ALGORITHMS FOR GROUP ISOMORPHISM VIA GROUP EXTENSIONS AND COHOMOLOGY[*]

JOSHUA A. GROCHOW[†] AND YOUMING QIAO[‡]

Abstract. The isomorphism problem for finite groups of order n (GPI) has long been known to be solvable in $n^{\log n + O(1)}$ time, but only recently were polynomial-time algorithms designed for several interesting group classes. Inspired by recent progress, we revisit the strategy for GPI via the extension theory of groups. The extension theory describes how a normal subgroup N is related to G/N via G, and this naturally leads to a divide-and-conquer strategy that "splits" GPI into two subproblems: one regarding group actions on other groups, and one regarding group cohomology. When the normal subgroup N is abelian, this strategy is well known. Our first contribution is to extend this strategy to handle the case when N is not necessarily abelian. This allows us to provide a unified explanation of all recent polynomial-time algorithms for special group classes. Guided by this strategy, to make further progress on GPI, we consider *central-radical* groups, proposed in Babai et al. [*Code equivalence and group isomorphism*, in Proceedings of the 22nd Annual ACM–SIAM Symposium on Discrete Algorithms (SODA'11), SIAM, Philadelphia, 2011, ACM, New York, pp. 1395–1408]: the class of groups such that G modulo its center has no abelian normal subgroups. This class is a natural extension of the group class considered by Babai et al. [*Polynomial-time isomorphism test for groups with no abelian normal subgroups (extended abstract)*, in International Colloquium on Automata, Languages, and Programming (ICALP), 2012, pp. 51–62], namely those groups with no abelian normal subgroups. Following the above strategy, we solve GPI in $n^{O(\log \log n)}$ time for central-radical groups, and in polynomial time for several prominent subclasses of central-radical groups. We also solve GPI in $n^{O(\log \log n)}$ time for groups whose solvable normal subgroups are elementary abelian but not necessarily central. As far as we are aware, this is the first time there have been worst-case guarantees on an $n^{o(\log n)}$-time algorithm that tackles both aspects of GPI—actions and cohomology—simultaneously. Prior to this work, the best proven upper bounds on algorithms for groups with central radicals were $n^{O(\log n)}$, even for groups with a central radical of constant size, such as $\mathrm{Rad}(G) = Z(G) = \mathbb{Z}_2$. To develop our new algorithms we utilize several mathematical results on the detailed structure of cohomology classes, as well as algorithmic results for code equivalence, coset intersection, and cyclicity testing of modules over finite-dimensional associative algebras. We also suggest several promising directions for future work.

Key words. group isomorphism, algorithm, group cohomology, polynomial-time

AMS subject classifications. 20C40, 68Q25

DOI. 10.1137/15M1009767

1. Introduction. [§]   The group isomorphism problem (GPI) is to determine whether two finite groups are isomorphic. For groups of order n, the easy $n^{\log n + O(1)}$-time algorithm [34, 64][1] for the general case of GPI has barely seen any asymptotic improvement over the past four decades; it was improved recently to $n^{1/4 \log n + O(1)}$ by

---

[†]Department of Computer Science, University of Colorado, Boulder, CO 80309, and the Santa Fe Institute, Santa Fe, NM 87501 (joshua.grochow@cs.colorado.edu).

[‡]Centre for Quantum Software and Information, University of Technology Sydney, New South Wales 2007, Australia (youming.qiao@uts.edu.au).

[§]The introduction may serve as an extended abstract: Subsection 1.3 contains an informal exposition of sections 2 to 4; subsection 1.4 gives a brief overview of sections 5 to 7. An extended abstract based on the introduction appeared as [41].

[1]Miller [64] attributes this algorithm to Tarjan.

Rosenbaum [75] (see [59, sect. 2.2]), but even the extensive body of work on practical algorithms led by Eick, Holt, Leedham-Green, and O'Brien (e. g., [19, 31, 18, 26])—resulting in most of the functional algorithms in use today—was recently found [88] to only improve the constant in the exponent, still resulting in an $n^{\Theta(\log n)}$-time algorithm for the general case. The past few years have witnessed a resurgence of activity on worst-case guaranteed algorithms for this problem [58, 11, 70, 86, 60, 16, 12, 76, 75, 23, 42].

Before introducing these works and our results, we recall why group isomorphism is an intriguing problem from the complexity-theoretic perspective, even when the groups are given by their multiplication tables. (See Remark 1.1 below.) We call this version of the problem CayleyGpI. As CayleyGpI reduces to Graph Isomorphism (GraphI) (see, e. g., the book [56]), CayleyGpI currently has an intermediate status: It is not NP-complete unless PH collapses [20, 15], and is not known to be in P. In addition to its intrinsic interest, resolving the exact complexity of GpI is a tantalizing question. Further, there is a surprising connection between CayleyGpI and the Geometric Complexity Theory program (see, e. g., [65] and references therein): Techniques from CayleyGpI were used to solve cases of Lie Algebra Isomorphism that have applications in Geometric Complexity Theory [40].

In a survey article [5] in 1995, after enumerating several isomorphism-type problems, including GraphI and GpI, Babai expressed the belief that CayleyGpI might be the only one expected to be in P. Indeed, in many ways CayleyGpI seems easier than GraphI: There is a simple $n^{\log n + O(1)}$-time algorithm for GpI, whereas the best known algorithm for GraphI takes time $2^{(\log n)^c}$ for some $c \geq 3$ [8] and is quite complicated.[2] There is a polynomial-time reduction from CayleyGpI to GraphI, yet there is provably no AC^0 reduction in the opposite direction [27]. The reduction CayleyGpI $\leq$ GraphI means that CayleyGpI stands as an obstacle to putting GraphI into P; in light of the recent quasi-polynomial-time algorithm for GraphI [8], this obstacle has become much more salient (the previous best algorithm for GraphI was so far from quasi-polynomial that there were clearly obstacles to be overcome before CayleyGpI became a serious road block, but many of those obstacles have now been overcome). Further, GraphI is as hard as its counting version, whereas no such counting-to-decision reduction is known for GpI. Finally, whereas the smallest standard complexity class known to contain GraphI is NP∩coAM, Arvind and Torán [1] showed that CayleyGpI for solvable groups is in NP ∩ coNP under a plausible assumption, weaker than that needed to show GraphI ∈ coNP (recall that a group is solvable if all its composition factors are abelian, or equivalently if the derived series $G_0 = G$, $G_{i+1} = [G_i, G_i]$ terminates in the identity).

Despite this situation and considerable attention to GpI, prior to 2009 the actual developments toward algorithms with worst-case guaranteed running time polynomial in $|G|$ essentially stopped at abelian groups, although there have been impressive practical advances (see, e. g., the theses [80, 49] and references therein for nice overviews). Isomorphism of abelian groups has long been known to be solvable in polynomial time [78, 50, 85, 54]. The next natural group class after abelian groups—class 2 nilpotent groups, whose quotient by their center is abelian—turns out to be formidable (see [38, 60, 23] for efficient algorithms in some restricted cases).

Beginning in 2009 there were several advances on worst-case guaranteed algorithms, starting with Le Gall [58]. In [12], following [11], Babai et al. developed a

---

[2]At the time of writing, the paper [8] is still under peer review. The previous-best algorithm took time $2^{O(\sqrt{n \log n})}$ (see [14]) and was also quite complicated.

polynomial-time algorithm for groups with no abelian normal subgroups. This suggests the presence of abelian normal subgroups as a bottleneck. With this in mind, Babai and Qiao [16] developed a polynomial-time algorithm for a special class of nonnilpotent solvable groups, building on [58, 70]; this was recently extended by the present authors to the so-called groups of tame extensions [42]. In [60], Lewis and Wilson made intriguing progress on p-groups: They gave a polynomial-time algorithm for quotients of generalized Heisenberg groups, a decently large subclass of p-groups of class 2. Rosenbaum's recent works [76, 75] (some of them building on ideas of Wagner [86]) led to an $n^{1/4 \log n + O(1)}$-time algorithm for GPI (see [59, sect. 2.2]). To summarize, at present it is crucial to understand indecomposable groups with abelian normal subgroups to develop $n^{o(\log n)}$-time algorithms.

Given these developments, we are at an interesting crossroads: First, as several nontrivial algorithms with worst-case guarantees have recently been developed (or worst-case guarantees proven for previously known algorithms), it is reasonable to look back to see if there is some common pattern or structure to these results. Second, of course, we should continue to improve the state of the art by developing more $n^{o(\log n)}$-time algorithms for special group classes. Finally, class 2 nilpotent groups seem to remain the bottleneck, but despite heuristic evidence, it is still desirable to formalize a reduction from the general case to this seeming bottleneck.

In this paper we contribute to all three of the preceding aspects: (1) we show how a general strategy for group isomorphism from the mathematics literature can be used to bound the worst-case complexity; and (2) using that strategy, we develop an $n^{O(\log \log n)}$-time algorithm for a group class proposed in [11], and polynomial-time algorithms for some prominent subclasses thereof. The worst-case analysis of this strategy also helps to explain in a unified way the recent successes on other group classes [58, 70, 16, 11, 12, 60, 42], which can be viewed as adding class-specific tactics to the general strategy. We also explain how these results may help to reduce general GPI to the class 2 nilpotent case.

Remark 1.1 (efficient implementation versus computational complexity). There are naturally two audiences for this paper, who might view it quite differently: (A) computational complexity theorists / algorithms theory researchers, and (B) computational group theorists. To some in the latter group, much of the content of this paper is surely well known, and to them perhaps not even worth writing down at this point in history. For them, we would like to highlight our results on nonabelian cohomology (section 3) and the use of Guralnick–Kantor–Kassabov–Lubotzky [43], which we believe may be new even in light of the large body of work in that community. We would also like to point out that the remainder of the paper will likely seem new to computational complexity theorists, despite seeming trivial to computational group theorists. As so often happens when ideas from one area B (in this case, classical and computational group theory) are imported to another area A (computational complexity), area B may view the results as trivial, while area A may view them as a nontrivial advance. However, while much of this paper is targeted at computational complexity theorists and so has this flavor, we believe that some of the results (as mentioned above) will be new to both communities.

As discussed above, there are good reasons (e. g., related to GRAPHI) to be interested in getting the asymptotic complexity of GPI down to polynomial in the order of the group, independently of its (ir)relevance to practical implementation. Combined with the fact that all of the practical algorithms put together still leave the worst-case bound at $|G|^{\Theta(\log |G|)}$ [88], we thus state our runtime bounds as a function of $n = |G|$.

Nonetheless, our algorithms are fairly agnostic to the method of input and in particular *do not depend on the input being a Cayley table.* That is, when the runtime of an algorithm is at least quadratic in the order of the group, even if the group is input as a black box, the algorithm may first enumerate the entire multiplication table and then proceed from there.

Except in a few cases, we do not bother to estimate multiplicative constants, or even constant exponents, in most of our algorithms. The current inability to get a general algorithm with runtime $|G|^{o(\log |G|)}$ [88] suggests that this is a potentially deep research topic, and getting algorithms with guaranteed runtime $|G|^{O(1)}$ can be seen as a first step in understanding the difficulties involved (not to mention its complexity-theoretic interest). Furthermore, there is nontrivial precedent for synergistic interaction between worst-case analyses and practical algorithms, as reflected in, e.g., Seress's pioneering implementations in GAP [37] of theoretically proven fast algorithms for permutation groups (see, e.g., [79]). Therefore, although we do not deal with issues of practical algorithms in this paper, it is our sincere hope that this great tradition will continue in future work.

1.1. Main results. The classes of groups we consider are natural extensions of the class of groups considered in [11, 12] and are additionally motivated by the Babai–Beals filtration [9] (resurrecting certain ideas that go back to Fitting [35]) and the Cannon–Holt approach to group isomorphism in the practical setting [26]. We go into the details of the Babai–Beals filtration and the Cannon–Holt approach in section 8. Here we merely give enough of a flavor to help motivate the classes of groups we consider.

