(\mathfrak{P} \cap A^N) \to 1$$

and

$$1 \to I_N(\mathfrak{P}) \to I_G(\mathfrak{P}) \xrightarrow{\psi} I_{G/N}(\mathfrak{P} \cap A^N) \to 1$$

are exact in the first and second positions by the definitions; we have to prove surjectivity of φ and ψ.

Consider φ first. Suppose $g \in G$ is such that its image $\bar{g}$ in G/N lies in $D_{G/N}(\mathfrak{P} \cap A^N)$. Then, setting $\mathfrak{Q} = g\mathfrak{P}$, we have

$$\mathfrak{Q} \cap A^N = \mathfrak{P} \cap A^N.$$

All primes of A that intersect A^N in $\mathfrak{P} \cap A^N$ lie in the same orbit of N. Thus there exists $n \in N$ with $n\mathfrak{Q} = \mathfrak{P}$. Therefore $ng\mathfrak{P} = \mathfrak{P}$, i.e., $ng \in D_G(\mathfrak{P})$, and we have $\varphi(ng) = \bar{g}$. So φ is surjective.

We establish the surjectivity of ψ with a diagram chase. Let $\mathfrak{P}' = \mathfrak{P} \cap A^N$ and let $\mathfrak{P}^\star = \mathfrak{P} \cap A^G$. We have the following commutative diagram:

$$
\begin{array}{ccccc}
1 & 1 & 1 & & \\
\downarrow & \downarrow & \downarrow & & \\
1 \longrightarrow I_N(\mathfrak{P}) \longrightarrow I_G(\mathfrak{P}) \overset{\psi}{\longrightarrow} I_{G/N}(\mathfrak{P}') & & \\
\downarrow & \downarrow{\scriptstyle i_G} & \downarrow{\scriptstyle i_{G/N}} & & \\
1 \longrightarrow D_N(\mathfrak{P}) \overset{\iota_D}{\longrightarrow} D_G(\mathfrak{P}) \overset{\varphi}{\longrightarrow} D_{G/N}(\mathfrak{P}') \longrightarrow 1 & & \\
{\scriptstyle p_N}\downarrow & \downarrow{\scriptstyle p_G} & \downarrow{\scriptstyle p_{G/N}} & & \\
1 \longrightarrow \mathrm{Aut}_{\kappa(\mathfrak{P}')}(\kappa(\mathfrak{P})) \overset{\iota_\kappa}{\longrightarrow} \mathrm{Aut}_{\kappa(\mathfrak{P}^\star)}(\kappa(\mathfrak{P})) \overset{\xi}{\longrightarrow} \mathrm{Aut}_{\kappa(\mathfrak{P}^\star)}(\kappa(\mathfrak{P}')) \longrightarrow 1 & & \\
\downarrow & \downarrow & \downarrow & & \\
1 & 1 & 1 & &
\end{array}
$$

where $\kappa(\mathfrak{P})$, $\kappa(\mathfrak{P}')$, and $\kappa(\mathfrak{P}^\star)$ are the residue fields. The first and second row are exact by what we have just done. The third row is exact by consideration of the definitions and the fact that $\kappa(\mathfrak{P})$ is normal over $\kappa(\mathfrak{P}')$ (by [Bourbaki 1964, Chapitre V §2.2, Théorème 2(ii)], as recalled above), since field automorphisms always extend to normal extensions. The columns are also exact by the same theorem.

Let $g \in I_{G/N}(\mathfrak{P}')$ be arbitrary and consider $i_{G/N}(g)$. Since φ is surjective, there is a $y \in D_G(\mathfrak{P})$ with $\varphi(y) = i_{G/N}(g)$. Then

$$
1 = p_{G/N} \circ i_{G/N}(g) = p_{G/N} \circ \varphi(y) = \xi \circ p_G(y),
$$

so that $p_G(y) \in \ker \xi = \mathrm{im}\, \iota_\kappa$. Thus there is a $z \in \mathrm{Aut}_{\kappa(\mathfrak{P}')}(\kappa(\mathfrak{P}))$ with $\iota_\kappa(z) = p_G(y)$. Since p_N is surjective, we have a $z' \in D_N(\mathfrak{P})$ with $p_N(z') = z$. Now consider

$$
y^\star = \iota_D(z')^{-1} y \in D_G(\mathfrak{P}).
$$

We have

$$
p_G(y^\star) = p_G \circ \iota_D(z')^{-1} p_G(y) = \iota_\kappa \circ p_N(z')^{-1} p_G(y) = \iota_\kappa(z)^{-1} p_G(y) = p_G(y)^{-1} p_G(y) = 1.
$$

Thus $y^\star \in \ker p_G = \mathrm{im}\, i_G$, so there exists $g' \in I_G(\mathfrak{P})$ with $i_G(g') = y^\star$. Then

$$
i_{G/N} \circ \psi(g') = \varphi \circ i_G(g') = \varphi(y^\star) = \varphi(\iota_D(z')^{-1} y) = \varphi \circ \iota_D(z')^{-1} \varphi(y) = 1^{-1} i_{G/N}(g) = i_{G/N}(g).
$$

Since $i_{G/N}$ is injective, we can conclude $\psi(g') = g$. Thus ψ is surjective. $\qquad\square$

The inertia group of a prime that survives a base change remains stable under that base change, and the decomposition group can only shrink:

Lemma 2.11. *Let C be an arbitrary A^G-algebra, and let $B := A \otimes_{A^G} C$. Let G act on B through its action on A and trivial action on C. If there is a prime $\mathfrak{Q}$ of B pulling back to $\mathfrak{P}$ in A, then $D_G(\mathfrak{Q}) \subset D_G(\mathfrak{P})$, and $I_G(\mathfrak{Q}) = I_G(\mathfrak{P})$.*

Proof. Let $\tau : A \to B$ be the canonical map. By construction, τ is G-equivariant. Thus if $g \in G$ stabilizes $\mathfrak{Q} \lhd B$ setwise, it also stabilizes the preimage $\mathfrak{P} \lhd A$ setwise, and it follows that $D_G(\mathfrak{Q}) \subset D_G(\mathfrak{P})$.

When $g \in D_G(\mathfrak{Q})$ and therefore $\in D_G(\mathfrak{P})$, it has an induced action on both $B/\mathfrak{Q}$ and $A/\mathfrak{P}$, and the G-equivariance of τ then implies that the induced map

$$\bar{\tau} : A/\mathfrak{P} \to B/\mathfrak{Q}$$

is $\langle g \rangle$-equivariant. If also $g \in I_G(\mathfrak{Q})$, then its action on $B/\mathfrak{Q}$ is trivial. Since $\mathfrak{P}$ is the full preimage of $\mathfrak{Q}$, $\bar{\tau}$ is an injective map, and it follows that g's action on $A/\mathfrak{P}$ is also trivial, i.e., $g \in I_G(\mathfrak{P})$. Thus $I_G(\mathfrak{Q}) \subset I_G(\mathfrak{P})$.

In the other direction, suppose $g \in I_G(\mathfrak{P})$. By [Liu 2002, Chapter 1, Corollary 1.13], we have a canonical isomorphism

$$B/\tau(\mathfrak{P})B \cong A/\mathfrak{P} \otimes_{A^G} C. \tag{3}$$

Using only the fact that $g \in D_G(\mathfrak{P})$ and the G-equivariance of τ, we already know that g fixes $\mathfrak{P}$ and $\tau(\mathfrak{P})$ setwise, and thus has well-defined actions on $A/\mathfrak{P}$ and $B/\tau(\mathfrak{P})B$ that coincide via (3). But because g is actually in $I_G(\mathfrak{P})$, the action on $A/\mathfrak{P}$ is trivial, and therefore, by (3), the action of g on $B/\tau(\mathfrak{P})B$ is also trivial.

In other words, g fixes the cosets of the additive subgroup $\tau(\mathfrak{P})B$ of B setwise. Since $\mathfrak{Q}$ pulls back to $\mathfrak{P}$, it contains the image of $\mathfrak{P}$, thus we have $\mathfrak{Q} \supset \tau(\mathfrak{P})B$. Then the cosets of $\mathfrak{Q}$ are unions of cosets of $\tau(\mathfrak{P})B$, and therefore g fixes these setwise as well. In other words, g acts trivially on $B/\mathfrak{Q}$, i.e., $g \in I_G(\mathfrak{Q})$. Thus $I_G(\mathfrak{P}) \subset I_G(\mathfrak{Q})$, and we conclude $I_G(\mathfrak{P}) = I_G(\mathfrak{Q})$. $\qquad \square$

Remark 2.12. Examining the proof of Lemma 2.11, we see why the analogous equality to $I_G(\mathfrak{P}) = I_G(\mathfrak{Q})$ may fail for decomposition groups. If $g \in D_G(\mathfrak{P})$, then we do have the $\langle g \rangle$-equivariant isomorphism (3), and therefore g does act on the cosets of $\tau(\mathfrak{P})B$ in B, but the only one we know it fixes is $\tau(\mathfrak{P})B$ itself. In particular, $\mathfrak{Q}$, which may be the union of many of these cosets, need not be fixed setwise, so that $g \notin D_G(\mathfrak{Q})$.

Henceforth, let $\mathfrak{p}$ be a prime of A^G. Our goal is to show that, in a suitable sense, the local structure of A^G at $\mathfrak{p}$ is determined by the inertia group of a prime of A lying over $\mathfrak{p}$. The precise statement is Lemma 2.14 below. It is stated by Michel Raynaud [1970, Chapitre X §1, Corollaire 1], with lines of proof indicated. Because it is central to our results, we develop in detail the notation and tools that will be required to state and prove this lemma.

Let $C_{\mathfrak{p}}^{\mathrm{hs}}$ be the strict henselization (see [Raynaud 1970, Chapitre VIII, Definition 4] or [EGA I 1960, Definition 18.8.7]) of A^G at $\mathfrak{p}$, with respect to some embedding of $\kappa(\mathfrak{p})$ in its separable closure. Then $C_{\mathfrak{p}}^{\mathrm{hs}}$ is faithfully flat over $(A^G)_{\mathfrak{p}}$, and of relative dimension zero [EGA I 1960, Proposition 18.8.8(iii)]. Furthermore, $C_{\mathfrak{p}}^{\mathrm{hs}}$ and $(A^G)_{\mathfrak{p}}$ are simultaneously noetherian [EGA I 1960, Proposition 18.8.8(iv)], and

$$A_{\mathfrak{p}}^{\mathrm{hs}} := A \otimes_{A^G} C_{\mathfrak{p}}^{\mathrm{hs}}$$

is integral over $C_{\mathfrak{p}}^{\mathrm{hs}}$ (as it is a base change of the integral morphism $A^G \to A$). Moreover, G acts on $A_{\mathfrak{p}}^{\mathrm{hs}}$ via the first component of the tensor product, so that the map $A \to A_{\mathfrak{p}}^{\mathrm{hs}}$ is G-equivariant, and

$$(A_{\mathfrak{p}}^{\mathrm{hs}})^G = C_{\mathfrak{p}}^{\mathrm{hs}}$$

since $A^G \to (A^G)_{\mathfrak{p}} \to C_{\mathfrak{p}}^{\mathrm{hs}}$ is flat and the functor of invariants commutes with flat base change.