Important in both the Babai–Beals filtration and the Cannon–Holt approach is the solvable radical. The solvable radical $\mathrm{Rad}(G)$ of a group G is the unique maximum solvable normal subgroup of G. Note that the center $Z(G)$, as an abelian normal subgroup, is contained in $\mathrm{Rad}(G)$. $G/\mathrm{Rad}(G)$ contains no solvable normal subgroups, side-stepping the currently intractable obstacle of solvable groups. Babai et al. [12] give a polynomial-time algorithm for isomorphism of groups with no solvable normal subgroups (equivalently, no abelian normal subgroups); following them, we call such groups "semisimple."

We mainly consider the class of groups whose solvable radical coincides with its center, that is, $\mathrm{Rad}(G) = Z(G)$; in subsection 6.2 we also consider groups whose solvable radical is abelian but need not be contained in the center. The former class, which we refer to as *groups with central radicals* or *central-radical groups* for short, is a natural extension of the class of semisimple groups and a natural stepping stone toward general groups (see Figure 1 below). Note that for such groups the solvable radical is necessarily abelian. Besides the motivations mentioned above, central-radical groups also cover a class of groups that is well studied in finite group theory (see Appendix B). In the theory of Lie groups, central-radical groups correspond to the well-studied and important class of reductive Lie groups, which are important throughout mathematics and physics, often because of their nice representation-theoretic properties.

We use the strategy outlined in subsection 1.3 below to achieve the following results. For groups with central radicals, we give an $n^{O(\log \log n)}$-time algorithm in general, and for several subclasses of groups with central radicals we give polynomial-time algorithms. We also give similarly efficient algorithms for groups with elementary abelian, but not necessarily central, radicals. Prior to this work, the best proven upper bounds on algorithms for groups with central radicals were $n^{O(\log n)}$, even for groups with a central radical of constant size, such as $\mathrm{Rad}(G) = Z(G) = \mathbb{Z}_2$.

Recall that for any groups G and H, the set of isomorphisms between them is either empty or a coset of the automorphism group $\mathrm{Aut}(G)$ in the group of permutations of the disjoint union $G \sqcup H$. We say that the coset of isomorphisms can be found if one isomorphism $G \stackrel{\cong}{\to} H$ and a generating set for $\mathrm{Aut}(G)$ can be found. Finding the full coset of isomorphisms—rather than just deciding GPI or finding a single isomorphism—is often useful in recursively building algorithms for larger group classes from those for smaller classes.

THEOREM A(=COROLLARY 6.2). *Isomorphism of central-radical groups of order* n *can be decided in time* $n^{c \log \log n + O(1)}$ *for* $c = 1/\log_2(60) \approx 0.169$. *Furthermore, if the radical is elementary abelian, the coset of isomorphisms can be found in the same time bound.*

The algorithm in the above theorem in fact runs in polynomial time when the order or structure of the semisimple quotient $G/\mathrm{Rad}(G)$ is bounded as follows. Recall that a normal subgroup of G is *minimal* if it is nontrivial and does not contain any smaller normal subgroups of G. The number of minimal normal subgroups of $G/\mathrm{Rad}(G)$ is always at most $\log_{60} n$; if it happens to be just slightly smaller, then we have the following corollary.

COROLLARY 6.3. *Let G and H be central-radical groups of order n. If $G/\mathrm{Rad}(G)$ has $O(\frac{\log n}{\log \log n})$ minimal normal subgroups, isomorphism between G and H can be decided in* poly(n) *time. Furthermore, if the radical is elementary abelian, the coset of isomorphisms can be found in the same time bound.*

In particular, this includes groups G satisfying $|G/\mathrm{Rad}(G)| \leq n^{O(1/\log \log n)}$, but also many groups where $G/\mathrm{Rad}(G)$ is much larger. Both of these theorems are in fact corollaries of our more general Theorem 6.1 together with previous results on semisimple groups [11], but we defer the statement of Theorem 6.1 until section 6, as the above results make its significance clearer.

For groups with elementary abelian radicals—even if they are not central—we get the same conclusions. This requires us to simultaneously solve ACTION COMPATIBILITY and COHOMOLOGY CLASS ISOMORPHISM. We combine the above techniques with a novel reduction to known representation-theoretic algorithms [28] to get the following theorem.

THEOREM B(=COROLLARIES 6.12 AND 6.13). *Isomorphism of groups of order n with elementary abelian radicals can be decided, and the coset of isomorphisms found, in time* $n^{c \log \log n + O(1)}$ *for* $c = 1/\log_2(60) \approx 0.169$.

If, furthermore, $G/\mathrm{Rad}(G)$ has $O(\frac{\log n}{\log \log n})$ normal subgroups, isomorphism can be decided, and the coset of isomorphisms found, in poly(n) *time.*

We then consider central-radical groups with $G/\mathrm{Rad}(G)$ a direct product of nonabelian simple groups. Although this may seem restrictive, this class of groups is quite natural. In group theory, this class is closely related to the generalized Fitting subgroups (see, e.g., [83, Chap. 6, section 6], [2, Chap. 11], and Appendix B). Also, within central-radical groups, this class has two characterizations: (1) the class of groups in which the last two of the four levels of the Babai–Beals filtration are trivial (see subsection 8.2); or (2) the class of groups that are equal to their generalized Fitting subgroup (see Appendix B). Our next result gives polynomial-time algorithms for this group class, which includes, for example, central extensions of $\mathbb{Z}_p^{\Theta(\log n)}$ by $A_k^{\Theta(\log n)}$, which do not satisfy the conditions of the results above.

THEOREM C(=THEOREM 7.1). *Isomorphism between two groups G_1, G_2 with central, elementary abelian radicals can be decided, and the coset of isomorphisms found, in* poly($|G_i|$) *time if either*

1. $G_1/\operatorname{Rad}(G_1)$ *is a direct product of simple groups; or*
2. $G_1/\operatorname{Rad}(G_1)$ *is a direct product of perfect groups, each of order $O(1)$.*

More importantly, we believe the techniques that go into proving this theorem are worth noting: we rely on a detailed analysis of the structure of the cohomology classes specific to this group class (see subsection 7.1, and the use of the powerful results of [43]) to allow for the application of known algorithmic techniques, including singly exponential-time algorithms for LINEAR CODE EQUIVALENCE [7] (see [11, Thm. 7.1]) and COSET INTERSECTION [4, 63] (see also [6, 13]).

1.2. Motivation for the classes of groups considered. Aside from the motivations already mentioned above, Figure 1 gives the general idea of where this paper fits in the picture of a larger approach toward putting GPI into P. The figure is neither complete nor 100% accurate in terms of the landscape of groups and algorithms for GPI but is more or less correct for algorithms with worst-case guarantees at a large scale.
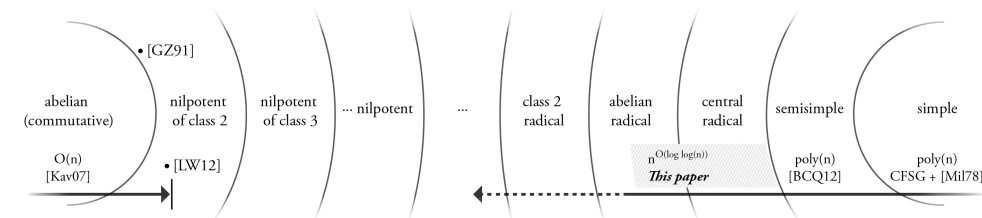

FIG. 1. *Some progress on group isomorphism. The arrow on the left indicates that the techniques for abelian groups hit a wall before class 2 nilpotent groups. The arrow on the right indicates the recent progress at the "opposite end of the spectrum," including this paper and prospects for future progress.*

While central-radical groups may seem to be only a slight extension of semisimple groups, they in fact differ significantly from previous classes of groups with $n^{o(\log n)}$-time isomorphism algorithms. In particular, most previous $n^{o(\log n)}$-time algorithms for GPI of special group classes only consider one of the two main aspects of GPI, namely actions (in subsection 4.2, we briefly indicate how actions are used in [58, 70, 16, 11, 12]). On the other hand, to work with groups with central radicals, we need to focus on the other main aspect of the problem, namely cohomology (see subsection 1.3).

Our results also suggest one more step toward a formal reduction from the general case to nilpotent groups of class 2. In particular, in Proposition 3.13 and Remark 3.14, we show that for groups G where the action of G on $\operatorname{Rad}(G)$ by conjugation is essentially trivial (technically, the action is by inner automorphisms of $\operatorname{Rad}(G)$) and $\operatorname{Aut}(Z(\operatorname{Rad}(G)))$ is small, GPI reduces to isomorphism of central-radical groups and isomorphism of solvable groups, separately. Thus, if isomorphism of central-radical groups could be decided in polynomial time, then GPI for this class of groups would reduce to GPI for solvable groups. Note that central-radical groups arise here naturally, by considering cases in which the relationship between $G/\operatorname{Rad}(G)$ and $\operatorname{Rad}(G)$ is simple. This is just one example of how we believe our ideas point toward a possible

reduction from the general case to nilpotent groups; some other ideas in this paper may also be useful in this regard (see subsection 8.3).

1.3. A strategy via group extensions and cohomology. In this paper we use the theory of group extensions (which we describe briefly in this section and in more detail in section 2, and give expository preliminaries in Appendix A; for textbook treatments, see, e. g., [73, Chap. 11] and [77, Chap. 7]) to show that the group isomorphism problem "splits" into two subproblems—one coming from actions of groups on other groups (Action Compatibility), and the other coming from group extensions and cohomology (Cohomology Class Isomorphism), which we explain below. We note that Besche and Eick have proposed this splitting in a slightly different setting, under the name "strong isomorphism" [18]. In the abstract theory of finite groups this splitting is standard material; the contribution here is the observation that this standard material can be used to prove worst-case algorithmic guarantees and that doing so is useful and even formally necessary to resolve the complexity of GpI. We also extend this approach to the setting where the normal subgroup can be nonabelian. For the converse direction, we observe that special cases of these subproblems reduce to GpI under polynomial-time reductions (subsection 4.3). We summarize these results in the following.

"Splitting" GpI into actions and cohomology.
- *Fact 4.1: For coprime extensions* Action Compatibility $\equiv_m^p$ GpI.
- *Fact 4.3: For p-groups of class 2, when $p > 2$,* Cohomology Class Isomorphism $\equiv_m^p$ GpI.
- *Lemmas 2.3 and 3.7: GpI for groups with a normal subgroup from one class of groups $\mathcal{N}$ and a quotient from a class of groups $\mathcal{Q}$ reduces to solving GpI in $\mathcal{N}$, solving it in $\mathcal{Q}$, and simultaneously solving* Action Compatibility *and* Cohomology Class Isomorphism.

A "simultaneous solution" to these two problems is possible because they have the same space of potential solutions, namely certain automorphisms of certain groups; a simultaneous solution is an automorphism that is simultaneously a solution for both problems. See subsection 1.4 and Lemma 2.3 for more details.

Most previous complexity-theoretic results on GpI have focused on some combination of algorithmic techniques and Action Compatibility. In this paper, for the first time from the worst-case complexity perspective, we make progress on Cohomology Class Isomorphism.

We now explain this "splitting" and the problems mentioned above informally. Consider the following natural strategy for testing whether G is isomorphic to H. If G is simple, then isomorphism can be tested in polynomial time as G is generated by at most two elements (Fact 5.1). If G is not simple, then it has some normal subgroup $N \trianglelefteq G$, and we may try to use a divide-and-conquer strategy by first solving the isomorphism problem for N and G/N. However, even if we find $M \trianglelefteq H$ such that $N \cong M$ and $G/N \cong H/M$, this is typically not sufficient to conclude that $G \cong H$: e.g., $\mathbb{Z}_4$ and $\mathbb{Z}_2 \times \mathbb{Z}_2$ both have $\mathbb{Z}_2$ as a normal subgroup, with corresponding quotient $\mathbb{Z}_2$. We must then understand how the groups N and G/N "glue" back together to get G. G is called an *extension* of N by G/N (some authors use the terminology in the opposite order); given N and Q, understanding the collection of groups G which are extensions of N by Q—that is, where $N \trianglelefteq G$ and $G/N \cong Q$—is known as the *extension problem*. The extension problem is considered quite difficult in general, but the theory of group cohomology exactly captures this problem and provides useful

tools for its study, including connections with other cohomology theories, such as in algebraic topology. One of the main technical achievements of the present paper is to make some aspects of group cohomology effective in the setting of worst-case complexity.