Let $\mathfrak{P}$ be a prime ideal of A lying over $\mathfrak{p}$, and let $\mathfrak{Q}$ be a prime ideal of $A_{\mathfrak{p}}^{\mathrm{hs}}$ lying over the maximal ideal of $C_{\mathfrak{p}}^{\mathrm{hs}}$ corresponding to $\mathfrak{p}$, and pulling back to $\mathfrak{P}$ in A.

From Lemma 2.11, we have that

$$I_G(\mathfrak{Q}) = I_G(\mathfrak{P}).$$

The action of G on $A_{\mathfrak{p}}^{\mathrm{hs}}$ induces an action on its ideals. Since $A_{\mathfrak{p}}^{\mathrm{hs}}$ is integral over $C_{\mathfrak{p}}^{\mathrm{hs}}$, all of its maximal ideals lie over the one maximal of $C_{\mathfrak{p}}^{\mathrm{hs}}$. Because $C_{\mathfrak{p}}^{\mathrm{hs}}$ is the invariant ring under the action of G, this implies [Bourbaki 1964, Chapitre V §2.2, Théorème 2(i)] that the maximal ideals of $A_{\mathfrak{p}}^{\mathrm{hs}}$ comprise a single orbit for the action on ideals. The maximals are therefore finite in number. We denote them by $\mathfrak{M}_1(= \mathfrak{Q}), \ldots, \mathfrak{M}_s$.

The product of canonical localization homomorphisms

$$\phi : A_{\mathfrak{p}}^{\mathrm{hs}} \to \prod_{j=1}^{s} (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j} \tag{4}$$

is an isomorphism. Indeed, $A_{\mathfrak{p}}^{\mathrm{hs}}$ is the inductive limit of $C_{\mathfrak{p}}^{\mathrm{hs}}$-finite subalgebras (since it is integral over $C_{\mathfrak{p}}^{\mathrm{hs}}$). Since $A_{\mathfrak{p}}^{\mathrm{hs}}$ has only s maximals, there exists a finite subalgebra containing s maximals. Now view $A_{\mathfrak{p}}^{\mathrm{hs}}$ as the inductive limit just of the finite subalgebras that contain this one. For each of them, the analogous product of canonical localization morphisms is an isomorphism because $C_{\mathfrak{p}}^{\mathrm{hs}}$ is henselian [Raynaud 1970, Chapitre I, §1 Définition 1 and Proposition 3]; then the statement about (4) follows because inductive limits commute with finite products.

Lemma 2.13. *If A is a noetherian ring, then $A_{\mathfrak{p}}^{\mathrm{hs}}$ is noetherian too.*

Proof. Because of the isomorphism (4), it suffices to show that the localizations of $A_{\mathfrak{p}}^{\mathrm{hs}}$ at its maximal ideals $\mathfrak{M}_j$ are noetherian rings, and because the action of G on $A_{\mathfrak{p}}^{\mathrm{hs}}$ by automorphisms is transitive on these maximals, it suffices to show this for a single maximal. We will do this by showing that there is a maximal ideal $\mathfrak{M}_j$ of $A_{\mathfrak{p}}^{\mathrm{hs}}$ such that

$$(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}$$

is isomorphic to the strict henselization of the noetherian local ring $A_{\mathfrak{P}}$, whereupon the result will follow because strict henselization preserves noetherianity [EGA I 1960, Proposition 18.8.8(iv)].

Consider the local ring $(A^G)_{\mathfrak{p}}$. By slight abuse of notation, let us call its maximal ideal $\mathfrak{p}$. Note that the residue field $\kappa(\mathfrak{p})$ is the same whether $\mathfrak{p}$ refers to the prime in A^G or in $(A^G)_{\mathfrak{p}}$, so we can write $\kappa(\mathfrak{p})$ without ambiguity. Then the maximal ideals in the ring

$$B := A \otimes_{A^G} (A^G)_{\mathfrak{p}}$$

are in bijection with the prime ideals of A lying over $\mathfrak{p} \lhd A^G$. There are finitely many of these since they are subject to a transitive action by G, so B is semilocal. It is also integral as an extension of $(A^G)_\mathfrak{p}$ since this is a base change of the integral extension $A^G \subset A$. One of the prime ideals over $\mathfrak{p}$ in A is $\mathfrak{P}$. By the same abuse of notation, let $\mathfrak{P}$ also refer to the corresponding ideal in B; again, this does not introduce ambiguity when writing $\kappa(\mathfrak{P})$. Note that $B_\mathfrak{P} = A_\mathfrak{P}$ because B is obtained from A by inverting some but not all of the elements in the complement of $\mathfrak{P}$.

Because B is semilocal and integral over $(A^G)_\mathfrak{p}$ (and $\mathfrak{P}$ and $\mathfrak{p}$ are maximal ideals of these rings respectively), if we can show that the extension of residue fields $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ has finite separable degree, then it will follow from [EGA I 1960, Proposition 18.8.10 and its proof, and Remarque 18.8.11] that the strict henselization

$$(B_\mathfrak{P})^{\mathrm{hs}}$$

of the localization $B_\mathfrak{P}$ (with respect to some embedding of its residue field in a separable closure) is isomorphic to the localization of

$$B \otimes_{(A^G)_\mathfrak{p}} C_\mathfrak{p}^{\mathrm{hs}}$$

at some maximal ideal, since $C_\mathfrak{p}^{\mathrm{hs}}$ is a strict henselization of $(A^G)_\mathfrak{p}$. But we also have

$$B \otimes_{(A^G)_\mathfrak{p}} C_\mathfrak{p}^{\mathrm{hs}} = A \otimes_{A^G} (A^G)_\mathfrak{p} \otimes_{(A^G)_\mathfrak{p}} C_\mathfrak{p}^{\mathrm{hs}} = A \otimes_{A^G} C_\mathfrak{p}^{\mathrm{hs}} = A_\mathfrak{p}^{\mathrm{hs}}.$$

Thus the conclusion from [EGA I 1960, 18.8.10 and 18.8.11] will actually be that

$$(A_\mathfrak{P})^{\mathrm{hs}} = (B_\mathfrak{P})^{\mathrm{hs}} \cong (A_\mathfrak{p}^{\mathrm{hs}})_{\mathfrak{M}_j}$$

for some maximal ideal $\mathfrak{M}_j$ of $A_\mathfrak{p}^{\mathrm{hs}}$. This is the desired conclusion, so it remains to show that $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ has finite separable degree.

Now return $\mathfrak{p}$ and $\mathfrak{P}$ to the setting of A^G and A, recalling that the residue fields $\kappa(\mathfrak{p})$ and $\kappa(\mathfrak{P})$ do not change. From [Bourbaki 1964, Chapitre V, §2.2(ii)] we have that $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is a normal field extension, and the group of $\kappa(\mathfrak{p})$-automorphisms of $\kappa(\mathfrak{P})$ is isomorphic to

$$D_G(\mathfrak{P})/I_G(\mathfrak{P}).$$

This is a subquotient of the finite group G and is therefore finite. For a normal field extension, infinite separable degree would imply infinitely many automorphisms. Thus $\kappa(\mathfrak{P})/\kappa(\mathfrak{p})$ is an extension of finite separable degree, and the proof is complete. $\square$

The action of G on $A_\mathfrak{p}^{\mathrm{hs}}$ induces, via the isomorphism ϕ of (4), an action on $\prod_1^s (A_\mathfrak{p}^{\mathrm{hs}})_{\mathfrak{M}_j}$: it is the unique action on this ring such that ϕ is G-equivariant. Because ϕ is the product of the canonical localization maps

$$\phi_j : A_\mathfrak{p}^{\mathrm{hs}} \to (A_\mathfrak{p}^{\mathrm{hs}})_{\mathfrak{M}_j},$$

it is possible to write down this action explicitly. Via the isomorphism ϕ of (4) we associate uniquely to $a \in A_{\mathfrak{p}}^{\mathrm{hs}}$ the s-tuple

$$\phi(a) = (a_{\mathfrak{M}_1}, \ldots, a_{\mathfrak{M}_s}) \in \prod_{j=1}^{s} (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j} \tag{5}$$

where each $a_{\mathfrak{M}_j}$ is the image in $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}$ of a under ϕ_j. If $g \in G$ maps $\mathfrak{M}_i$ to $\mathfrak{M}_j$, then it also induces an isomorphism

$$(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_i} \xrightarrow{g} (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}$$

$$a/s \mapsto ga/gs$$

of the localizations that makes the following square commutative:

$$
\begin{array}{ccc}
A_{\mathfrak{p}}^{\mathrm{hs}} & \xrightarrow{\ \ g\ \ } & A_{\mathfrak{p}}^{\mathrm{hs}} \\
\phi_i \downarrow & & \downarrow \phi_j \\
(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_i} & \xrightarrow[\ \ g\ \ ]{} & (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}
\end{array}
$$

By such isomorphisms, G acts on the disjoint union of the localizations $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}$. Given an $\alpha \in (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_i}$, if one chooses $a \in A_{\mathfrak{p}}^{\mathrm{hs}}$ with $\phi_i(a) = \alpha$, then the commutativity of this square can be rewritten as

$$g\alpha = \phi_j(ga).$$

Note that this statement is true regardless of the choice of a. For any such choice, writing $\alpha = a_{\mathfrak{M}_i}$ and $\phi_j(ga) = (ga)_{\mathfrak{M}_j} = (ga)_{g(\mathfrak{M}_i)}$, this becomes

$$g(a_{\mathfrak{M}_i}) = (ga)_{g(\mathfrak{M}_i)},$$

or equivalently,

$$g(a_{g^{-1}(\mathfrak{M}_j)}) = (ga)_{\mathfrak{M}_j}. \tag{6}$$

Thus, for any $a \in A_{\mathfrak{p}}^{\mathrm{hs}}$, the i-th coordinate of $\phi(a)$ determines the j-th coordinate of $\phi(ga)$, without requiring additional information about a. Then the action of G on $\prod_1^s (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{M}_j}$ induced by ϕ may be written

$$g(a_{\mathfrak{M}_1}, \ldots, a_{\mathfrak{M}_s}) = (g(a_{g^{-1}(\mathfrak{M}_1)}), \ldots, g(a_{g^{-1}(\mathfrak{M}_s)})). \tag{7}$$

Indeed, if $a \in A_{\mathfrak{p}}^{\mathrm{hs}}$, then the left side of this formula is $g\phi(a)$, and the right side is $\phi(ga)$ by (6).

Because $I_G(\mathfrak{Q})$ stabilizes $\mathfrak{Q} = \mathfrak{M}_1$, it acts on $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$. In this setting, we have the following lemma. As mentioned above, this lemma was stated by Michel Raynaud [1970, Chapitre X §1, Corollaire 1], with the proof sketched. It is the needed statement that the local structure of A^G is determined by the inertia groups. Because it is critical to our results, we give a detailed proof.