When N is abelian, the extension theory is conceptually easier and technically cleaner. Coincidentally, due to the polynomial-time algorithm for semisimple groups [12], abelian normal subgroups are exactly the subject of interest at present. So for the rest of this subsection, *we assume N is abelian*; the theory for the general case is similar but more complicated, and is covered in section 3.

The extensions of N by Q are governed by two pieces of data: (1) an action of Q on N and (2) a cohomology class. We explain each of these in turn.

The action. If G is an extension of N by Q, then $N \trianglelefteq G$, and thus G acts on N by conjugation, giving a homomorphism $\theta' : G \to \mathrm{Aut}(N)$. As we have assumed N is abelian, N lies in the kernel of θ', so the conjugation action of G on N induces an action θ of $G/N \cong Q$ on N. Two such actions are *compatible* if they become equal after applying some element of $\mathrm{Aut}(N) \times \mathrm{Aut}(Q)$, giving rise to the first problem.

DEFINITION 1.2 (ACTION COMPATIBILITY). *Given two actions $\varphi_i : Q \to \mathrm{Aut}(N)$ of a group Q on a group N—specified by giving, for each $q \in Q$, $\varphi_i(q)$ as a permutation on the set N—decide whether the actions are compatible, that is, whether there is an element of $\mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ whose application to φ_1 makes it equal to φ_2.*

The cohomology class. Informally speaking, the simplest examples of extensions are when Q can be "lifted" to a subgroup of G that is compatible with the isomorphism $G/N \cong Q$ (the extension is said to be *split*). However, it is possible to have an extension G of N by Q in which this cannot happen. For example, consider the additive group of real numbers $\mathbb{R}$ and its normal subgroup $2\pi\mathbb{Z}$. (There are similar examples in finite groups, but we believe this example has more intuitive appeal. For readers familiar with group extensions, the goal here is to exhibit a nonsplit extension; $\mathbb{Z}_2 \trianglelefteq \mathbb{Z}_4$ is a familiar example.) The quotient $\mathbb{R}/2\pi\mathbb{Z}$ is isomorphic to the "circle group" S^1 of unit complex numbers under multiplication, yet S^1 is not even a subgroup of $\mathbb{R}$, let alone "liftable to $\mathbb{R}$." Contrast this with the group $G = 2\pi\mathbb{Z} \times S^1$, which also has $2\pi\mathbb{Z} \trianglelefteq G$ and $G/2\pi\mathbb{Z} \cong S^1$, yet S^1 is a subgroup of G. Note that as both $\mathbb{R}$ and G are abelian, the conjugation action of $\mathbb{R}$ or G on any normal subgroup is trivial. So the actions cannot explain the fact that S^1 is not a subgroup of $\mathbb{R}$; instead, it is group cohomology that exactly captures this phenomenon.

Specifically, if G is an extension of N by Q, the failure of Q to be "liftable" to G (a split extension) is measured by a *cohomology class* as follows. Consider any function $s : Q \to G$ such that $s(q)$ is in the coset of N corresponding to q under the identification $G/N \cong Q$. Q is "liftable" if and only if there is some such s which is also a group homomorphism. The failure of any given s to be a homomorphism is measured by the function

$$f_s(q, p) := s(q)s(p)s(qp)^{-1}.$$

Then s is a homomorphism if and only if $f_s(q, p) = 1$ for all $p, q \in Q$. The cohomology class corresponding to G, viewed as an extension of N by Q, is then the set $\{f_s | s : Q \to G$ as above$\}$. Two cohomology classes are *isomorphic* if they become equal after applying some element of $\mathrm{Aut}(N) \times \mathrm{Aut}(Q)$, giving rise to the second problem.

DEFINITION 1.3 (COHOMOLOGY CLASS ISOMORPHISM). *Given two functions $f_i : Q \times Q \to N$ as above, decide whether there is an element α of $\mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ whose application to f_1 makes $[f_1^\alpha] = [f_2]$.*

Toward a formal strategy. Let us now see how the action and cohomology class just introduced can be used in isomorphism testing. We refer to the pair (θ, f) of the corresponding action and (a representative of) a cohomology class as the *extension data* of the extension. Suppose we are given two groups G_1 and G_2. We cleverly choose some $N_1 \trianglelefteq G_1$ and $N_2 \trianglelefteq G_2$, and (somehow we are lucky to find that) $N_1 \cong N_2$ and $G_1/N_1 \cong G_2/N_2$. Viewing G_i as extensions of N_i by G_i/N_i, we extract the action θ_i and cohomology classes f_i for $i = 1, 2$. If there is a simultaneous solution (one single $(\alpha, \beta) \in \mathrm{Aut}(N) \times \mathrm{Aut}(G/N)$) to ACTION COMPATIBILITY for θ_1, θ_2 and COHOMOLOGY CLASS ISOMORPHISM for f_1, f_2, we say the extension data are *pseudo-congruent*.[3]

DEFINITION 1.4 (EXTENSION DATA PSEUDO-CONGRUENCE). *Given two extension data (θ_i, f_i) for extensions of N by Q—that is, $\theta_i \colon Q \to \mathrm{Aut}(N)$ and $f_i \colon Q \times Q \to N$—decide whether they are pseudo-congruent (see preceding paragraph).*

If the extension data are pseudo-congruent, then $G_1 \cong G_2$ and we are done. However, it is possible that $G_1 \cong G_2$ but the extension data are not pseudo-congruent (we thank Naik [66] for providing Example A.12). The difficulty is that G may contain two normal subgroups $M, M' \trianglelefteq G$ such that $M \cong M'$ and $G/M \cong G/M'$, but no automorphism of G sends M to M'. To resolve this problem, the Main Lemma 3.7 shows that it is enough to take N_1 and N_2 to be the center or the radical, or more generally any characteristic subgroups that are preserved under isomorphisms. (We note that in Besche and Eick [18], they get around the pitfall by introducing the related concept of "strong isomorphism," which is more natural for their purpose, namely the construction of finite groups.)

Now we state the Main Lemma informally. Let us remark that, since in this section we mostly discuss the case of abelian normal subgroups, the Main Lemma is presented in the abelian case here, which is well known (see, e.g., [47, sect. 2.7.4]). We shall develop a general Main Lemma 3.7 (including the case of *nonabelian* normal subgroups) in section 3.

LEMMA 1.5 (Main Lemma, abelian case (Lemma 2.3), informal). *Given two groups G_1 and G_2, let A_i be the abelian characteristic subgroup of G_i of a given type (e.g., the center), θ_i the action of G_i/A_i on A_i, and f_i the cohomology class of the extension of A_i by G_i/A_i given by G_i. Suppose $A_1 \cong A_2$ (identified as A) and $G_1/A_1 \cong G_2/A_2$ (identified as Q).*

Then $G_1 \cong G_2$ if and only if $\theta_1 \equiv \theta_2$, and $f_1 \equiv f_2$ up to the action of $\mathrm{Aut}(A) \times \mathrm{Aut}(Q)$.

As evidence of the usefulness of the Main Lemma beyond this paper, we note that the polynomial-time algorithms for a special class of solvable groups in [58, 70, 16] follow this strategy: They use a theorem of Taunt [84] to reduce isomorphism testing to a problem about linear representations of finite groups (see Problem 1 in [70]), and they solve that problem with additional tactics. In retrospect, Taunt's Theorem is a special case of the Main Lemma, and Problem 1 in [70] is essentially ACTION COMPATIBILITY. Taunt's Theorem applies regardless of whether the normal subgroup

---

[3]We take this terminology from Naik [67], who gives a different definition of pseudo-congruence of extensions that is more standard from the group-theoretic point of view, but less well adapted to the computational setting. We give the other definition and show that the two are formally equivalent in Appendix A.4.1. Robinson [73] uses the term "isomorphism" for this notion; we prefer "pseudo-congruence" to avoid confusion with the several other notions of isomorphism floating around. Theoretical investigations of some aspects of this concept can be found in Robinson [72, sect. 4].

N is abelian or not, though the works [58, 70, 16] only used the case when N is abelian. The general Main Lemma 3.7 additionally covers and extends the nonabelian case of Taunt's Theorem.

Similarly, in retrospect the polynomial-time algorithm for semisimple groups [11, 12] can be viewed as taking advantage of the nonabelian Main Lemma 3.7. We cover these examples in more detail in subsection 4.2.

Due to the structure of the group classes considered in [58, 70, 16], COHOMOLOGY CLASS ISOMORPHISM does not appear in these works. On the other hand, for p-groups of class 2—currently believed to be the bottleneck—COHOMOLOGY CLASS ISOMORPHISM is well known to be necessary (see Fact 4.3). We thus turn to study the COHOMOLOGY CLASS ISOMORPHISM problem in the following. As far as we know, this is the first time group cohomology has been used to get worst-case bounds for GPI.

1.3.1. The general Main Lemma. As mentioned, the Main Lemma in the abelian case is well known, and one contribution of this paper is to extend it to the case when the normal subgroup can be nonabelian. The extension theory in the nonabelian case is classical; a nice introduction can be found in Suzuki's book [82, sect. 2.7]. In this paper, we shall adapt this theory explicitly to the setting of isomorphism testing. The general strategy is the same as in the abelian case, but several technical details need to be taken care of. Consider the extension $N \hookrightarrow G \twoheadrightarrow Q$, where N need not be abelian. An obvious difference from the abelian case is that the conjugation action of G on N no longer induces an action of Q on N, so one needs to consider the homomorphism $Q \to \mathrm{Aut}(N)/\mathrm{Inn}(N)$ instead. As another example, it is important to note that the set of 2-cocycles is no longer a group, let alone an abelian group. Due to all these complications, a careful treatment is needed for the formulation and the proof for the general Main Lemma 3.7, and we refer the interested readers to section 3.

1.4. Overview of our algorithms. Here we give an overview of the structure of our algorithms, as well as some of the more salient details. We first consider the case when the solvable radical is abelian to see how the strategy in the above section is applied. We then focus on central-radical groups to outline some key steps in the algorithms.

Given groups G_1, G_2, we first compute their solvable radicals $A_i = \mathrm{Rad}(G_i)$ and the corresponding semisimple quotients $Q_i = G_i/\mathrm{Rad}(G_i)$. Then we apply the algorithm from [54] to A_1 and A_2, and the algorithm from [12] to Q_1 and Q_2. If either of them returns nonisomorphic, then $G_1 \not\cong G_2$. If both algorithms return isomorphic, they also yield isomorphisms. Thus, without loss of generality, for $i = 1, 2$, we use A to denote $\mathrm{Rad}(G_i)$ and Q to denote $G_i/\mathrm{Rad}(G_i)$, identifying G_i as an extension of A by Q.

Next, we compute the corresponding actions θ_1, θ_2 and representatives f_1, f_2 of the corresponding cohomology classes. As mentioned in subsection 1.3, G_1 and G_2 are isomorphic if and only if there is an element of $\mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ which simultaneously turns θ_1 into θ_2, and $[f_1]$ into $[f_2]$ (as cohomology classes).

For groups with elementary abelian radicals (Theorem B=Corollaries 6.12 and 6.13). Babai et al. [11] showed that all automorphisms of a semisimple group can be enumerated in time $n^{O(\log \log n)}$. So if $n^{O(\log \log n)}$ time is allowed, we can use that algorithm to enumerate all $\beta \in \mathrm{Aut}(Q)$. Then for each such β, search for some $\alpha \in \mathrm{Aut}(A)$ such that $\theta_1 = \theta_2^{(\alpha, \beta)}$, and $[f_1] = [f_2^{(\alpha, \beta)}]$. When $A \cong \mathbb{Z}_p^k$ is elementary abelian, this task can be reduced to MODULE CYCLICITY TESTING over finite-dimensional

algebras, in almost the same way as the reduction from MODULE ISOMORPHISM to MODULE CYCLICITY TESTING [28]. Here we only mention that the algebra is $U = \{\gamma \in M(k,p) \mid \forall q \in Q, \gamma\theta_2(q) = \theta_2(q)\gamma$, and $\exists a \in \mathbb{Z}_p, [\gamma f_2] = [af_2]\}$, and we consider the U-module $V = \{\alpha \in M(k,p) \mid \forall q \in Q, \alpha\theta_1(q) = \theta_2(q)\alpha$, and $\exists a \in \mathbb{Z}_p, [\alpha f_1] = [af_2]\}$. What is left is to verify that V is a cyclic U-module if and only if there exists some desired $\alpha \in \mathrm{GL}(k,p)$. See subsection 6.2 for the details.

For general central-radical groups (Theorem A=Corollary 6.2). For groups with central radicals, $A = Z(G_i)$, so the actions θ_i are trivial and we need only solve COHOMOLOGY CLASS ISOMORPHISM. As before, since $n^{O(\log\log n)}$ time is allowed, we can use the algorithm of Babai et al. [11] to enumerate all $\beta \in \mathrm{Aut}(Q)$. Then for each such β, we need to search for some $\alpha \in \mathrm{Aut}(A)$ such that $[f_1] = [f_2^{(\alpha,\beta)}]$. We solve this problem using linear algebra over abelian groups, as follows. To ease the exposition let us assume $A = \mathbb{Z}_p^k$. Then we shall view any map $f : Q \times Q \to A$ as a $k \times |Q|^2$-size matrix over $\mathbb{Z}_p$, with $\mathrm{Aut}(A)$ acting on the rows and $\mathrm{Aut}(Q)$ inducing an action on the columns. The main difficulty at this point has to do with identifying which cohomology class f is in, in a way that is $\mathrm{Aut}(A)$-invariant. Viewing f as a $\mathbb{Z}_p$-linear vector (of dimension $k \times |Q|^2$), by Proposition 6.10 we can compute a projection π in this vector space such that $\pi(f)$ identifies the cohomology class of f—that is, $\pi(f) = \pi(f')$ if and only if f and f' are in the same cohomology class— and such that π commutes with every $\alpha \in \mathrm{Aut}(A)$ (i.e., π is $\mathrm{Aut}(A)$-invariant). With fixed β, this allows us to compute $\pi(f_1)$ and $\pi(f_2^{(\mathrm{id},\beta)})$, and then determine whether, as $k \times |Q|^2$-size matrices, their row spans are the same, which is a standard task in linear algebra. For central-radical groups with elementary abelian radicals, this approach allows us to compute the coset of isomorphisms. We also give an alternative proof (subsection 6.1.1) that allows us to decide isomorphism for general central-radical groups (where the radical need not be elementary abelian), but the alternative approach does not yield the coset of isomorphisms.

For central-radical groups with $G/\mathrm{Rad}(G)$ a direct product of nonabelian simple groups (Theorem C=Theorem 7.1). In this case $Q = \prod_{i\in[\ell]} T_i$, T_i nonabelian simple. To ease the exposition let us assume T_i's are all mutually isomorphic to some T, and $A = \mathbb{Z}_p^k$. For a function $f : Q \times Q \to A$, a key fact is that the cohomology class of f is completely determined by the *restrictions* of f to the direct factors T_i (Lemma 7.4). Several group-theoretic facts lead to this cohomological proposition, including (1) the direct product decomposition of Q into nonabelian simple factors is unique (not just up to isomorphism); (2) if U_i is the preimage of T_i under the projection $G \to G/Z(G)$, then $u_i u_j = u_j u_i$ whenever $u_i \in U_i$, $u_j \in U_j$, and $i \neq j$ (see [83, Chap. 6, Prop. 6.5]; see also Proposition 7.3). Another useful fact is the well-known description of $\mathrm{Aut}(Q)$ as $\mathrm{Aut}(T) \wr S_\ell$.

Instead of considering $f : Q \times Q \to A$, we can thus consider $f_i : T_i \times T_i \to A$, $i \in [\ell]$; and instead of working with a $k \times |Q|^2$-size matrix, we can work with a $k \times (\sum_{i\in[\ell]} |T_i|^2)$-size matrix. This difference between $|Q|^2 = \prod_{i\in[\ell]} |T_i|^2$ and $\sum_{i\in[\ell]} |T_i|^2$ leads to major savings. When T is a nonabelian simple group, we use the powerful theorem of Guralnick, Kantor, Kassabov, and Lubotzky [43] (reproduced as Theorem 5.3 below); when T is more generally only centerless and perfect, we restrict our attention to the setting where $|T| = O(1)$. In either case, we combine algorithms for LINEAR CODE EQUIVALENCE and COSET INTERSECTION. We need several technical ingredients (including Lemma 7.5) to make the above procedure work, though. As in the previous setting, we give two proofs, one of which handles the general abelian case, and the other of which allows us to compute the coset of isomorphisms in the elementary

abelian case.

1.5. Organization of the paper. In section 2 we collect basic concepts from extension theory. Appendix A contains a gentle introduction to extensions and cohomology, designed to be digestible without first preparing ones palette with half of a textbook. In section 3 we prove our (nonabelian) Main Lemma. We develop our strategy in section 4, which uses the Main Lemma to expand the ideas in subsection 1.3 into a formal framework. Section 5 contains preliminaries and previous algorithmic results to prepare for the algorithms for central-radical groups. In section 6 we describe the $n^{O(\log \log n)}$-time algorithm for general central-radical groups (Theorem A=Corollary 6.2); this is also the algorithm for Corollary 6.3. We also give the algorithms for groups with elementary abelian radicals that need not be central (Theorem B). In section 7, we prove Theorem C, giving the polynomial-time algorithms for central-radical groups with $G/\operatorname{Rad}(G)$ a direct product of nonabelian simple groups (or centerless perfect groups of constant size). Finally, section 8 contains future directions, some of which are motivated by the Cannon–Holt approach and the Babai–Beals filtration.

For any of our results which currently depend on an abelian group being *elementary* abelian, we discuss what obstacles toward this generalization we know how to overcome, and which remain to be overcome: See the end of subsection 6.1.2, subsection 7.4, and section 8.

2. Preliminaries on abelian cohomology. The material in this section is standard group theory; for group theorists, this section serves primarily to fix notation. For computer scientists, we provide a gentle introduction to this material, with motivation and proofs, in Appendix A.

General notation. For $n \in \mathbb{N}$, $[n] = \{1, \dots, n\}$. In this paper, all groups are finite. We use id to denote the identity element, or the group of order 1. For a group G, $|G|$ denotes the order of G. We write $H \leq G$ if H is a subgroup of G. The (right) coset of H in G containing $g \in G$ is $Hg = \{hg \mid h \in H\}$. Given two groups G_1 and G_2, $\operatorname{Iso}(G_1, G_2)$ denotes the set of $G_1 \to G_2$ isomorphisms. $\operatorname{Aut}(G) = \operatorname{Iso}(G, G)$ is the group of automorphisms of G. The set $\operatorname{Iso}(G_1, G_2)$ is either empty or a coset of $\operatorname{Aut}(G_1)$. For $g \in G$, conjugation by g is the automorphism $\theta_g : G \to G$ defined by $\theta_g(x) := gxg^{-1}$. For $g \in G$, the maps θ_g are the inner automorphisms of G, and they form a subgroup $\operatorname{Inn}(G) \leq \operatorname{Aut}(G)$. A subgroup $N \leq G$ is normal if it is invariant under all inner automorphisms, and we write $N \trianglelefteq G$. $N \leq G$ is a characteristic subgroup of G if it is invariant under all automorphisms of G. $Z(G)$ denotes the center of G. For $K, L \leq G$, $[K, L]$ denotes the subgroup generated by all elements of the form $[x, y] := xyx^{-1}y^{-1}$, $x \in K$, and $y \in L$. $[G, G]$ is called the commutator subgroup of G.

Group extension data. Given a finite group G and an abelian normal subgroup $A \trianglelefteq G$, when we consider G as an extension of A by $Q := G/A$, we denote this by $A \overset{\iota}{\hookrightarrow} G \overset{\pi}{\twoheadrightarrow} Q$, where ι is an injective homomorphism and π a surjective homomorphism such that $\operatorname{Ker}(\pi) = \operatorname{Im}(\iota)$. In this paper, we mostly use the "inner" perspective, by identifying A with its image $\iota(A) \trianglelefteq G$. We sometimes refer to G as the "total group" of the extension.

When $A \leq Z(G)$, the action of G on A by conjugation induces an action of Q on A by conjugation, which is the action associated to the extension $A \hookrightarrow G \twoheadrightarrow Q$.

As A is abelian, we write the group operation in A additively, despite the fact that when considering general elements of G we write the group operation in G mul-

tiplicatively (this mixed notation is fairly standard in this setting). Even though A is a subgroup of G, we tend to only use these notations in separate contexts, and it should not cause confusion.

Let $\pi\colon G \to G/A \cong Q$ be the natural projection; then any function $s\colon Q \to G$ such that $\pi(s(q)) = q$ for all $q \in Q$ is called a *section* of π. Any such section s gives rise to a function $f_s\colon Q \times Q \to A$ defined by $f_s(p, q) := s(p)s(q)s(pq)^{-1}$ (by applying π, it is readily verified that the image of f_s is in fact contained in A). We are free to choose $s(1) = \mathrm{id}_G$, and then $f_s(1, q) = f_s(q, 1) = 0$ for all $q \in Q$. Such f are called *normalized*. In the following all sections are normalized unless stated otherwise.

The fact that the group operation in G is associative implies that for all $p, q, r \in Q$,

$$f_s(p, q) + f_s(pq, r) = \theta_p(f_s(q, r)) + f_s(p, qr) \qquad \text{(the 2-cocycle identity)}.$$

Any function $f\colon Q \times Q \to A$ is called a 2-*cochain*; any 2-cochain satisfying the 2-cocycle identity with respect to θ is a 2-*cocycle* (with respect to θ). Given any homomorphism $\theta\colon Q \to \mathrm{Aut}(A)$, every 2-cocycle with respect to θ arises as f_s for some section s of some extension $A \hookrightarrow G \twoheadrightarrow Q$ with action θ.

Given a function $u\colon Q \to A$, the 2-*coboundary* associated to u is the function $b_u\colon Q \times Q \to A$ defined by $b_u(p, q) := u(p) + \theta_p(u(q)) - u(pq)$. Any two 2-cocycles associated to the same extension differ by a coboundary.

The 2-cochains form an abelian group $C^2(Q, A)$ defined by pointwise addition: $(f + g)(p, q) := f(p, q) + g(p, q)$. It is readily visible that the 2-cocycle identity is $\mathbb{Z}$-linear, and hence the 2-cocycles form a subgroup of the 2-cochains, denoted by $Z^2(Q, A, \theta)$. It is similarly verified that the 2-coboundaries form a subgroup of the 2-cocycles, denoted $B^2(Q, A, \theta)$.

A 2-*cohomology class* is a coset of $B^2(Q, A, \theta)$ in $Z^2(Q, A, \theta)$, and any element of this coset is a representative of the cohomology class. If $f \in Z^2(Q, A, \theta)$, we denote the corresponding cohomology class by $[f]$. The group of 2-cohomology classes is denoted $H^2(Q, A, \theta) := Z^2(Q, A, \theta)/B^2(Q, A, \theta)$. By the above discussion, each extension $A \hookrightarrow G \twoheadrightarrow Q$ determines a single cohomology class $[f] \in H^2(Q, A, \theta)$.