Lemma 2.14 (Raynaud). *We have a ring isomorphism*

$$(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})} \cong C_{\mathfrak{p}}^{\mathrm{hs}}.$$

Proof. Recall that $I_G(\mathfrak{P}) = I_G(\mathfrak{Q})$. Let $g_1, \ldots, g_s \in G$ be a set of left coset representatives for $G/I_G(\mathfrak{P})$, with g_1 the identity. Since $C_{\mathfrak{p}}^{\mathrm{hs}}$ is strictly henselian, its residue field is separably closed, so there are no nontrivial automorphisms of $\kappa(\mathfrak{Q})$ over it. Since the group of automorphisms of $\kappa(\mathfrak{Q})/\kappa(C_{\mathfrak{p}}^{\mathrm{hs}})$ is isomorphic to $D_G(\mathfrak{Q})/I_G(\mathfrak{Q})$, we have $D_G(\mathfrak{Q}) = I_G(\mathfrak{Q})$, so that $I_G(\mathfrak{Q})$, which equals $I_G(\mathfrak{P})$, is the stabilizer of $\mathfrak{Q}$. Thus, if we put $\mathfrak{M}_j := g_j\mathfrak{Q}$, then the ideals $\mathfrak{M}_1, \ldots, \mathfrak{M}_s$ are exactly the maximal ideals of $A_{\mathfrak{p}}^{\mathrm{hs}}$, and all of the above discussion applies.

We claim that if one restricts the canonical localization map

$$\phi_1 : A_{\mathfrak{p}}^{\mathrm{hs}} \to (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$$

to $C_{\mathfrak{p}}^{\mathrm{hs}}$, one obtains an isomorphism onto $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{Q})}$. We see this as follows:

The map ϕ_1 is the composition of ϕ with projection to the first coordinate. Because (7) makes ϕ a G-equivariant isomorphism, $a \in A_{\mathfrak{p}}^{\mathrm{hs}}$ is in $C_{\mathfrak{p}}^{\mathrm{hs}} = (A_{\mathfrak{p}}^{\mathrm{hs}})^G$ if and only if

$$(g(a_{g^{-1}(\mathfrak{M}_1)}), \ldots, g(a_{g^{-1}(\mathfrak{M}_s)})) = (a_{\mathfrak{M}_1}, \ldots, a_{\mathfrak{M}_s}) \tag{8}$$

for all $g \in G$. From (8), we will deduce the following:

(a) If $a \in C_{\mathfrak{p}}^{\mathrm{hs}}$ is an arbitrary G-invariant, then $\phi_1(a)$ is invariant under $I_G(\mathfrak{P})$. Thus $\phi_1(C_{\mathfrak{p}}^{\mathrm{hs}})$ is contained in $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$.

(b) If $a \in C_{\mathfrak{p}}^{\mathrm{hs}}$ is an arbitrary G-invariant, then all the coordinates of $\phi(a)$ are determined by the first coordinate. Thus a itself is determined by $\phi_1(a)$. In other words, the restriction of ϕ_1 to $C_{\mathfrak{p}}^{\mathrm{hs}}$ is injective.

(c) If $\alpha \in (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$ is arbitrary, there exists an $a \in C_{\mathfrak{p}}^{\mathrm{hs}}$ with $\phi_1(a) = \alpha$. Thus the restriction of ϕ_1 to $C_{\mathfrak{p}}^{\mathrm{hs}}$ is surjective.

This will suffice to establish the lemma.

To prove (a), take $g \in I_G(\mathfrak{P})$. The condition in the first coordinate of (8) is

$$g(a_{g^{-1}(\mathfrak{M}_1)}) = a_{\mathfrak{M}_1}.$$

For $g \in I_G(\mathfrak{P}) = D_G(\mathfrak{Q})$, we have $g^{-1}(\mathfrak{M}_1) = \mathfrak{M}_1 = \mathfrak{Q}$, and this condition becomes

$$g(a_{\mathfrak{Q}}) = a_{\mathfrak{Q}}.$$

Thus for the G-invariant a, we have that $a_{\mathfrak{Q}} = \phi_1(a)$ is an $I_G(\mathfrak{P})$-invariant. Therefore, $\phi_1(C_{\mathfrak{p}}^{\mathrm{hs}})$ is contained in $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$.

For (b), consider $g = g_j$ for $j = 1, \ldots, s$. The condition in the jth coordinate of (8) is

$$g(a_{g^{-1}(\mathfrak{M}_j)}) = a_{\mathfrak{M}_j}.$$

Since $g_j^{-1}(\mathfrak{M}_j) = \mathfrak{Q}$, this becomes

$$g_j(a_{\mathfrak{Q}}) = a_{\mathfrak{M}_j}.$$

Letting $j = 1, \ldots, s$, this shows that if a is a G-invariant, then all the coordinates of $\phi(a)$ are determined by $a_\mathfrak{Q}$, which is $\phi_1(a)$, so a itself is determined by $\phi_1(a)$. Therefore, the restriction of ϕ_1 to $C_\mathfrak{p}^{hs}$ is injective.

Lastly, for (c), let $\alpha \in (A_\mathfrak{p}^{hs})_\mathfrak{Q}^{I_G(\mathfrak{P})}$ be arbitrary. We construct a specific $a \in A_\mathfrak{p}^{hs}$ with $\phi_1(a) = \alpha$, and show it lies in $C_\mathfrak{p}^{hs}$. Set

$$a_{\mathfrak{M}_j} := g_j(\alpha)$$

for $j = 1, \ldots, s$, and let

$$a := \phi^{-1}(a_{\mathfrak{M}_1}, \ldots, a_{\mathfrak{M}_s}) \in A_\mathfrak{p}^{hs}.$$

Note that this a satisfies $\phi_1(a) = a_{\mathfrak{M}_1} = g_1(\alpha) = \alpha$ since g_1 is the identity. To show that it also lies in $C_\mathfrak{p}^{hs} = (A_\mathfrak{p}^{hs})^G$, it is necessary and sufficient to show that $\phi(a)$ satisfies (8) for all $g \in G$, i.e., that

$$g(a_{g^{-1}(\mathfrak{M}_j)}) = a_{\mathfrak{M}_j} \tag{9}$$

for all $g \in G$ and all $j = 1, \ldots, s$.

To do this, we first establish that

$$a_{g(\mathfrak{Q})} = g(a_\mathfrak{Q}) \tag{10}$$

for all $g \in G$, and then use this to show (9) for all g and all j.

To see (10), first recall that $\alpha = a_{\mathfrak{M}_1} = a_\mathfrak{Q}$, and then use this and $\mathfrak{M}_j = g_j(\mathfrak{Q})$ to rewrite the definition of each $a_{\mathfrak{M}_j}$:

$$a_{g_j(\mathfrak{Q})} = g_j(a_\mathfrak{Q}).$$

This establishes (10) in the particular case that g is one of $g_1, \ldots, g_s$. An arbitrary $g \in G$ has the form $g_j h$ for some g_j and some $h \in I_G(\mathfrak{P})$. Since $\mathfrak{Q}$ and $a_\mathfrak{Q} = \alpha$ are both $I_G(\mathfrak{P})$-invariant, we have

$$a_{g(\mathfrak{Q})} = a_{g_j h(\mathfrak{Q})} = a_{g_j(\mathfrak{Q})} = g_j(a_\mathfrak{Q}) = g_j h(a_\mathfrak{Q}) = g(a_\mathfrak{Q}),$$

and (10) is established for all $g \in G$.

Now we deduce (9). If $g \in G$ is arbitrary, then

$$a_{g^{-1}(\mathfrak{M}_j)} = a_{g^{-1}g_j(\mathfrak{Q})}$$

because $g_j(\mathfrak{Q}) = \mathfrak{M}_j$, and

$$a_{g^{-1}g_j(\mathfrak{Q})} = g^{-1}g_j(a_\mathfrak{Q})$$

by (10). Thus $a_{g^{-1}(\mathfrak{M}_j)} = g^{-1}g_j(a_\mathfrak{Q})$, and applying g to the left on both sides yields

$$g(a_{g^{-1}(\mathfrak{M}_j)}) = g_j(a_\mathfrak{Q}) = a_{\mathfrak{M}_j},$$

so condition (9) is met for all g and all j, i.e., (8) is met for all g. Thus

$$a \in (A_\mathfrak{p}^{hs})^G = C_\mathfrak{p}^{hs}.$$

Since $\alpha \in (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$ was arbitrary, this shows that the restriction of ϕ_1 to $C_{\mathfrak{p}}^{\mathrm{hs}}$ is surjective onto $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$, completing the proof of isomorphism. $\square$

3. Inertia groups and Cohen–Macaulayness of invariant rings

Using Lemma 2.14, we can show that the Cohen–Macaulayness of a ring of invariants at a prime ideal $\mathfrak{p}$ can always be tested in a faithfully flat neighborhood of $\mathfrak{p}$, and only depends on the action of the inertia group considered around this neighborhood. The precise statement is Theorem 3.1.

We use this to derive an obstruction to Cohen–Macaulayness for a characteristic p ring that will apply in the situation of Theorem 1.2 to prove the "only if" direction. The statement is Proposition 3.11.

In all of what follows, we use the notation of Section 2C: A is a commutative, unital ring endowed with a faithful action of a finite group G, if $\mathfrak{p}$ is a prime ideal of A^G, then $C_{\mathfrak{p}}^{\mathrm{hs}}$ is the strict henselization of A^G at $\mathfrak{p}$, and $A_{\mathfrak{p}}^{\mathrm{hs}}$ is

$$A \otimes_{A^G} C_{\mathfrak{p}}^{\mathrm{hs}},$$

with G acting through its action on A (and trivially on $C_{\mathfrak{p}}^{\mathrm{hs}}$).

Theorem 3.1. *Assume that A^G is noetherian. Then the following assertions are equivalent*:

(1) A^G *is Cohen–Macaulay.*

(2) *For every prime ideal $\mathfrak{p}$ of A^G, and for every prime ideal $\mathfrak{Q}$ of $A_{\mathfrak{p}}^{\mathrm{hs}}$ lying over $\mathfrak{p}C_{\mathfrak{p}}^{\mathrm{hs}}$ and pulling back to a prime $\mathfrak{P}$ of A lying over $\mathfrak{p}$,*

$$(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$$

is Cohen–Macaulay.

(3) *For every maximal ideal $\mathfrak{p}$ of A^G, there is some prime ideal $\mathfrak{Q}$ of $A_{\mathfrak{p}}^{\mathrm{hs}}$ lying over $\mathfrak{p}C_{\mathfrak{p}}^{\mathrm{hs}}$ and pulling back to a prime $\mathfrak{P}$ of A lying over $\mathfrak{p}$, such that*

$$(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}$$

is Cohen–Macaulay.

Proof. Clearly (2)$\Rightarrow$(3). We will show that (3)$\Rightarrow$(1) and (1)$\Rightarrow$(2).

(3)$\Rightarrow$(1): Lemma 2.14 states that for each maximal ideal $\mathfrak{p}$ of A^G and for any choice of $\mathfrak{P}$ and $\mathfrak{Q}$ as in (3),

$$C_{\mathfrak{p}}^{\mathrm{hs}} \cong (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}.$$

Thus (3) implies that $C_{\mathfrak{p}}^{\mathrm{hs}}$ is Cohen–Macaulay for each $\mathfrak{p}$. The homomorphism of local noetherian rings

$$(A^G)_{\mathfrak{p}} \to C_{\mathfrak{p}}^{\mathrm{hs}}$$

is flat, so by [Bruns and Herzog 1993, Theorem 2.1.7] quoted in Section 2A, Cohen–Macaulayness of $C_{\mathfrak{p}}^{\mathrm{hs}}$ is equivalent to that of $(A^G)_{\mathfrak{p}}$ plus that of $C_{\mathfrak{p}}^{\mathrm{hs}}/\mathfrak{p}C_{\mathfrak{p}}^{\mathrm{hs}}$. In particular, since $C_{\mathfrak{p}}^{\mathrm{hs}}$ is Cohen–Macaulay, so is $(A^G)_{\mathfrak{p}}$. Since this holds for all maximal ideals $\mathfrak{p}$ of A^G, A^G is Cohen–Macaulay.

(1)$\Rightarrow$(2): Suppose A^G is Cohen–Macaulay. Let $\mathfrak{p}$ be any prime ideal of A^G. It suffices to prove that $C_{\mathfrak{p}}^{\mathrm{hs}}$ is Cohen–Macaulay since, by Lemma 2.14, for any $\mathfrak{P}$ and $\mathfrak{Q}$ as in (2), we have

$$C_{\mathfrak{p}}^{\mathrm{hs}} \cong (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}^{I_G(\mathfrak{P})}.$$

Since $(A^G)_{\mathfrak{p}} \to C_{\mathfrak{p}}^{\mathrm{hs}}$ is flat, we again have by [Bruns and Herzog 1993, Theorem 2.1.7] that the Cohen–Macaulayness of $C_{\mathfrak{p}}^{\mathrm{hs}}$ is equivalent to that of $(A^G)_{\mathfrak{p}}$ plus that of $C_{\mathfrak{p}}^{\mathrm{hs}}/\mathfrak{p}C_{\mathfrak{p}}^{\mathrm{hs}}$. The former ring is Cohen–Macaulay since A^G is, by the hypothesis (1), and the latter is Cohen–Macaulay since it is a field (see Section 2A), namely, the residue field of the local ring $C_{\mathfrak{p}}^{\mathrm{hs}}$. $\square$

Theorem 3.1 allows us to test Cohen–Macaulayness of an invariant ring A^G locally, prime by prime, in terms of the local ring $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ and the local group action $I_G(\mathfrak{P})$. For the application we have in mind in Section 4, we will need to carry information about A and G to $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ and $I_G(\mathfrak{P})$, so we enunciate a few more lemmas to accomplish this:

Lemma 3.2. *If A is Cohen–Macaulay, then $A_{\mathfrak{p}}^{\mathrm{hs}}$ is Cohen–Macaulay for any prime ideal $\mathfrak{p}$ of A^G.*

Proof. Suppose A is Cohen–Macaulay, thus noetherian, and $\mathfrak{p}$ is a prime of A^G. By Lemma 2.13, $A_{\mathfrak{p}}^{\mathrm{hs}}$ is noetherian.

Let $\mathfrak{Q}$ be any maximal ideal of $A_{\mathfrak{p}}^{\mathrm{hs}}$ and let $\mathfrak{P}$ be its contraction in A. (Note that $\mathfrak{Q}$ lies over $\mathfrak{p}C_{\mathfrak{p}}^{\mathrm{hs}}$, per Section 2C, and therefore $\mathfrak{P}$ lies over $\mathfrak{p}$.) Now

$$A^G \to (A^G)_{\mathfrak{p}} \to C_{\mathfrak{p}}^{\mathrm{hs}}$$

is a flat map. Therefore, base changing by $A^G \to A_{\mathfrak{P}}$,

$$A_{\mathfrak{P}} \to A_{\mathfrak{P}} \otimes_{A^G} C_{\mathfrak{p}}^{\mathrm{hs}} = A_{\mathfrak{P}} \otimes_A A_{\mathfrak{p}}^{\mathrm{hs}}$$

is also a flat map. Since $\mathfrak{Q} \lhd A_{\mathfrak{p}}^{\mathrm{hs}}$ pulls back to $\mathfrak{P}$ in A, $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ is a localization of $A_{\mathfrak{P}} \otimes_A A_{\mathfrak{p}}^{\mathrm{hs}}$; thus

$$A_{\mathfrak{P}} \to (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$$

is also flat. Therefore, again by [Bruns and Herzog 1993, Theorem 2.1.7], the Cohen–Macaulayness of $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ is equivalent to that of $A_{\mathfrak{P}}$ plus that of $(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}/\mathfrak{P}(A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$. The former is Cohen–Macaulay since A is, while the latter is Cohen–Macaulay since it is an artinian local ring (see Section 2A), which in turn is because $A_{\mathfrak{P}} \to (A_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ is of relative dimension zero. This itself is because this map is a localization of the base change $A_{\mathfrak{P}} \otimes_{(A^G)_{\mathfrak{p}}}-$ of the map $(A^G)_{\mathfrak{p}} \to C_{\mathfrak{p}}^{\mathrm{hs}}$, which is flat of relative dimension zero because it is a strict henselization [EGA I 1960, Proposition 18.8.8(iii)]. $\square$

For a natural number t, an element $g \in G$ is called a t-*reflection* if the ideal generated by $(g-1)A$ in A is contained in a prime of height $\leq t$. A prime $\mathfrak{P}$ contains $(g-1)A$ if and only if $g \in I_{\mathfrak{P}}(A)$, so another way to say this is that g is a t-reflection if it is in the inertia group of some prime of height $\leq t$.

In the geometric situation (where A is a finitely generated algebra over a field), the ideal generated by $(g-1)A$ corresponds to the fixed point locus of g, so this definition makes a group element a t-reflection if this fixed point locus has codimension at most t. Thus if G is a linear group acting on the coordinate

ring of affine space, a 1-reflection is either the identity or a reflection in the classical sense. A 2-reflection has a fixed point locus of codimension 0, 1, or 2. In particular, if G acts by permutations of a basis, then the 2-reflections are exactly the identity, the transpositions, the double transpositions, and the 3-cycles.

Lemma 3.3. *If an element $g \in I_G(\mathfrak{P})$ acts as a t-reflection on $A_{\mathfrak{P}}$, then it acts as a t-reflection on A.*

Proof. Since $g \in I_G(\mathfrak{P})$, we have $(g-1)A \subset \mathfrak{P}$. The primes of A contained in $\mathfrak{P}$ are in containment-preserving bijection with the primes of $A_{\mathfrak{P}}$, with the bijection given by extension along the canonical localization map, and $(g-1)A_{\mathfrak{P}}$ is the image of $(g-1)A$ along this map. Thus if a prime of height t in $A_{\mathfrak{P}}$ contains $(g-1)A_{\mathfrak{P}}$, then its pullback in A is also of height t and contains $(g-1)A$. $\square$

Lemma 3.4. *If A is noetherian, and $g \in I_G(\mathfrak{P}) = I_G(\mathfrak{Q})$ acts as a t-reflection on $A_{\mathfrak{p}}^{\mathrm{hs}}$, then it acts as a t-reflection on A.*

Proof. If g is a t-reflection on $A_{\mathfrak{p}}^{\mathrm{hs}}$, then there is a prime ideal $\mathfrak{S}$ of $A_{\mathfrak{p}}^{\mathrm{hs}}$ of height $\leq t$ and containing $(g-1)A_{\mathfrak{p}}^{\mathrm{hs}}$. Let $\mathfrak{R}$ be $\mathfrak{S}$'s pullback in A. Then $\mathfrak{R}$ contains $(g-1)A$. Since by Section 2C and Lemma 2.13,

$$A \to A_{\mathfrak{p}}^{\mathrm{hs}}$$

is a flat extension of noetherian rings, going-down applies [Eisenbud 1995, Lemma 10.11], so that the height of $\mathfrak{S}$ is at least that of $\mathfrak{R}$. In particular, the height of $\mathfrak{R}$ is $\leq t$, so that g is a t-reflection on A. $\square$

We will also need to take an element of G acting on A but not as a t-reflection, and conclude that it does not act on a certain subring as a t-reflection either:

Lemma 3.5. *If N is the normal subgroup of G generated by the t-reflections, then no element of $G \setminus N$ acts on A^N as a t-reflection.*

Remark 3.6. This lemma does not require a noetherian hypothesis on A.

Proof. Let $g \in G$. We will show that if its image $\bar{g} \in G/N$ acts on A^N as a t-reflection, then actually $g \in N$.

If $\bar{g}$ acts on A^N as a t-reflection, then there is a prime $\mathfrak{p}$ of A^N of height $\leq t$ with $\bar{g} \in I_{G/N}(\mathfrak{p})$. Let $\mathfrak{P}$ be any prime of A lying over $\mathfrak{p}$. The height of $\mathfrak{P}$ is equal to that of $\mathfrak{p}$ (e.g., by [Gordeev and Kemper 2003, Lemma 5.3], which is stated for noetherian A but the argument holds in general); in particular it is $\leq t$. By Lemma 2.10, we have

$$I_{G/N}(\mathfrak{p}) = I_G(\mathfrak{P})/I_N(\mathfrak{P}).$$

In particular, $I_G(\mathfrak{P})$ surjects onto $I_{G/N}(\mathfrak{p})$, so there is an element $g' \in I_G(\mathfrak{P})$ whose image in G/N is $\bar{g}$. Since $\mathfrak{P}$ has height $\leq t$, g' is a t-reflection, so it is contained in N by construction. Then its image $\bar{g}$ must actually be the identity. So g (with the same image) lies in the kernel of $G \to G/N$, i.e., $g \in N$. $\square$

The following lemma allows us to detect a failure of Cohen–Macaulayness locally.

Lemma 3.7. *Let A be a ring containing the prime field $\mathbb{F}_p$, and let G be a p-group. Suppose that A is Cohen–Macaulay, A^G is noetherian, and A is finite over A^G. Further, suppose there is a prime ideal $\mathfrak{P}$ of A such that $G = I_G(\mathfrak{P})$. Then A^G is not Cohen–Macaulay unless G is generated by its 2-reflections.*

Remark 3.8. This statement is closely related to [Gordeev and Kemper 2003, Theorem 5.5], which also applies to non-p-groups and gives some control over how far A^G can be from Cohen–Macaulay. However, a key step in the proof of that result requires the rings to be normal rings that are localizations of algebras finitely generated over fields. As our application will be to rings that do not fulfill this hypothesis, we give an independent proof.

Proof of Lemma 3.7. Let N be the normal subgroup of G generated by the 2-reflections.