We thus arrive at one of the central notions in this paper.

DEFINITION 2.1. *For A an abelian group and Q any group, a pair (θ, f) of an action $\theta\colon Q \to \mathrm{Aut}(A)$ and a 2-cocycle $f\colon Q \times Q \to A$, $f \in Z^2(Q, A, \theta)$, are* extension data*. Two extension data for the pair (Q, A) are* equivalent *if they have the exact same action and if the two* 2*-cocycles are cohomologous (differ by a coboundary).*

Given an extension $A \hookrightarrow G \twoheadrightarrow Q$, the extension data associated to this extension are the action θ as defined above, and any 2-cocycle f_s for any section $s\colon Q \to G$. Note that extension data are nonunique, as we may choose any representative of the corresponding 2-cohomology class. Furthermore, if the action is trivial, then this extension is called *central*. If the 2-cohomology class is trivial, then this extension is called *split*; in this case G is a semidirect product of A by P for some subgroup P.

2.1. Pseudo-congruent extensions versus isomorphic total groups. Recall that a characteristic subgroup is a subgroup invariant under all automorphisms. The analogous notion for isomorphisms (rather than automorphisms) is a function $\mathcal{S}$ that assigns to each group G a subgroup $\mathcal{S}(G) \leq G$ such that any isomorphism $\varphi\colon G_1 \to G_2$ restricts to an isomorphism $\varphi|_{\mathcal{S}(G_1)}\colon \mathcal{S}(G_1) \to \mathcal{S}(G_2)$. We call such a function a *characteristic subgroup function*. Note that if $G_1 = G_2$, this says that $\mathcal{S}(G_1)$ is sent to itself by every automorphism of G_1, that is, $\mathcal{S}(G_1)$ is a characteristic subgroup of G_1. Most natural characteristic subgroups encountered are characteristic

subgroup functions, for example the center $Z(G)$, the commutator subgroup $[G, G]$, or the radical $\text{Rad}(G)$.

DEFINITION 2.2. *Let A be an abelian group and Q any group, and let (θ_1, f_1) and (θ_2, f_2) be two extension data for A-by-Q. Then the extension data are* pseudo-congruent[4] *if there exists $(\alpha, \beta) \in \text{Aut}(A) \times \text{Aut}(Q)$ such that*

(1) $$\theta_1(q)(a) = \alpha^{-1}(\theta_2(\beta(q))(\alpha(a))) =: \theta_2^{(\alpha, \beta)}(q)(a)$$

for all $q \in Q, a \in A$, and

(2) $$f_1(p, q) = \alpha^{-1}(f_2(\beta(p), \beta(q))) + b_u(p, q)$$

for all $p, q \in Q$, and for some 2-coboundary b_u. In this case we write $(\theta_1, f_1) \cong (\theta_2, f_2)$.

LEMMA 2.3 (see, e.g., [47, sect. 2.7.4]). *Let $\mathcal{S}$ be a characteristic subgroup function. Given two finite groups G_1 and G_2, suppose $\mathcal{S}(G_1)$ and $\mathcal{S}(G_2)$ are abelian. Then $G_1 \cong G_2$ if and only if both of the following conditions hold:*
   1. *$\mathcal{S}(G_1) \cong \mathcal{S}(G_2)$ (which we denote by A) and $G_1/\mathcal{S}(G_1) \cong G_2/\mathcal{S}(G_2)$ (which we denote by Q).*
   2. *$(\theta_1, f_1) \cong (\theta_2, f_2)$, where (θ_i, f_i) is the extension data of the extension $A \hookrightarrow G_i \twoheadrightarrow Q$.*

For a detailed proof that does not require reading half a textbook on group theory first, see Appendix A. In the next section we generalize this to the case where the normal subgroup need not be abelian.

3. Nonabelian cohomology and its applications. Here we consider extensions $N \hookrightarrow G \twoheadrightarrow Q$ where N need not be abelian, i.e., the general case. We show that Lemma 2.3 extends to the case when N comes from a characteristic subgroup function—not necessarily abelian—showing the usefulness of the extension theory perspective in its full generality. The results of this section may be of independent interest and of further use in the future, but in this paper they will only be needed for the applications in this section and the next (namely, to show how the results of [12] can be viewed in the same cohomological light as Lemma 2.3, and to reduce the case where G acts by inner automorphisms on its radical to natural subproblems). Suzuki's book [82, sect. 2.7] contains a nice introduction to the extension theory in the nonabelian case, while our contribution here is to adapt this theory explicitly to the setting of isomorphism testing. While some of the basics here are already laid out in Suzuki [82]—so the first part of this section can be viewed as a review to fix notation—here we consider how the extension theory with nonabelian kernels can be used to understand isomorphism of the total groups, which was only considered in a very special case there (reproduced as Theorem 3.15 below).

The action. The first difference to notice when N is nonabelian is that the conjugation map $\theta' \colon G \to \text{Aut}(N)$, defined by $\theta'(g) = c_g$ where $c_g(n) = gng^{-1}$, no longer contains N in its kernel, and hence no longer descends to a map $Q \to \text{Aut}(N)$. However, the action of N on itself by conjugation is by inner automorphisms (by definition) so that we do get a well-defined map $G/N \to \text{Aut}(N)/\text{Inn}(N)$, that is, $Q \to \text{Out}(N)$. For ease of reference, we give a name to such maps.

DEFINITION 3.1. *An* outer action *of a group Q on a group N is a group homomorphism $Q \to \text{Out}(N) = \text{Aut}(N)/\text{Inn}(N)$.*

---

[4]See Footnote 3 on page 1161.

In computations, rather than represent an outer automorphism as a coset of $\mathrm{Inn}(N)$ in $\mathrm{Aut}(N)$, we simply give it by a representative automorphism, and we must remember when we may need to multiply by an arbitrary element of $\mathrm{Inn}(N)$. Throughout this section we use T to denote an action rather than θ, to remind the reader that the essential object here is the outer action represented by T, despite the fact that we work directly with actions $T\colon Q \to \mathrm{Aut}(N)$.

Correspondingly, in the setting of general N, the problem ACTION COMPATIBILITY must be generalized to OUTER ACTION COMPATIBILITY, which is defined as follows. Two actions $T_1, T_2\colon Q \to \mathrm{Aut}(N)$ are said to be "outer equivalent" if there is a function $t'\colon Q \to \mathrm{Inn}(N)$ and an automorphism $\alpha \in \mathrm{Aut}(N)$ such that $T_1(q) = \alpha^{-1} \circ t'(q) \circ T_2(q) \circ \alpha$ for all $q \in Q$.

DEFINITION 3.2 (OUTER ACTION COMPATIBILITY). *Given two actions $T_1, T_2\colon$ $Q \to \mathrm{Aut}(N)$, decide whether there an element $\beta \in \mathrm{Aut}(Q)$ such that T_1 and $T_2 \circ \beta$ are outer equivalent, that is, whether there exists $(\beta, \alpha, t') \in \mathrm{Aut}(Q) \times (\mathrm{Aut}(N) \ltimes \mathrm{Inn}(N)^Q)$ such that $T_1(q) = \alpha^{-1} \circ t'(\beta(q)) \circ T_2(\beta(q)) \circ \alpha$ for all $q \in Q$.*

Although this formulation of OUTER ACTION COMPATIBILITY is more complicated than if we had represented an outer automorphism as a full coset $\theta\mathrm{Inn}(N)$, it will be useful when we formulate EXTENSION DATA PSEUDO-CONGRUENCE below.

Note that when N is abelian there are no inner automorphisms, so $\mathrm{Out}(N) = \mathrm{Aut}(N)$, the only choice for t' above is trivial, and OUTER ACTION COMPATIBILITY then becomes ACTION COMPATIBILITY.

Remark 3.3. We note that, unlike the case when N is abelian, when N is nonabelian it is possible that some outer actions $\theta\colon Q \to \mathrm{Out}(N)$ may not be induced by *any* extension of N by Q. When there *is* such an extension, the outer action θ is called *extendable*. Eilenberg and MacLane [33, sect. 7–9] characterize which outer actions are extendable in terms of the third cohomology group $H^3(Q, Z(N))$. As our interest is primarily in GPI, whenever it matters (e. g., in the definition of OUTER ACTION COMPATIBILITY) we happily restrict our attention to extendable outer actions.

The characterization in terms of cohomology with coefficients in $Z(N)$ allows one to use linear algebra over the abelian group $Z(N)$ to test in polynomial time whether a given outer action is extendable. In particular, the outer action is extendable if and only if an associated third cohomology class vanishes, i.e., is a 3-coboundary. A basis for $B^3(Q, Z(N))$ can be computed analogously to that in Proposition 6.5, by enumerating over a natural basis of the 2-cochains $C^2(Q, Z(N))$ (which has dimension $|Q|^2$) and applying the coboundary operator. Then we just have to compute the 3-cocycle corresponding to the given outer action and check that it lies in the linear subspace spanned by (the subgroup generated by) the 3-coboundaries. Note that we may treat 3-cochains as $k \times |Q|^3$ matrices, analogously to how we treat 2-cochains.

The cohomology class. As in the case when N is abelian, when N is nonabelian, one may still choose a set-theoretic section $s\colon Q \to G$ and get a 2-cocycle $f_s\colon Q \times Q \to N$. This section s gives an action (not just an outer action) $T_s\colon Q \to \mathrm{Aut}(N)$ by conjugation, namely $T_s(q)(n) = s(q)ns(q)^{-1}$. Starting from associativity in G, one may then work out, as in the abelian case, the 2-cocycle condition:

$$f_s(q_1, q_2)f_s(q_1q_2, q_3) = T_s(q_1)(f_s(q_2, q_3))f_s(q_1, q_2q_3) \text{ (nonabelian 2-cocycle identity).}$$

However, this condition depends not just on the action T_s and the 2-cochain f_s, but also on some relationship between T_s and f_s (in this case, that they come from the *same* section s). We would much prefer a condition that does not depend on

the ambient extension group G. To get this condition, note that the action satisfies $T_s(q_1)T_s(q_2) = c_{f_s(q_1,q_2)}T_s(q_1q_2)$, where $c_n \colon N \to N$ denotes the inner automorphism given by conjugation by $n \in N$: $c_n(m) = nmn^{-1}$. This leads us to the definition of extension data for general N.

DEFINITION 3.4 (extension data for general Q, N). *Let Q and N be groups. We say that a pair (T, f) of an action $T \colon Q \to \mathrm{Aut}(N)$ and a function $f \colon Q \times Q \to N$ is extension data if, for all $q_i \in Q$,*

$$(3) \qquad\qquad T(q_1)T(q_2) = c_{f(q_1,q_2)}T(q_1q_2)$$

and

$$(4) \qquad\qquad f(q_1,q_2)f(q_1q_2,q_3) = T(q_1)\left(f(q_2,q_3)\right)f(q_1,q_2q_3).$$

In this case, we refer to f as a 2-cocycle with respect to the action T.

Note that (3) very nearly determines f: It determines $f(q_1, q_2)$ up to an element of $Z(N)$. In particular, when $Z(N) = 1$, (3) actually does determine f completely, a fact we will take advantage of when discussing the polynomial-time algorithm for groups with no abelian normal subgroups [12] (see also Theorem 3.15).