Since A is finite over the noetherian ring A^G, it is noetherian as an A^G-module. Since it also contains $\mathbb{F}_p$, [Lorenz and Pathak 2001, Corollary 4.3] applies, which, when specialized to the situation that G is a p-group, states that if both A and A^G are Cohen–Macaulay, then the map

$$\mathrm{Tr}_{G/N} : A^N \to A^G$$

given by

$$x \mapsto \sum_{g \in G/N} gx$$

is surjective onto A^G, where we think of each g as an element of G and the sum is taken over coset representatives of N.

We will show that this map cannot be surjective unless $N = G$. Since A is Cohen–Macaulay by assumption, this will show A^G is not Cohen–Macaulay if $N \neq G$.

If $\mathrm{Tr}_{G/N}$ is surjective, then we have

$$1 = \sum_{g \in G/N} gx$$

for some $x \in A^N$. Since $G = I_G(\mathfrak{P})$, all $g \in G$ satisfy $gx = x \mod \mathfrak{P}$ in A, thus

$$1 = \sum_{g \in G/N} x = [G : N]x \mod \mathfrak{P}$$

in A. Since G is a p-group and A contains $\mathbb{F}_p$, $[G : N]x = 0$ in A unless $N = G$. In particular, $[G : N]x$ cannot be $1 \mod \mathfrak{P}$ unless $N = G$. $\qquad\square$

Remark 3.9. The map $\mathrm{Tr}_{G/N}$ is called the *relative trace* or *relative transfer*; see Remark 2.3.

Remark 3.10. The proof uses a result of Lorenz and Pathak [2001, Lemma 4.3], which has as a hypothesis that A is noetherian as an A^G-module; call this $(\star)$. Above, we deduced $(\star)$ from the assumptions that (1) A^G is noetherian and (2) A is finite over it. Actually, $(\star)$ also implies (1) and (2), hence is equivalent to them. Since any ideal of A^G is also an A^G-submodule of A (since A^G embeds in A), $(\star)$ implies that all these ideals are finitely generated, thus (1). Meanwhile, A itself is an A^G-submodule of A, so $(\star)$ implies it is finitely generated as an A^G-module, thus (2). More generally, if a module M over a ring R has an injective R-module map from R, then noetherianity of M as R-module is equivalent to noetherianity of R as a ring plus finite generation of M over R, by the same arguments.

Combining all of these results, we get an obstruction to Cohen–Macaulayness for a characteristic p ring expressed entirely in terms of the presence of a certain inertia group. The proof of the "only if" direction of Theorem 1.2 will be an application of this proposition.

Proposition 3.11. *Let A be a ring containing $\mathbb{F}_p$ and let G be a finite group of automorphisms of A. Let N be the normal subgroup of G generated by the 2-reflections. Suppose that A^N is Cohen–Macaulay, A^G is noetherian, and A^N is finite over A^G. If there is an inertia group for the action of G/N on A^N that is a nontrivial p-group, then A^G is not Cohen–Macaulay.*

Proof. Note that

$$A^G = (A^N)^{G/N}.$$

Since A^G is noetherian, Theorem 3.1 applies.

Suppose $\mathfrak{P}$ is a prime of A^N whose inertia group $I_{G/N}(\mathfrak{P})$ is a p-group, per the hypothesis. Let

$$\mathfrak{p} = \mathfrak{P} \cap (A^N)^{G/N},$$

let $C_{\mathfrak{p}}^{\mathrm{hs}}$ be the strict henselization of $(A^G)_{\mathfrak{p}} = ((A^N)^{G/N})_{\mathfrak{p}}$, and let

$$(A^N)_{\mathfrak{p}}^{\mathrm{hs}} = A^N \otimes_{A^G} C_{\mathfrak{p}}^{\mathrm{hs}},$$

as in Section 2C.

By assumption, A^N is Cohen–Macaulay. Thus $(A^N)_{\mathfrak{p}}^{\mathrm{hs}}$ is Cohen–Macaulay, by Lemma 3.2, and thus so is

$$((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$$

for any $\mathfrak{Q} \lhd (A^N)_{\mathfrak{p}}^{\mathrm{hs}}$, and in particular any $\mathfrak{Q}$ as described in Theorem 3.1.

As A^N is finite over the noetherian ring A^G by assumption, its base change $(A^N)_{\mathfrak{p}}^{\mathrm{hs}}$ is finite over $C_{\mathfrak{p}}^{\mathrm{hs}}$, which is noetherian by [EGA I 1960, Proposition 18.8.8(iv)], as discussed in Section 2C. The localization $((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$ is a homomorphic image of $(A^N)_{\mathfrak{p}}^{\mathrm{hs}}$ by the isomorphism (4), so it too is finite over $C_{\mathfrak{p}}^{\mathrm{hs}}$.

By Lemma 2.14, $C_{\mathfrak{p}}^{\mathrm{hs}}$ is the invariant ring for the action of $I_{G/N}(\mathfrak{P})$ on $((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$. Since A contains $\mathbb{F}_p$ and therefore so do A^N and $((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$, and since $I_{G/N}(\mathfrak{P})$ is a p-group that is equal to $I_{G/N}(\mathfrak{Q})$ which is an inertia group of $((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$, we have now verified all the hypotheses of Lemma 3.7 for the action of $I_{G/N}(\mathfrak{P})$ on $((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}$. We can conclude from that lemma that the invariant ring cannot be Cohen–Macaulay unless $I_{G/N}(\mathfrak{P})$ is generated by 2-reflections.

However $I_{G/N}(\mathfrak{P})$ is not so generated. By Lemma 3.5, no nontrivial element of G/N acts on A^N as a 2-reflection. In particular, no nontrivial element of $I_{G/N}(\mathfrak{P})$ acts on A^N as a 2-reflection. Since A^N is Cohen–Macaulay, it is noetherian, so Lemma 3.4 applies, and no nontrivial element of $I_{G/N}(\mathfrak{P})$ acts on $(A^N)_{\mathfrak{p}}^{\mathrm{hs}}$ as a 2-reflection either. By Lemma 3.3, the same is true for the action of $I_{G/N}(\mathfrak{P}) = I_{G/N}(\mathfrak{Q})$ on

$$((A^N)_{\mathfrak{p}}^{\mathrm{hs}})_{\mathfrak{Q}}.$$

In particular, the p-group $I_{G/N}(\mathfrak{P})$ is not generated by 2-reflections on this ring, since it is nontrivial. Then Lemma 3.7 implies that

$$((A^N)^{\mathrm{hs}}_{\mathfrak{p}})^{I_{G/N}(\mathfrak{P})}_{\mathfrak{Q}}$$

is not Cohen–Macaulay. Therefore, by Theorem 3.1, neither is $(A^N)^{G/N} = A^G$.   □

4. Permutation invariants

In this section we prove the two directions of Theorem 1.2. A schematic diagram of the proof is found in Figure 5.

4A. *The* **if** *direction.* In this section we prove:

Proposition 4.1. *If G is generated by transpositions, double transpositions, and 3-cycles, then $k[\boldsymbol{x}]^G$ is Cohen–Macaulay regardless of the field k.*

The groundwork has been laid in Section 2B. The remaining piece of the proof is supplied by a recent, beautiful result of Christian Lange, building on earlier work of Marina Mikhaîlova. Let H be a finite subgroup of the orthogonal group $O_d(\mathbb{R})$, acting on $\mathbb{R}^d$. Endow $\mathbb{R}^d$ with its standard piecewise-linear (PL) structure. The topological quotient $\mathbb{R}^d/H$ carries a PL structure such that the quotient map $\mathbb{R}^d \to \mathbb{R}^d/H$ is a PL map, and the main result of [Lange 2016] is that it is a PL manifold (possibly with boundary) if and only if H is generated by 2-reflections. (Lange calls elements of $O_d(\mathbb{R})$ fixing a codimension-2 subspace *rotations* since they rotate a plane and fix its orthogonal complement, so he calls groups generated this way *rotation-reflection groups*.) The bulk of the work in this result lies in the "if" direction. The proof is a delicate induction on the group order, based on a complete classification of rotation-reflection groups. This classification was proven in joint work of Lange and Mikhaîlova [2016].

Proof of Proposition 4.1. Let G act on $\mathbb{R}^n$ by permutations of the axes. Let $x_1, \ldots, x_n$ be the coordinates on $\mathbb{R}^n$. The subspace

$$T = \left\{ \sum_{i=1}^n x_i = 0 \right\}$$

is G-invariant. Transpositions in G act as reflections on T, while double transpositions and 3-cycles act as rotations. Thus under the hypothesis of the proposition, G acts on T as a rotation-reflection group. By Lange's work [2016], T/G is a PL manifold.

Recall the Δ of Section 2B: it is the order complex of $B_n \setminus \{\varnothing\}$, which is the first barycentric subdivision of an $(n-1)$-simplex. Embed the underlying topological space $|\Delta|$ of Δ in T as follows. First, map the vertices of Δ to the barycenters of the standard simplex

$$\left\{ x_i \geq 0, \ \sum x_i = 1 \right\}$$

$$\Delta/G \text{ satisfies (2)} \xleftarrow{\text{Reisner etc. (§2B)}} k[\Delta/G] \text{ is CM} \xleftarrow{\text{Reiner (§2B)}} k[\Delta]^G \text{ is CM}$$

Lange's theorem (§4A)   Garsia–Stanton/Reiner (§2B)

$$G \text{ generated by 2-reflections} \qquad\qquad k[\boldsymbol{x}]^G \text{ is CM}$$

$$G \text{ not generated by 2-reflections} \qquad\qquad \text{If char } k = p, \ k[\boldsymbol{x}]^G \text{ is not CM}$$

Lemma 4.5 (§4B)

$$k[\boldsymbol{x}]^N \text{ is CM}$$

$$\exists p \text{ with } G_\pi^B N/N \cong \mathbb{Z}/p \qquad\qquad \text{Proposition 3.11 (§3)}$$

Lemma 4.4 (§4B)

$$\exists \mathfrak{P} \lhd k[\boldsymbol{x}]^N \text{ s.t. } I_{G/N}(\mathfrak{P}) \cong \mathbb{Z}/p$$
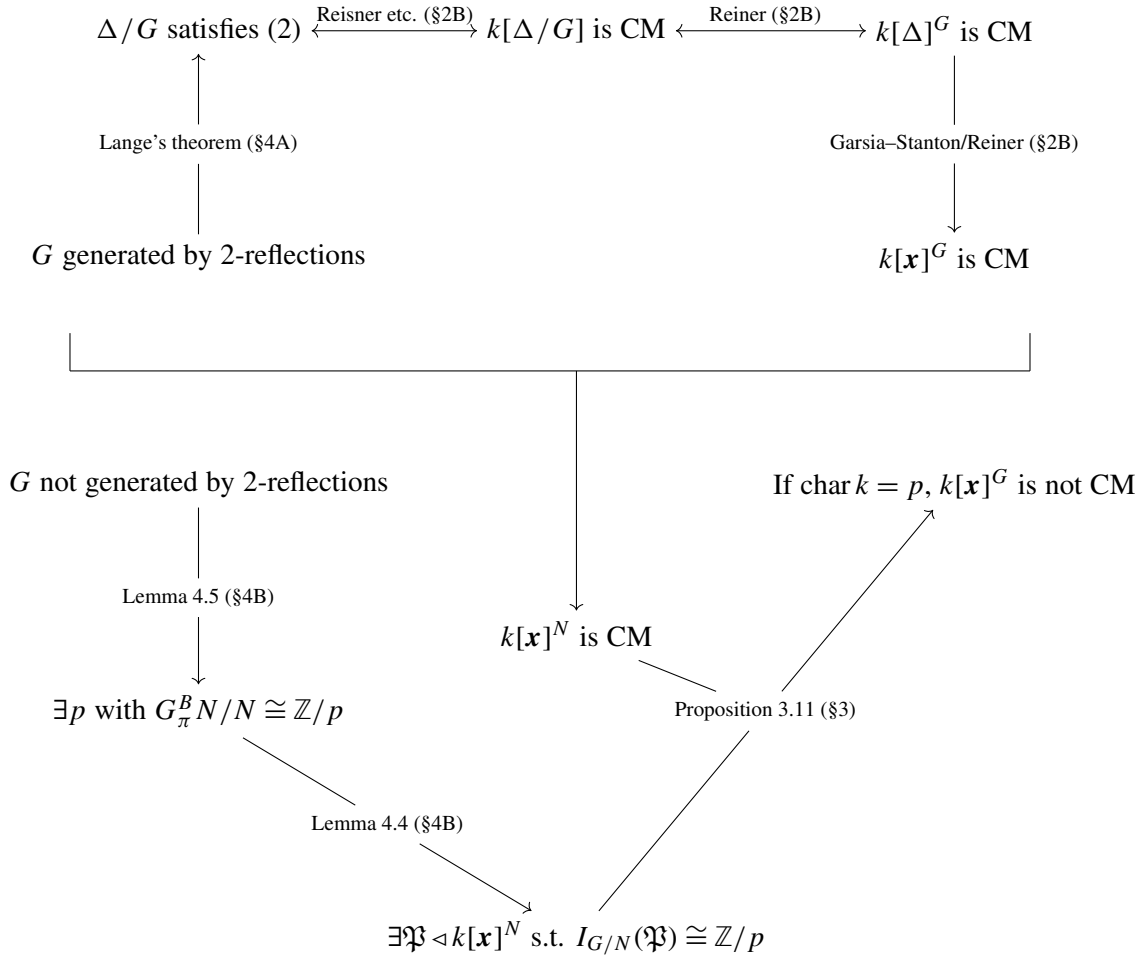
Figure 5. Schematic diagram of the proof of Theorem 1.2. Arrows are implications, and small print above or interrupting an arrow names a result needed for the implication to go through. The §-references indicate where to look for statements and notation definitions. The top half is the "if" direction (Proposition 4.1). The bottom half is the "only if" direction (Proposition 4.2). The group N is the subgroup of G generated by the 2-reflections, so the "if" direction is required to conclude that $k[\boldsymbol{x}]^N$ is Cohen–Macaulay in the bottom half.

in $\mathbb{R}^n$ by mapping each vertex, which by definition is an element $\alpha \in B_n \setminus \{\varnothing\}$, which is itself a nonempty subset of $[n]$, to the barycenter

$$\frac{1}{|\alpha|} \sum_{i \in \alpha} e_i$$

of the set of standard basis vectors $\{e_i\}_{i \in \alpha}$ corresponding to that subset. Then, extend this map to all of Δ by extending linearly from the vertices to each simplex in Δ. Finally, project the affine hyperplane

$\{\sum x_i = 1\}$ containing the image orthogonally onto T via $(x_1, \ldots, x_n) \mapsto (x_1 - 1/n, \ldots, x_n - 1/n)$. This embedding is G-equivariant for the action of G on $|\Delta|$ induced from its action on $[n]$, and the present action of G on T.

The embedded complex $|\Delta| \subset T$ is evidently a polyhedron, and it is a star of the origin in T since it is the union of closed line segments from the origin to its compact boundary, these segments are disjoint except for the origin itself, and it is a neighborhood of the origin in T (see the definition of a star in Section 2B). Since the action of G is linear, it permutes these segments. Thus $|\Delta|/G = |\Delta/G|$ is also a union of line segments from the (image of the) origin to its compact boundary, and these segments are disjoint except for the origin itself. Also, $|\Delta/G|$ is a neighborhood of the (image of the) origin since $T \to T/G$ is the quotient map by a group of homeomorphisms and is therefore an open map. It is additionally a polyhedron since the quotient map $T \to T/G$ is PL, and the image of a compact polyhedron under a PL map is a compact polyhedron [Rourke and Sanderson 1972, Corollary 2.5]. In other words, $|\Delta/G|$ is a polyhedral star of the image of the origin in the PL $(n-1)$-manifold T/G. It is therefore (per [Rourke and Sanderson 1972, pp.20–21], see the discussion at the end of Section 2B) homeomorphic to a ball. In particular, it is contractible, thus

$$\tilde{H}_i(|\Delta/G|; k) = 0$$

for all i, regardless of the field k; and it is a manifold (with boundary), thus

$$H_i(|\Delta/G|, |\Delta/G| - q; k) = 0$$

for all $i < n - 1$ and all $q \in |\Delta/G|$, regardless of k. Thus it satisfies (2) for all $i < \dim \Delta/G$ and all $q \in |\Delta/G|$, so by the discussion in Section 2B, $k[\boldsymbol{x}]^G$ is Cohen–Macaulay. □

4B. *The* only if *direction*. In this section we complete the proof of Theorem 1.2 by proving the converse of Proposition 4.1.

Proposition 4.2. *If G is not generated by transpositions, double transpositions, and 3-cycles, then there exists a prime p such that for any k of characteristic p, $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay.*

The proof is at the end of the section. Actually we prove somewhat more: for a group G not generated by transpositions, double transpositions, and 3-cycles, we give an explicit construction yielding the prime p. The precise statement is given below as Proposition 4.2b.

In this section, p is conceptually prior to the field k. Our proof will first construct p and then prove that when $\operatorname{char} k = p$, $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay.

We develop the needed machinery for the proof. Let Π_n be the poset of partitions of the set $[n]$, with the order relation given, for any $\pi, \tau \in \Pi_n$, by

$$\pi \leq \tau \iff \pi \text{ refines } \tau.$$

An element $g \in G \subset S_n$ partitions $[n]$ into orbits, and thus determines an element $\pi \in \Pi_n$. This gives a map

$$\varphi : G \to \Pi_n$$
$$g \mapsto \pi.$$

If $\pi \in \Pi_n$, we write G_π^B for the blockwise stabilizer of π in G, i.e., the set of elements of G that act separately on each block of π.

For a given $\pi \in \Pi_n$, let $\mathfrak{P}_\pi^\star$ be the prime ideal of $k[\boldsymbol{x}]$ generated by the binomials $x_i - x_j$ for every pair $i, j \in [n]$ lying in the same block of π. The dimension of $\mathfrak{P}_\pi^\star$ (i.e., the dimension of $k[\boldsymbol{x}]/\mathfrak{P}_\pi^\star$) is the number of blocks of π.

Lemma 4.3. *With this notation, we have*

$$I_G(\mathfrak{P}_\pi^\star) = G_\pi^B.$$

Proof. The ring $k[\boldsymbol{x}]/\mathfrak{P}_\pi^\star$ is the polynomial ring obtained by identifying x_i with x_j for each i and j in the same block of π, so its indeterminates are in bijection with the blocks of π. If $h \in G_\pi^B$, then h acts separately on the x_i's in each block, and therefore h fixes $\mathfrak{P}_\pi^\star$ setwise and the induced action on $k[\boldsymbol{x}]/\mathfrak{P}_\pi^\star$ is trivial. Thus $h \in I_G(\mathfrak{P}_\pi^\star)$. Conversely, if $h \notin G_\pi^B$, then either h fixes π but not blockwise, in which case h fixes $\mathfrak{P}_\pi^\star$ setwise but the action of h on $k[\boldsymbol{x}]/\mathfrak{P}_\pi^\star$ is not trivial, so that $h \in D_G(\mathfrak{P}_\pi^\star)$ but not $I_G(\mathfrak{P}_\pi^\star)$; or else h does not fix π at all, in which case it does not act on $\mathfrak{P}_\pi^\star$, and is not contained in $D_G(\mathfrak{P}_\pi^\star)$, let alone $I_G(\mathfrak{P}_\pi^\star)$. $\qquad\square$

If N is a normal subgroup of G, denote by $G_\pi^B N/N$ the image of G_π^B in the quotient G/N, and let

$$\mathfrak{P}_\pi = \mathfrak{P}_\pi^\star \cap k[\boldsymbol{x}]^N.$$

Lemma 4.4. *With this notation, we have*

$$I_{G/N}(\mathfrak{P}_\pi) = G_\pi^B N/N.$$

Proof. We have from Lemma 2.10 that

$$I_{G/N}(\mathfrak{P}_\pi) = I_G(\mathfrak{P}_\pi^\star)/I_N(\mathfrak{P}_\pi^\star) = I_G(\mathfrak{P}_\pi^\star)/(N \cap I_G(\mathfrak{P}_\pi^\star)),$$

and from Lemma 4.3 that

$$I_G(\mathfrak{P}_\pi^\star)/(N \cap I_G(\mathfrak{P}_\pi^\star)) = G_\pi^B/(N \cap G_\pi^B) = G_\pi^B N/N. \qquad\square$$

The following lemma is the device we use to find the characteristic p in which we can prove that $k[\boldsymbol{x}]^G$ fails to be Cohen–Macaulay.

Lemma 4.5. *Let $N \triangleleft G$ be a proper normal subgroup. Let π be minimal in Π_n among partitions associated (via φ) with elements of G that are not in N. Then:*

(1) *The group $G_\pi^B N/N$ is cyclic of prime order, say p.*

(2) *Any element g of $G \setminus N$ whose orbits are given by π has order a power of p.*

(3) *The image of g in G/N generates $G_\pi^B N/N$.*

Proof. Let g be an element of $G \setminus N$ whose orbits are given by π, and let h be any other nontrivial element of G_π^B, in other words a nontrivial element of G whose orbits refine π. (Note that, by minimality of π, either $\varphi(h) = \pi$ or else $h \in N$.) Pick any element $a \in [n]$ acted on nontrivially by h. Then g acts nontrivially on a as well since h's orbits refine g's.