Another difference in the case when N is nonabelian is that although we might denote the set of 2-cocycles by $Z^2(Q, N, T)$, this set will not in general be a group in any natural way, let alone an abelian group. (Also note that it depends on the action T, whereas we know that the action is not intrinsic to the extension—only the corresponding outer action is.) However, the quotient of two 2-cocycles with respect to the same action T will land in the center $Z(N)$, allowing us to reduce part of the question back to the case of N abelian. To see this, let f_1, f_2 be two 2-cocycles with respect to an action $T \colon Q \to \mathrm{Aut}(N)$, and consider conjugating by their quotient $f_1(q_1, q_2)f_2(q_1, q_2)^{-1}$:

$$c_{f_1(q_1,q_2)f_2(q_1,q_2)^{-1}} = T(q_1)T(q_2)T(q_1q_2)^{-1}T(q_1q_2)T(q_2)^{-1}T(q_1) = \mathrm{id}_N.$$

As the center $Z(N)$ consists exactly of those $n \in N$ such that $c_n = \mathrm{id}_N$, we find that the quotient $f_1(q_1, q_2)f_2(q_1, q_2)^{-1}$ lies in $Z(N)$. So although there is not really a cohomology group "$H^2(Q, N, T)$," we can often nonetheless reduce to questions about cohomology classes in $H^2(Q, Z(N), T|_{Z(N)})$. Everything up to this point in this section has been classical (although it has not been leveraged much for algorithms).

Equivalence. As in the case of N abelian, two extensions $N \hookrightarrow G_i \twoheadrightarrow Q$ are said to be *equivalent* if there is an isomorphism $\gamma \colon G_1 \to G_2$ such that γ induces the identity map on both N and Q. However, since $Z^2(Q, N, T)$ is no longer a group and $B^2(Q, N, T)$ no longer its subgroup, the notion of equivalent extensions does not translate so easily to a notion of equivalence for extension data. Hence we derive this condition more or less from scratch. That is, we derive what it means for two extension data to be equivalent by analyzing how two extension data coming from the same extension may differ when a different choice of section $s \colon Q \to G$ is made.

Fix an extension $N \hookrightarrow G \twoheadrightarrow Q$ and two sections $s_1, s_2 \colon Q \to G$. Let $t(q) := s_1(q)s_2(q)^{-1}$; as the s_i are both sections, $s_1(q)$ and $s_2(q)$ are in the same coset of N, so that $t(q) \in N$ for all $q \in Q$. Then the actions $T_1 = \theta_{s_1}$ and $T_2 = \theta_{s_2}$ differ by the inner automorphism $c_{t(q)}$: $T_1(q) = c_{t(q)} \circ T_2(q)$. Recall that we set

$f_i(q,r) := s_i(q)s_i(r)s_i(qr)^{-1}$ for $q,r \in Q$. Then we have that

$$\begin{aligned}
f_1(q,r) &= s_1(q)s_1(r)s_1(qr)^{-1} \\
&= t(q)s_2(q)t(r)s_2(r)s_2(qr)^{-1}t(qr)^{-1} \\
&= t(q)T_2(q)(t(r))s_2(q)s_2(r)s_2(qr)^{-1}t(qr)^{-1} \\
&= t(q)T_2(q)(t(r))f_2(q,r)t(qr)^{-1} \\
&= t(q)T_2(q)(t(r))c_{f_2(q,r)}(t(qr)^{-1})f_2(q,r) =: f_2^{t,T_2}(q,r).
\end{aligned}$$

For future reference, we denote this final expression by $f_2^{t,T_2}(q,r)$.

DEFINITION 3.5. *Two extension data* (T_i, f_i) *are* equivalent *if there is a map* $t\colon Q \to N$ *such that* $T_1(q) = c_{t(q)} \circ T_2(q)$ *for all* $q \in Q$, *and* $f_1 = f_2^{t,T_2}$.

There are several aspects of this definition to take note of:
- By definition, two extension data can be equivalent only if $T_1(q)$ and $T_2(q)$ represent the same outer automorphism of N, in accord with our discussion above.
- When N is abelian, this definition agrees with the previous definition of equivalent extensions. For, in this case, $c_{t(q)} = \mathrm{id}_N$ and the condition $f_1 = f_2^{t,T_2}$ says exactly that f_1 and f_2 differ by the 2-coboundary b_t defined by t.
- $T_1 = T_2$ if and only if $s_1(q)$ and $s_2(q)$ differ by an element of the center $Z(N)$, that is, t is a map $Q \to Z(N)$. In this case, let $T = T_1 = T_2$; then (T, f_1) and (T, f_2) are equivalent if and only if f_1 and f_2 differ by the coboundary $b_t \in B^2(Q, Z(N), T)$. Again, this will be relevant for our discussion below of the polynomial-time algorithm for semisimple groups [12].

Pseudo-congruence. As before, pseudo-congruence is defined as "equivalence up to the action of $\mathrm{Aut}(N) \times \mathrm{Aut}(Q)$."

DEFINITION 3.6 (pseudo-congruence of extension data for general Q, N). *Let* Q *and* N *be two groups. Two extension data* $(T_i, f_i) \in (Q \to \mathrm{Aut}(N), Q \times Q \to N)$ *are* pseudo-congruent *if there exists* $(\alpha, \beta) \in \mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ *such that* (T_1, f_1) *and* $(T_2^{(\alpha,\beta)}, f_2^{(\alpha,\beta)})$ *are equivalent.*

In more detail, the extension data are pseudo-congruent if there exists $(\alpha, \beta) \in \mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ *and* $t\colon Q \to N$ *such that, for all* $q \in Q$ *and all* $n \in N$,

$$(5) \qquad T_1(q)(n) = (\alpha^{-1} \circ c_{t(\beta(q))} \circ T_2(\beta(q)) \circ \alpha)(n)$$

and

$$(6) \qquad f_1(q_1, q_2) = \alpha^{-1}\left[f_2^{t,T_2}(\beta(q_1), \beta(q_2))\right] =: f_2^{(\alpha,\beta,t,T_2)}(q_1, q_2).$$

The problem EXTENSION DATA PSEUDO-CONGRUENCE *is to decide, given two extension data* (T_i, f_i), *whether they are pseudo-congruent.*

LEMMA 3.7 (Main Lemma[5]). *Let* $\mathcal{S}$ *be a characteristic subgroup function. Given two finite (or Lie; see Remark 3.8) groups* G_1 *and* G_2, $G_1 \cong G_2$ *if and only if both of the following conditions hold:*

---

[5] Although the statement of the Main Lemma may not surprise experts, and follows from standard constructions in group cohomology, we have been unable to find a reference for it, and it seems not widely known even among mathematicians, despite the abelian case being very well known. As an example of our Main Lemma not being well known, we point out that the main theorem of a 2003 paper [44] in *L'Enseignment Mathématique*, whose proof takes approximately seven pages, there even assuming knowledge of group cohomology, is a short corollary of our Main Lemma, as shown in Remark 3.8. In that paper, it is even asked whether there are larger classes of groups for which its main theorem holds [44, Rem. 4.2]; our Main Lemma gives quite a general answer to this question.

1. $\mathcal{S}(G_1) \cong \mathcal{S}(G_2)$ (which we denote by N) and $G_1/\mathcal{S}(G_1) \cong G_2/\mathcal{S}(G_2)$ (which we denote by Q).
2. $(T_1, f_1) \cong (T_2, f_2)$, where (T_i, f_i) is the extension data of the extensions $N \hookrightarrow G_i \twoheadrightarrow Q$.

Proof. First suppose that $\gamma\colon G_1 \to G_2$ is an isomorphism. Since $\mathcal{S}$ is a characteristic subgroup function, γ restricts to an isomorphism between the copy $\mathcal{S}(G_1)$ of N in G_1 and the copy $\mathcal{S}(G_2)$ of N in G_2, i.e., an automorphism $\alpha \in \mathrm{Aut}(N)$. Consequently, γ induces an automorphism $\beta := \overline{\gamma} \in \mathrm{Aut}(Q)$. After twisting by these automorphisms, the discussion preceding Definition 3.5 shows that the extension data become equivalent.

Conversely, suppose $(T_1, f_1) \cong (T_2, f_2)$, via $(\alpha, \beta, t) \in \mathrm{Aut}(N) \times \mathrm{Aut}(Q) \times (Q \to N)$. As in the abelian case, we have a standard reconstruction procedure; we construct groups H_i from (T_i, f_i) such that $H_i \cong G_i$, and then we show how the pseudo-congruence of the extension data easily yields an isomorphism $H_1 \cong H_2$.

The underlying set of H_i will be $N \times Q$, with multiplication defined by

$$(n, p) \circ_{H_i} (m, q) = (n \cdot T_i(p)(m) \cdot f_i(p, q), pq).$$

Note that this is the same as in the abelian case, just being careful about the order. Here we have started using dots "$\cdot$" to denote multiplication, as the expressions below get somewhat complicated and this helps to keep things clear. Let $s_i\colon Q \to G_i$ denote the sections used to construct the extension data (T_i, f_i). Then it is readily verified that the map $(n, q) \mapsto ns_i(q)$ gives an isomorphism $H_i \overset{\cong}{\to} G_i$.

Finally, we claim that the map $\varphi(n, p) := (\alpha(n) \cdot t(\beta(p)), \beta(p))$ is an isomorphism from H_1 to H_2. The main fact to check is that this is even a homomorphism. Consider (n, p) and $(m, q) \in H_1$. On the one hand, we have

$$\begin{aligned}
\varphi((n, p) \circ_{H_1} (m, q)) &= \varphi(n \cdot T_1(p)(m) \cdot f_1(p, q), pq) \\
&= (\alpha(n \cdot T_1(p)(m) \cdot f_1(p, q)) \cdot t(\beta(pq)), \beta(pq)) \\
&= (\alpha(n) \cdot \alpha(T_1(p)(m)) \cdot \alpha(f_1(p, q)) \cdot t(\beta(pq)), \beta(pq)).
\end{aligned}$$

On the other hand, we have (here we will sometimes use square brackets $[\,]$ to denote application of an automorphism to help keep all the parentheses straight):

$$\begin{aligned}
\varphi(n, p) \circ_{H_2} \varphi(m, q) &= (\alpha(n) \cdot t(\beta(p)), \beta(p)) \circ_{H_2} (\alpha(m) \cdot t(\beta(q)), \beta(q)) \\
&= (\alpha(n) \cdot t(\beta(p)) \cdot T_2(\beta(p))[\alpha(m) \cdot t(\beta(q))] \cdot f_2(\beta(p), \beta(q)), \beta(p)\beta(q)) \\
&= (\alpha(n) \cdot t(\beta(p)) \cdot T_2(\beta(p))[\alpha(m)] \cdot T_2(\beta(p))[t(\beta(q))] \cdot f_2(\beta(p), \beta(q)), \\
&\qquad \beta(pq)) \\
&= (\alpha(n) \cdot (c_{t(\beta(p))} \circ T_2(\beta(p)))[\alpha(m)] \cdot t(\beta(p)) \cdot T_2(\beta(p))[t(\beta(q))] \cdot f_2(\beta(p), \\
&\qquad \beta(q)), \beta(pq)).
\end{aligned}$$

Let us work through these two expressions bit by bit. We can dispense easily with the second coordinate, as $\beta(pq) = \beta(p)\beta(q)$ since $\beta \in \mathrm{Aut}(Q)$. Both of the first coordinates begin with $\alpha(n)$. Next we have $\alpha(T_1(p)(m))$ on the one hand and $(c_{t(\beta(p))} \circ T_2(\beta(p)))[\alpha(m)]$ on the other. From the definition of pseudo-congruence, we have that $T_1(p)(m) = \alpha^{-1}(c_{t(\beta(p))} \circ T_2(\beta(p)))[\alpha(m)]$. Applying α to both sides of this equation, we see that these two terms are equal.

The remainder of the first coordinate is then $\alpha(f_1(p, q)) \cdot t(\beta(pq))$ in the first case.

From the definition of pseudo-congruence we have

$$
\begin{aligned}
\alpha(f_1(p,q))t(\beta(pq)) &= f_2^{t,T_2}(\beta(p),\beta(q)) \cdot t(\beta(pq)) \\
&= t(\beta(p)) \cdot T_2(\beta(p))[t(\beta(q))] \cdot f_2(\beta(p),\beta(q)) \cdot t(\beta(pq))^{-1} \cdot t(\beta(pq)) \\
&= t(\beta(p)) \cdot T_2(\beta(p))[t(\beta(q))] \cdot f_2(\beta(p),\beta(q)),
\end{aligned}
$$

which is exactly the remainder of the first coordinate in the second case, as desired. Hence φ is a homomorphism.