Since h preserves π and g acts transitively on each block of π, there is an $m \in \mathbb{Z}$ such that $g^m(a) = h(a)$. Then $h^{-1}g^m(a) = a$, so that $h^{-1}g^m$ both preserves π and has a fixed point a that g does not have. Thus its orbits properly refine π, and minimality of π among partitions associated to elements of $G \setminus N$ implies that $h^{-1}g^m \in N$. Thus $hN = g^m N$. This shows that g generates the image of G_π^B in G/N, proving (3); thus $G_\pi^B N/N$ is cyclic. Meanwhile, for any prime p dividing the order of g, the orbits of g^p also properly refine those of g, so g^p is in N too; thus the image of g in G/N has order dividing p. Since $g \notin N$ by construction, the order of the image of g in G/N is exactly p. This completes the proof of (1). If q is a hypothetical second prime dividing the order of g in G, then the order of the image of g in G/N is q, for the same reason it is p, and it follows that $q = p$ after all, so there is no such second prime. Therefore g has p-power order in G. This proves (2). $\square$

Proof of Proposition 4.2. Let N be the subgroup of G generated by the transpositions, double transpositions, and 3-cycles (i.e., 2-reflections). By Proposition 4.1, $k[\boldsymbol{x}]^N$ is a Cohen–Macaulay ring. Since $k[\boldsymbol{x}]$ is a finitely generated algebra over k, $k[\boldsymbol{x}]^G$ is also finitely generated as an algebra over k [Bourbaki 1964, Chapitre V §1.9, Théorème 2], so in particular it is noetherian. By the same logic, $k[\boldsymbol{x}]^N$ is finitely generated as an algebra over k, and therefore over $k[\boldsymbol{x}]^G$. Since it is a subring of $k[\boldsymbol{x}]$, which is integral over $k[\boldsymbol{x}]^G$ by [Bourbaki 1964, Chapitre V §1.9, Proposition 22], it is integral over $k[\boldsymbol{x}]^G$ as well, which, together with finite generation as an algebra, implies it is actually finite over the noetherian ring $k[\boldsymbol{x}]^G$. Thus if k is a field of positive characteristic p, then Proposition 3.11 applies, and we can show $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay by exhibiting an inertia group for the action of G/N on $k[\boldsymbol{x}]^N$ that is a nontrivial p-group.

Now if N is a proper subgroup of G per the hypothesis, then we can find a $\pi \in \Pi_n$ that is minimal among all partitions associated (via φ) with elements of $G \setminus N$. Then Lemma 4.5 gives us a prime number p such that $G_\pi^B N/N$ is cyclic of order p, and then Lemma 4.4 gives us a prime ideal $\mathfrak{P}_\pi$ of $k[\boldsymbol{x}]^N$ such that

$$I_{G/N}(\mathfrak{P}_\pi) = G_\pi^B N/N.$$

Thus, for any k of this specific characteristic, we can conclude by Proposition 3.11 that $k[\boldsymbol{x}]^G$ fails to be Cohen–Macaulay. $\square$

An examination of the proof in view of conclusion (2) of Lemma 4.5 shows that we have actually proven the following constructive version of Proposition 4.2 with no additional work:

Proposition 4.2b. *Let N be the subgroup of G generated by the transpositions, double transpositions, and 3-cycles. If $N \subsetneq G$, then for any $g \in G \setminus N$ whose orbits are not refined by the orbits of any other*

$g \in G \setminus N$, *the order of g is a prime power p^ℓ, where p has the property that $k[\boldsymbol{x}]^G$ is not Cohen–Macaulay if* char $k = p$. $\square$

5. Conclusion and further questions

In this section we note some implications of the results above, and pose questions for further exploration. Throughout, let N be the subgroup of $G \subset S_n$ generated by the transpositions, double transpositions, and 3-cycles, as at the end of Section 4B.

5A. *Bad primes; relation to previous work.* Given a permutation group $G \subset S_n$, let us refer to the set of prime numbers p for which $k[\boldsymbol{x}]^G$ fails to be Cohen–Macaulay if char $k = p$ as the *bad primes* of G.

It was mentioned in the introduction that the "if" direction of Theorem 1.2 implies that the bad primes of G are a subset of the primes dividing $[G : N]$. We see this as follows: the "if" direction implies that $k[\boldsymbol{x}]^N$ is Cohen–Macaulay. Then, since

$$k[\boldsymbol{x}]^G = (k[\boldsymbol{x}]^N)^{G/N},$$

it follows from the Hochster–Eagon theorem [Hochster and Eagon 1971, Proposition 13] that $k[\boldsymbol{x}]^G$ is Cohen–Macaulay in any characteristic not dividing the order of G/N. Meanwhile, the "only if" direction of Theorem 1.2 implies that if the set of primes dividing $[G : N]$ is nonempty, then so is G's set of bad primes.

It was also mentioned in the introduction that the present work unites and generalizes several previously known results: Reiner's theorem [1992] that the invariant rings of A_n and the diagonally embedded $S_n \hookrightarrow S_n \times S_n$ are Cohen–Macaulay over all fields; Hersh's similar theorem [2003a; 2003b] for the wreath product $S_2 \wr S_n \subset S_{2n}$, and Kemper's theorems [1999] that in the p-group case, the "only if" direction of Theorem 1.2 holds, and that the invariant ring of a *regular* permutation group G is Cohen–Macaulay over all fields if and only if $G = C_2$, C_3, or $C_2 \times C_2$, and in all other cases, every prime dividing $|G|$ is a bad prime for G. Most of these results are immediate implications of the "if" direction of Theorem 1.2:

- The group A_n is generated by 3-cycles.
- The diagonal $S_n \hookrightarrow S_n \times S_n$ is generated by the double transpositions $(i, i+1)(i+n, i+n+1)$ for $i = 1, \ldots, n-1$.
- The wreath product $S_2 \wr S_n$ is generated by the transpositions $(2i-1, 2i)$ and the double transpositions $(2i-1, 2i+1)(2i, 2i+2)$ for $i = 1, \ldots, n-1$.
- The regular representations of C_2, C_3, and $C_2 \times C_2$ are generated by (in fact, their only nontrivial elements are) transpositions, 3-cycles, and double transpositions, respectively.

Recovering the other half of Kemper's result on regular permutation groups (that *every* prime dividing $|G|$ is bad for G) from the present work requires the constructive version of the "only if" direction given in Proposition 4.2b. Recall that if G acts regularly, i.e., freely and transitively, on $[n]$, then this action is isomorphic to G's left-translation action on its own elements. Then we have $|G| = n$, and every element

g of G splits $[n]$ into orbits of equal length the order of g, because these orbits are in bijection with the right cosets $\langle g \rangle h$, $h \in G$.

If G acts regularly and $|G| = n \geq 5$, then G does not contain any transpositions, double transpositions, or 3-cycles, so N is trivial. If p is any prime dividing $|G|$, then G has an element g of order p, which, by the discussion in the last paragraph, partitions $[n]$ into orbits of equal length p. This partition cannot be refined by any nontrivial partition with parts of equal length since p is prime; thus no element of $G \setminus N = G \setminus \{1\}$ can have orbits refining g's. It follows from Proposition 4.2b that p is a bad prime for G.

The remaining case is $n = 4$ and $G = C_4$. In this case, G is a 2-group not generated by its lone double transposition, so it follows from Theorem 1.2 that 2 is a bad prime for G.

5B. *Groups generated by transpositions, double transpositions, and 3-cycles.* Theorem 1.2 calls attention to the family of permutation groups generated by transpositions, double transpositions, and 3-cycles. One may wonder how extensive is this family of groups. It turns out to be very limited. One can extract a classification from Lange and Mikhaîlova's classification of all rotation-reflection groups [2016], but this is more power than is needed. In the case that G is transitive, such groups were already classified in 1979 by W. Cary Huffman [1980, Theorem 2.1]:

(1) If G's transpositions generate a transitive subgroup, then $G = S_n$.

(2) If G contains a transposition but the transpositions do not act transitively, then $n = 2m$ is even and G is isomorphic to the wreath product $S_2 \wr S_m$.

(3) If G does not contain a transposition but does contain a three-cycle, then $G = A_n$.

(4) Otherwise, G contains no transpositions or 3-cycles and is generated by double transpositions. Then we have:

    (a) If G contains a subgroup acting transitively on 5 points and fixing the rest, then either $n = 5$ and $G \cong D_5$ in its usual action on the vertices of a regular pentagon, or else $n = 6$ and $G \cong A_5 \cong \mathrm{PSL}(2, 5)$ in its transitive action on 6 points, e.g., the six points of the projective line over $\mathbb{F}_5$.

    (b) If G contains a subgroup acting transitively on 7 points and fixing the rest, then either $n = 7$ and $G \cong GL(3, 2)$ acting on the nonzero vectors of $\mathbb{F}_2^3$, or else $n = 8$ and $G \cong \mathrm{AGL}(3, 2) = \mathbb{F}_2^3 \rtimes \mathrm{GL}(3, 2)$ acting on the points of $\mathbb{A}_{\mathbb{F}_2}^3$.

    (c) If G does not contain either of these kinds of subgroups, then $n = 2m$ is even, and G is isomorphic to the alternating subgroup of the wreath product $S_2 \wr S_m$.

When one considers intransitive groups G, one does not end up too far beyond direct products of the above, since transpositions and 3-cycles can only act in a single orbit, while double transpositions can only act in two orbits, as a transposition in each. For example, if G has two orbits, the classification begins as follows. If G is not a direct product of the above, it contains a double transposition that acts as a transposition in each orbit. Then its image in each orbit contains a transposition, so is either S_n or $S_2 \wr S_m$ by the above. The possibilities are then highly constrained by Goursat's lemma.

Thus Theorem 1.2 shows that most permutation groups G have at least one bad prime.

5C. *Further questions.* Since Theorem 1.2 implies that the set of bad primes of G is contained in the set of prime factors of $[G : N]$ and is nonempty exactly when the latter is nonempty, one might hope that these two sets are always equal. This is not the case. For example, let $G \subset S_7$ be the Frobenius group of order 21 generated by

$$(1234567), \ (124)(365).$$

All the nontrivial elements in this group are 7-cycles or double 3-cycles. Thus N is trivial in this case, and the candidate bad primes are 3 and 7.

Now $\pi = \{1, 2, 4\} \cup \{3, 5, 6\} \cup \{7\}$ is a minimal partition as in Lemma 4.5, and thus the corresponding $g = (124)(365)$ generates an inertia group of order 3 for the action of $G/N = G$ on $k[\boldsymbol{x}]^N = k[\boldsymbol{x}]$. Then Proposition 4.2b shows that if k has characteristic 3, $k[\boldsymbol{x}]^G$ fails to be Cohen–Macaulay; i.e., 3 is a bad prime for this G.

On the other hand, 7 is not a bad prime for this G. This can be seen using the criterion given by Kemper [2001, Theorem 3.3], since 7 divides $|G|$ just once. Thus, a prime can divide $[G : N]$ without being bad. (By a computer calculation, no example of this phenomenon occurs below degree 7.)

At the other extreme, one might hope that the bad primes of G are *only* those which are furnished by Proposition 4.2b. This is not true either. Take $G = D_7$, the dihedral group of order 14 acting on the vertices of a heptagon, which is also a Frobenius group. Now, all the nontrivial elements are 7-cycles and triple transpositions, so again, N is trivial, and the candidate bad primes are 2 and 7. This time, they both really are bad primes. One can see this using Kemper's criterion [2001, Theorem 3.3]. For 2 it also follows from Proposition 4.2b, but for 7 it does not, since the 7-cycles have orbits that are properly refined by the triple transpositions.