Finally, it suffices to show that φ is injective, for as $|H_1| = |N||Q| = |H_2|$, it will then follow that φ is bijective and hence an isomorphism. Consider the kernel of φ: $\varphi(n,p) = (1,1)$. As the second coordinate is 1, we have $\beta(p) = 1$ and hence $p = 1$. As the first coordinate is 1, we have $\alpha(n)t(\beta(p)) = \alpha(n) = 1$, so we also have $n = 1$. (In the first equality we use the fact that $t(1) = 1$, which follows from $T_i(1_Q) = \mathrm{id}_N$.) Hence φ is injective and thus an isomorphism. □

Remark 3.8 (the Main Lemma for Lie groups). Lemma 3.7 also extends to the case of Lie groups, allowing us to show that the main theorem of Hämmerli [44] follows from our Main Lemma as a quick corollary, as well as giving a very general answer to a question he posed. The main theorem of Hämmerli [44] is essentially the special case of our Main Lemma in which the characteristic subgroup is taken to be the connected component of the identity. Hämmerli asked [44, Rem. 4.2] whether his main theorem extended to other classes of groups; our Main Lemma extends it greatly and shows that the result has very little to do with Lie groups per se.

To see that our Main Lemma extends to Lie groups, note that the only place we used finiteness in the proof of Main Lemma 3.7 is in the final paragraph, to get surjectivity from injectivity. In the case of Lie groups, we also note that the map in the above proof is continuously differentiable (even smooth). The rest of the argument is essentially dimension-counting, but we give it here for completeness. First, because the homomorphism is differentiable, it descends to a map of Lie algebras, and injectivity of the map implies injectivity of the corresponding map of Lie algebras. As the Lie algebras are, in particular, finite-dimensional vector spaces of the same dimension, injectivity and linearity imply surjectivity, so we have an isomorphism of Lie algebras. For a Lie group G, let $G^{(0)}$ denote the connected component of the identity of G. Continuity of the homomorphism and the fact that it induces an isomorphism of Lie algebras implies that it induces an isomorphism of $H_1^{(0)} \cong H_2^{(0)}$. As $G/G^{(0)}$ is a finite group for any Lie group G, injectivity and continuity together imply that we have an isomorphism of the component groups $H_1/H_1^{(0)} \cong H_2/H_2^{(0)}$. As the map we started with was an injective homomorphism that induces an isomorphism of the identity components and of the component groups, it is an isomorphism.

3.1. Application to extensions with trivial outer action. We did not define Cohomology Class Isomorphism for general N and then proceed to pseudo-congruence, as in the abelian case, because it turns out that when the outer action is trivial, Cohomology Class Isomorphism for action-trivial extensions of N by Q reduces to Cohomology Class Isomorphism for extensions of $Z(N)$ by Q. To prove this we use one additional concept, that of a central product. Although this notion generalizes to an arbitrary number of factors, we only need the two-factor case.

Definition 3.9 (central decomposition; see, e.g., [89]). *A pair $\{G_1, G_2\}$ of subgroups of a finite group G is a* central decomposition *of G if G is generated by G_1 and G_2 ($G = \langle G_1, G_2 \rangle$) and G_1 and G_2 commute ($[G_1, G_2] = 1$).*

LEMMA 3.10. *Let $N \hookrightarrow G \xrightarrow{\pi} Q$ be an extension of N by Q which induces the trivial outer action $\theta(q) = \mathrm{id}_N \mathrm{Inn}(N)$ for all $q \in Q$. Then there is a subgroup H of G such that $H \cap N = Z(N)$, $\pi(H) = Q$, and $\{N, H\}$ is a central decomposition of G. We denote this subgroup H by $G|_{Z(N)}$.*

Proof. There is a section $s \colon Q \to G$ such that $c_{s(q)} = \mathrm{id}_N$ for all $q \in Q$. Let $f(p, q) = f_s(p, q) = s(p)s(q)s(pq)^{-1}$ be the 2-cocycle corresponding to s. As $c_{s(q)} = \mathrm{id}_N$ for all $q \in Q$, we also have that $c_{f(p,q)} = \mathrm{id}_N$ for all $p, q \in Q$. As $f(p, q) \in N$, this implies that $f(p, q) \in Z(N)$. Hence f is a 2-cocycle in $H^2(Q, Z(N))$ (for the trivial action of Q on $Z(N)$). Let H be the subgroup generated by $Z(N)$ and $s(Q)$. Since $f(p, q) \in Z(N)$ for all $p, q \in Q$, it follows that every element of H can be represented uniquely in the form $zs(q)$ for $z \in Z(N), q \in Q$.

From the uniqueness of the representation $zs(q)$, it follows immediately that $H \cap N = Z(N)$. Since H included $s(q)$ for all $q \in Q$, it follows that $\pi(H) = Q$. Finally, to see that $[H, N] = 1$, consider $zs(q)n(zs(q))^{-1}n^{-1} = zs(q)ns(q)^{-1}z^{-1}n^{-1}$. Since $c_{s(q)} = \mathrm{id}_N$, this equals $znz^{-1}n^{-1}$, but since $z \in Z(N)$, the latter is trivial.   ☐

The preceding lemma nearly allows us to reduce group isomorphism when the outer action of G on $\mathrm{Rad}(G)$ is trivial to isomorphism of a pair of central radical groups and a pair of solvable groups. However, up to this point we have brushed over the fact that central products are not uniquely determined by their factors. In a central decomposition $\{G_1, G_2\}$, as in a direct product decomposition, it is true that both G_i are normal subgroups. Unlike a direct decomposition, however, $G_1 \cap G_2$ need not be trivial. To make the discussion a little clearer, we introduce a standard alternative viewpoint on central decompositions.

DEFINITION 3.11 (central product; see, e.g., [2, (11.1)]). *Given two groups H_1, H_2 and an isomorphism $\varphi \colon Y_1 \to Y_2$ between two subgroups $Y_i \leq Z(H_i)$ of their centers, the quotient of $H_1 \times H_2$ by $\{(y^{-1}, \varphi(y)) : y \in Y_1\}$ is the* central product *of H_1 and H_2 along φ, denoted $H_1 \times_\varphi H_2$.*

Central products and central decompositions are essentially equivalent. More specifically, if $\{G_1, G_2\}$ is a central decomposition of G, then if we let $Y_i = G_1 \cap G_2 \leq G_i$ for $i = 1, 2$, and define $\varphi \colon Y_1 \to Y_2$ to be the map induced by the identity map on $G_1 \cap G_2$ (thinking of both G_1 and G_2 as subgroups of G), then we see that G is the central product of G_1 and G_2 along φ. Conversely, if $H = H_1 \times_\varphi H_2$ is a central product, then every element of H can be written (not uniquely!) as the equivalence class of (h_1, h_2) in the quotient $H_1 \times_\varphi H_2$ for some $h_1 \in H_1, h_2 \in H_2$. Let us denote this equivalence class by $\overline{(h_1, h_2)}$. Then it is readily verified that $G_1 = \{\overline{(h_1, 1)} : h_1 \in H_1\}$ is a subgroup of H isomorphic to H_1, that $G_2 = \{\overline{(1, h_2)} : h_2 \in H_2\}$ is a subgroup of H isomorphic to H_2, and that $\{G_1, G_2\}$ is a central decomposition of H.

When dealing with isomorphisms between central products, the fact that $G_1 \cap G_2$ is nontrivial becomes a source of difficulty, as in the following lemma. Although this lemma applies in a more general situation, we state it for the situation we are most interested in.

LEMMA 3.12. *Let A be a solvable group and B a central-radical group. Suppose that φ_1, φ_2 are two isomorphisms $Z(A) \to Z(B)$. Then $A \times_{\varphi_1} B$ is isomorphic to $A \times_{\varphi_2} B$ if and only if there are automorphisms $\alpha \in \mathrm{Aut}(A)$ and $\beta \in \mathrm{Aut}(B)$ such that $\varphi_1 = \beta^{-1} \circ \varphi_2 \circ \alpha$.*

Proof. Let G_i denote $A \times_{\varphi_i} B$, and let Z_i denote the copy of $Z(A) \cong Z(B)$ in G_i for $i = 1, 2$. Every element of G_i can be written—not necessarily uniquely—as

ab with $A \in A$ and $b \in B$; $ab = 1$ in G_i if and only if $a \in Z(A)$, $b \in Z(B)$, and $\varphi_i(a) = b^{-1}$.

First, suppose there are $\alpha \in \mathrm{Aut}(A)$ and $\beta \in \mathrm{Aut}(B)$ such that $\varphi_1 = \beta^{-1} \circ \varphi_2 \circ \alpha$. Then we claim that the map sending $ab \in G_1$ to $\alpha(a)\beta(b) \in G_2$ is both well defined and an isomorphism. To see that it is well defined, note that if $ab = 1$ in G_1, then $\varphi_1(a) = b^{-1} \in Z(G_1)$. By assumption, we then have $\beta^{-1}(\varphi_2(\alpha(a))) = b^{-1}$, or equivalently $\varphi_2(\alpha(a)) = \beta(b)^{-1}$, which means that $\alpha(a)\beta(b) = 1$ in G_2. From this, one concludes that if $ab = a'b'$, then $\alpha(a)\beta(b) = \alpha(a')\beta(b')$; it is then easily verified that this map is in fact a homomorphism. Since these are finite groups, injectivity then suffices to show it is an isomorphism; injectivity follows using the preceding argument run in reverse.

For the reverse direction, we will need to use the following fact. Let $\pi_i \colon A \times B \to A \times_{\varphi_i} B$ be the natural quotient map. Then we claim that $\pi_i(1 \times Z(B)) = Z(G_i)$. It is easy to see that $\pi_i(1 \times Z(B)) \leq Z(G_i)$. For the reverse inclusion, suppose $z \in Z(G_i)$. Then z must commute with both $\pi_i(1 \times B)$, and since π_i maps $1 \times B$ isomorphically onto its image, this means that $z \in \pi_i(1 \times Z(B))$.

Now, suppose there is an isomorphism $\gamma \colon G_1 \to G_2$. Note that as A is a solvable normal subgroup of G_i and $G_i/A \cong B/Z(B)$ is semisimple, A must be equal to $\mathrm{Rad}(G_i)$. In particular, since the radical is a characteristic subgroup function, $\gamma(A) = A$, and thus $\alpha := \gamma|_A$ is an automorphism of A.

Next we show that γ sends B to B, thus inducing an automorphism of B. Note that $G_i/Z(G_i) \cong (A/Z(A)) \times (B/Z(B))$. Since the center is a characteristic subgroup function, γ induces an isomorphism $G_1/Z(G_1) \to G_2/Z(G_2)$, that is, an automorphism of $(A/Z(A)) \times (B/Z(B))$. Since A is solvable and $B/Z(B)$ is semisimple, γ thus induces an automorphism of $B/Z(B)$ as well. Since $\gamma(Z(B)) = \gamma(Z(G_1)) = Z(G_2) = Z(B) \leq B$, by a counting argument γ must send B to B. Thus γ induces $\beta \in \mathrm{Aut}(B)$.

Thus, by construction, for $a \in A$ and $b \in B$, $\gamma(a) = \alpha(a)$ and $\gamma(b) = \beta(b)$. Since γ is a homomorphism, it follows that $\gamma(ab) = \alpha(a)\beta(b)$ for all $a \in A$ and all $b \in B$. To see that $\varphi_1 = \beta^{-1} \circ \varphi_2 \circ \alpha$, one uses the same calculation as in the first direction. □

The CENTRAL AMALGAM PROBLEM is as follows: Given two automorphisms $\varphi_i \in \mathrm{Aut}(Z)$, $i = 1, 2$, of an abelian group Z, two black-box groups G, H (think of these as $\mathrm{Aut}(A)$ and $\mathrm{Aut}(B)$ in the preceding lemma), and actions of G and H on Z—given by specifying the matrix actions of *generating sets* of G and H—decide whether there exist $\alpha \in G$ and $\beta \in H$ such that $\varphi_1 = \beta^{-1} \circ \varphi_2 \circ \alpha$.