Thus it remains to be determined, for a given G, exactly which primes are bad. Theorem 1.2 gives us a finite list of candidate bad primes (those dividing $[G : N]$), and, if this list is nonempty, Proposition 4.2b gives us some specific primes that are definitely bad. Among the remaining candidate bad primes, if any divide $|G|$ only once, [Kemper 2001, Theorem 3.3] can be used to determine if they are actually bad. What remains to be determined is whether p is a bad prime if $p^2 \mid |G|$ and p is not associated to a $g \in G \setminus N$ with minimal orbits as in Proposition 4.2b.

Question 5.1. How can Cohen–Macaulayness of $k[\boldsymbol{x}]^G$ be assessed when [Kemper 2001, Theorem 3.3] and the present work are both inapplicable, i.e., when $p \mid [G : N]$ and $p^2 \mid |G|$, but p does not come from a minimal $g \in G \setminus N$ as in Proposition 4.2b?

Another line of inquiry that flows from the present work has to do with the relationship between the arguments in the "if" and "only if" directions. The proof of the "if" direction is a mildly revised version of an argument given by the first author in his doctoral thesis [Blum-Smith 2017]. In that same work, he also proved the "only if" direction for $k[\Delta/G]$ (see Section 2B for notation), but not for $k[\boldsymbol{x}]^G$. There, the "only if" argument was framed in the same topological language as the "if" argument, which is why

it applied to $k[\Delta/G]$ (taking advantage of Stanley–Reisner theory) but not $k[\boldsymbol{x}]^G$. The second author suggested to transfer the "only if" argument from topological into commutative-algebraic language, and much of the present paper sprang from this suggestion.

This transfer was accomplished piecemeal, with an individual search for each commutative-algebraic fact needed to replace each topological fact. For example, Raynaud's theorem (Lemma 2.14) replaced an elementary principle about the relationship between point stabilizers and the local structure in a topological quotient. The well-behavedness of inertia groups with respect to normal subgroups (Lemma 2.10) replaced an elementary fact about group actions on a set. The observation that inertia p-groups obstruct Cohen–Macaulayness if they are not generated by 2-reflections (Lemma 3.7), based on [Lorenz and Pathak 2001, Corollary 4.3], replaced an argument about the homology of links in the quotient of a simplicial complex.

Nonetheless, the authors had the conviction throughout that an overarching principle was at play. It may be fruitful to seek a more comprehensive understanding of the relationship between the topology and the algebra. Stanley–Reisner theory gives a partial answer to this question, but it does not appear to account for the "only if" direction of Theorem 1.2, so a fuller picture is desirable.

Here are two more focused questions that approach this inquiry from various directions:

Question 5.2. Is there a purely algebraic proof of Theorem 1.2, making no use of Stanley–Reisner theory or Lange's result on PL manifolds?

Question 5.3. For a fixed $p = \operatorname{char} k$ as in Question 5.1, can $k[\boldsymbol{x}]^G$ be Cohen–Macaulay without $k[\Delta/G]$ being Cohen–Macaulay?

Acknowledgement

References

[Blum-Smith 2017] B. Blum-Smith, *Two inquiries about finite groups and well-behaved quotients*, Ph.D. thesis, New York University, 2017, Available at https://search.proquest.com/docview/1938647220. MR

[Bourbaki 1964] N. Bourbaki, *Éléments de mathématique: Algèbre commutative, Chapitres 5–6*, Actualités Scientifiques et Industrielles **1308**, Hermann, Paris, 1964. MR Zbl

[Bruns and Herzog 1993] W. Bruns and J. Herzog, *Cohen–Macaulay rings*, Cambridge Studies in Advanced Mathematics **39**, Cambridge University Press, 1993. MR Zbl

[Campbell et al. 1999] H. E. A. Campbell, A. V. Geramita, I. P. Hughes, R. J. Shank, and D. L. Wehlau, "Non-Cohen–Macaulay vector invariants and a Noether bound for a Gorenstein ring of invariants", *Canad. Math. Bull.* **42**:2 (1999), 155–161. MR Zbl

[Dufresne et al. 2009] E. Dufresne, J. Elmer, and M. Kohls, "The Cohen–Macaulay property of separating invariants of finite groups", *Transform. Groups* **14**:4 (2009), 771–785. MR Zbl

[Duval 1997] A. M. Duval, "Free resolutions of simplicial posets", *J. Algebra* **188**:1 (1997), 363–399. MR Zbl

[EGA I 1960] A. Grothendieck, "Eléments de géométrie algébrique, I: Le langage des schémas", *Inst. Hautes Études Sci. Publ. Math.* **4** (1960), 5–228. MR Zbl

[Eisenbud 1995] D. Eisenbud, *Commutative algebra with a view toward algebraic geometry*, Graduate Texts in Mathematics **150**, Springer, 1995. MR Zbl

[Ellingsrud and Skjelbred 1980] G. Ellingsrud and T. Skjelbred, "Profondeur d'anneaux d'invariants en caractéristique p", *Compositio Math.* **41**:2 (1980), 233–244. MR Zbl

[Garsia 1980] A. M. Garsia, "Combinatorial methods in the theory of Cohen–Macaulay rings", *Adv. in Math.* **38**:3 (1980), 229–266. MR Zbl

[Garsia and Stanton 1984] A. M. Garsia and D. Stanton, "Group actions of Stanley–Reisner rings and invariants of permutation groups", *Adv. in Math.* **51**:2 (1984), 107–201. MR Zbl

[Gordeev and Kemper 2003] N. Gordeev and G. Kemper, "On the branch locus of quotients by finite groups and the depth of the algebra of invariants", *J. Algebra* **268**:1 (2003), 22–38. MR Zbl

[Hamilton and Marley 2007] T. D. Hamilton and T. Marley, "Non-Noetherian Cohen–Macaulay rings", *J. Algebra* **307**:1 (2007), 343–360. MR Zbl

[Hersh 2003a] P. Hersh, "Lexicographic shellability for balanced complexes", *J. Algebraic Combin.* **17**:3 (2003), 225–254. MR Zbl

[Hersh 2003b] P. Hersh, "A partitioning and related properties for the quotient complex $\Delta(B_{lm})/S_l \wr S_m$", *J. Pure Appl. Algebra* **178**:3 (2003), 255–272. MR Zbl

[Hochster and Eagon 1971] M. Hochster and J. A. Eagon, "Cohen–Macaulay rings, invariant theory, and the generic perfection of determinantal loci", *Amer. J. Math.* **93** (1971), 1020–1058. MR Zbl

[Huffman 1980] W. C. Huffman, "Imprimitive linear groups generated by elements containing an eigenspace of codimension two", *J. Algebra* **63**:2 (1980), 499–513. MR Zbl

[Kac and Watanabe 1982] V. Kac and K. Watanabe, "Finite linear groups whose ring of invariants is a complete intersection", *Bull. Amer. Math. Soc.* (*N.S.*) **6**:2 (1982), 221–223. MR Zbl

[Kemper 1999] G. Kemper, "On the Cohen–Macaulay property of modular invariant rings", *J. Algebra* **215**:1 (1999), 330–351. MR Zbl

[Kemper 2001] G. Kemper, "The depth of invariant rings and cohomology", *J. Algebra* **245**:2 (2001), 463–531. MR Zbl

[Kemper 2012] G. Kemper, "The Cohen–Macaulay property and depth in invariant theory", pp. 53–63 in *Proceedings of the 33rd Symposium on Commutative Algebra* (Japan), 2012.

[Lange 2016] C. Lange, "Characterization of finite groups generated by reflections and rotations", *J. Topol.* **9**:4 (2016), 1109–1129. MR Zbl

[Lange and Mikhaîlova 2016] C. Lange and M. A. Mikhaîlova, "Classification of finite groups generated by reflections and rotations", *Transform. Groups* **21**:4 (2016), 1155–1201. MR Zbl

[Liu 2002] Q. Liu, *Algebraic geometry and arithmetic curves*, Oxford Graduate Texts in Mathematics **6**, Oxford University Press, 2002. MR Zbl

[Lorenz and Pathak 2001] M. Lorenz and J. Pathak, "On Cohen–Macaulay rings of invariants", *J. Algebra* **245**:1 (2001), 247–264. MR Zbl

[Munkres 1984] J. R. Munkres, "Topological results in combinatorics", *Michigan Math. J.* **31**:1 (1984), 113–128. MR Zbl

[Neusel and Smith 2002] M. D. Neusel and L. Smith, *Invariant theory of finite groups*, Mathematical Surveys and Monographs **94**, American Mathematical Society, Providence, RI, 2002. MR Zbl

[Raynaud 1970] M. Raynaud, *Anneaux locaux henséliens*, Lecture Notes in Math. **169**, Springer, 1970. MR Zbl

[Reiner 1992] V. Reiner, "Quotients of Coxeter complexes and P-partitions", pp. vi+134 Mem. Amer. Math. Soc. **460**, Amer. Math. Soc., Providence, RI, 1992. MR Zbl

[Reiner 2003] V. Reiner, Appendix to [Hersh 2003b], pages 269–271 in *J. Pure Appl. Algebra* **178**:3 (2003), 255–272.

[Reisner 1976] G. A. Reisner, "Cohen–Macaulay quotients of polynomial rings", *Advances in Math.* **21**:1 (1976), 30–49. MR Zbl

[Rourke and Sanderson 1972] C. P. Rourke and B. J. Sanderson, *Introduction to piecewise-linear topology*, Ergebnisse der Mathematik **69**, Springer, Berlin, 1972. Reprinted in the series Springer Study Edition **69**, Springer, Berlin, 1982. MR Zbl

[Smith 1995] L. Smith, *Polynomial invariants of finite groups*, Research Notes in Mathematics **6**, A K Peters, Ltd., Wellesley, MA, 1995. MR Zbl

[Smith 1996] L. Smith, "Some rings of invariants that are Cohen–Macaulay", *Canad. Math. Bull.* **39**:2 (1996), 238–240. MR Zbl

[Stanley 1986] R. P. Stanley, *Enumerative combinatorics, I*, Wadsworth & Brooks/Cole, Monterey, CA, 1986. MR Zbl

[Stanley 1991] R. P. Stanley, "f-vectors and h-vectors of simplicial posets", *J. Pure Appl. Algebra* **71**:2-3 (1991), 319–331. MR Zbl

ben@cims.nyu.edu   *Department of Natural Sciences and Mathematics, Eugene Lang College,*
                                    the New School for Liberal Arts, New York City, NY, United States

sophie.marques@uct.ac.za   *Department of Mathematics and Applied Mathematics,*
                                    University of Cape Town, Cape Town, South Africa

Algebra & Number Theory

msp.org/ant

See inside back cover or msp.org/ant for submission instructions.

Algebra & Number Theory

Volume 12   No. 7   2018