PROPOSITION 3.13. *Let $\mathcal{S}$ be a polynomial-time computable characteristic subgroup function. Suppose that G_1, G_2 are two groups for which the induced outer action of $G_i/\mathcal{S}(G_i)$ on $\mathcal{S}(G_i)$ by conjugation is trivial (equivalently, the induced action is by inner automorphisms of $\mathcal{S}(G_i)$). Then the group isomorphism problem for (G_1, G_2) reduces in polynomial time to finding a generating set for $\mathrm{Aut}(\mathcal{S}(G_i))$ and $\mathrm{Aut}(G_i|_{Z(\mathcal{S}(G_i))})$ and solving the CENTRAL AMALGAM PROBLEM.*

In particular, GPI for groups for which the outer action of $G/\mathrm{Rad}(G)$ on $\mathrm{Rad}(G)$ is trivial reduces in $n^{O(\log \log n)}$-time to finding generating sets of the automorphism group of solvable groups and solving the CENTRAL AMALGAM PROBLEM.

Remark 3.14. When $|\mathrm{Aut}(Z(\mathrm{Rad}(G)))|$ is bounded by $\mathrm{poly}(|G|)$, the CENTRAL AMALGAM PROBLEM can be solved in $\mathrm{poly}(|G|)$ time by standard permutation group algorithms. Note that the class of solvable groups S whose centers have automorphism groups of polynomial size may seem restrictive, but is in fact quite rich. In particular, it includes solvable groups whose centers are abelian groups with arbitrarily many

factors, as long as each prime appears in a bounded number of factors, and also includes all centerless solvable groups (itself quite a nontrivial class of groups).

Proof. We show how to construct a central decomposition as in Lemma 3.10 in polynomial time. Let $N = \mathcal{S}(G)$ and $Q = G/N$. By assumption, the subset $N = \mathcal{S}(G)$ can be identified in poly($|G|$) time. Next, choose any section $s\colon Q \to G$. It may be that some $s(q)$ acts nontrivially on N via conjugation. However, by the assumption that the outer action is trivial, $c_{s(q)}$ must be some inner automorphism of N, say $c_{n(q)}$ for some $n(q) \in N$. To find this $n(q)$, we may search through N exhaustively in at most $O(|N|^2) \le O(|G|^2)$ time: Essentially $|N|$ steps to check the action of a given n on N by conjugation, and there are $|N|$ possible n's to check. Then let $s'(q) = s(q)n(q)^{-1}$; as $n(q) \in N$, s' is another section, and by construction $c_{s'(q)} = \mathrm{id}_N$ for all N. Finally, let $f(p,q) = s'(p)s'(q)s'(pq)^{-1}$. Computing all the values of f takes essentially $O(|Q|^2) \le O(|G|^2)$ time, and then we construct $G|_{Z(\mathcal{S}(G))}$ as the subgroup of G generated by $s(Q)$ and $Z(N)$ in poly($|G|$) time. $\square$

Extensions of centerless groups. We have already mentioned a few useful properties of extensions of centerless groups, that is, when $Z(N) = 1$. One that is implicit in what we have already said is that every outer action $Q \to \mathrm{Out}(N)$ is extendable; that is, it is induced from some extension of N by Q. These properties culminate in the following very useful theorem.

THEOREM 3.15 (see, e.g., [82, Thm. 2.7.11]). *Let N be a centerless group, Q any group, and G an extension of N by Q. Then G is determined up to isomorphism by the induced outer action of Q on N.*

Furthermore, every such extension is equivalent to a subgroup $\Gamma \le Q \times \mathrm{Aut}(N)$ satisfying $\Gamma \cap \mathrm{Aut}(N) = \mathrm{Inn}(N)$ and $\pi_Q(\Gamma) = Q$, where $\pi_Q\colon Q \times \mathrm{Aut}(N) \to Q$ is the projection onto the first factor.

In particular, if $\mathcal{S}$ is a characteristic subgroup function computable in polynomial time, and $\mathfrak{C}$ is a class of groups for which $\mathcal{S}(G)$ is centerless for every $G \in \mathfrak{C}$, then isomorphism of groups in $\mathfrak{C}$ reduces to isomorphism of groups of the type $G/\mathcal{S}(G)$ for $G \in \mathfrak{C}$, groups of the type $\mathcal{S}(G)$ for $G \in \mathfrak{C}$, and OUTER ACTION COMPATIBILITY.

4. The strategy. Suppose we are given two groups G_1 and G_2 from some class of groups $\mathfrak{C}$. Our Main Lemma 3.7 suggests (and indeed was motivated by) a divide-and-conquer strategy to test isomorphism (subsection 4.1). This strategy highlights important structural features of GpI, which we show are formally necessary in subsection 4.3. It naturally suggests new group classes for which polynomial-time isomorphism tests might be within reach, and also suggests a priori many group classes for which polynomial-time algorithms have previously been achieved.

However, before we proceed, let us emphasize that the extension viewpoint only helps with a *conceptual* understanding of these previous works. Given this viewpoint, to tackle each group class may still require novel mathematical ideas and technically demanding algorithms. The extension viewpoint is mostly used to set the stage for applications of such mathematical and algorithmic techniques. In other words, instead of looking at a table encoding an *abstract* group, an application of the Main Lemma usually transfers us to a more *concrete* setting where we need to solve problems about, e.g., bilinear maps and permutation groups. We also stress that not all recent progress on GpI can be captured from this extension viewpoint; see, e.g., [92, 59].

4.1. A recipe for group isomorphism.
  1. Choose wisely a polynomial-time computable characteristic subgroup function

$\mathcal{S}$. Note that if $\mathcal{S}(G)$ is always abelian, then the technically simpler abelian Lemma 2.3 can be applied.

2. Test whether $\mathcal{S}(G_1) \cong \mathcal{S}(G_2)$ (which we henceforth refer to as N) and $G_1/\mathcal{S}(G_1) \cong G_2/\mathcal{S}(G_2)$ (which we refer to as Q). If either of these fails, then $G_1 \not\cong G_2$.

3. Extract the extension data (T_i, f_i) from the extension $N \hookrightarrow G_i \twoheadrightarrow Q$ for $i = 1, 2$ by picking arbitrary sections $s \colon Q \to G_i$ and computing the action and cohomology class.

4. Test pseudo-congruence of the two extension data. That is, find $(\alpha, \beta) \in \mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ and a function $t \colon Q \to N$ such that $T_1(q) = c_{t(q)} \circ T_2^{(\alpha,\beta)}(q)$ and $f_1 = (f_2^{(\alpha,\beta)})^{t,T_2}$. If the abelian version of the Main Lemma (Lemma 2.3) applies, then t is unnecessary.

Some general remarks are due for each of these steps:

1. A seemingly obvious requirement would be that $\mathcal{S}(G)$ should not be trivial for any $G \in \mathfrak{C}$. However, even if this is not the case, it may be fruitful to consider separately the class of groups for which $\mathcal{S}(G)$ is trivial. For example, semisimple groups arise this way, as those groups for which $\mathrm{Rad}(G)$ is trivial.

2. Due to the nature of the divide-and-conquer strategy, $\mathcal{S}(G)$ and $G/\mathcal{S}(G)$ should be from group classes with known efficient isomorphism tests. Alternatively, if, say, $\mathcal{S}(G)$ is not from such a class, it may be possible to use this strategy to reduce isomorphism of groups in $\mathfrak{C}$ to isomorphism of groups of the form $\mathcal{S}(G)$ for $G \in \mathfrak{C}$ (or similarly for $G/\mathcal{S}(G)$).

3. This step is easy ($\mathrm{poly}(|G|)$ time). Based on the group class $\mathfrak{C}$ the extension data will hopefully turn out to have a nice mathematical structure; indeed, looking for this nice mathematical structure is a nice heuristic that can help motivate and suggest various choices for $\mathfrak{C}$.

4. This pseudo-congruence test is the main bottleneck. Choosing $\mathcal{S}$ so that this step can take advantage of known cohomological results may be helpful. For example, if $\mathcal{S}(G) \le Z(G)$, then Extension Data Pseudo-congruence simplifies to Cohomology Class Isomorphism; at the opposite end of the spectrum, if $G = \mathcal{S}(G) \rtimes (G/\mathcal{S}(G))$, then Extension Data Pseudo-congruence simplifies to Action Compatibility. As another example, if $\mathcal{S}(G)$ is centerless, one may take advantage of Theorem 3.15, as in the case of semisimple groups (see below).

4.2. Some recent results from the point of view of the Main Lemma. As mentioned in the introduction, there have been some recent polynomial-time algorithms for several group classes: semisimple groups [11, 12], generalized Heisenberg groups [60], groups with abelian Sylow towers [16], and (in this paper) $n^{o(\log n)}$-time algorithms for central-radical groups. The experts would easily see how the perspective of group extension helps to open a venue of attack to devise efficient algorithms for these group classes. However, for readers who have not seen such a connection, the definitions of these group classes may at first seem obscure, and it is not a priori clear why we should have found efficient algorithms for these particular classes of groups, as opposed to others. We believe that the viewpoint of extensions and cohomology, especially in light of the Main Lemma, gives a unifying perspective to these works which helps to explain the progress on these group classes, thereby easing certain readers' understanding of these previous works.

In the following, we first summarize some basic information about these works, and then explain in detail how previous works on GpI fit into the general strategy

described as above.

References	Group class	Characteristic subgroup function	Extension type
[11, 12]	semisimple groups	socle[6]	extension of a centerless group
[60]	quotients of generalized Heisenberg groups	center	special type of central extensions of $\mathbb{Z}_p^k$ by $\mathbb{Z}_p^\ell$
[58, 70, 16]	groups with abelian Sylow towers	normal Hall subgroups	split extension of A by Q with $(|A|, |Q|) = 1$
this work	central-radical groups	solvable radical	(central) extension of abelian groups by semisimple groups
follow-up work [42]	groups with "tame" radicals	any "tame" abelian	tame extension of A by Q

Semisimple groups (groups with no abelian normal subgroups). In the polynomial-time algorithm for semisimple groups [12] we take $\mathcal{S} = \mathrm{Soc}$, i.e., $N = \mathrm{Soc}(G)$, which is a polynomial-time characteristic subgroup function. Hence the general Main Lemma 3.7 applies and isomorphism of semisimple groups reduces to EXTENSION DATA PSEUDO-CONGRUENCE. When G is semisimple, its socle is a direct product of nonabelian simple groups, and thus $Z(N) = 1$. N being centerless simplifies some of the results in the previous section, as captured in Theorem 3.15, which corresponds to the key lemma in [11] (Lemma 3.1 therein), and leads to the problems considered by Babai et al. [11, 12]. In particular, note that in the definition of pseudo-congruence for nonabelian N, after twisting by $(\alpha, \beta) \in \mathrm{Aut}(N) \times \mathrm{Aut}(Q)$ to make the actions T_1, T_2 become equivalent as outer actions, the condition on the 2-cocycles is simply that they differ by a 2-coboundary in $B^2(Q, Z(N), T)$. In particular, when N is centerless, $B^2(Q, Z(N), T)$ is trivial, so EXTENSION DATA PSEUDO-CONGRUENCE reduces to OUTER ACTION COMPATIBILITY.

In the case of semisimple groups, using the structure of these groups, one sees quickly that OUTER ACTION COMPATIBILITY reduces to TWISTED CODE EQUIVALENCE (introduced in [12]), where the "twisting" groups correspond to the action of $\mathrm{Out}(N) = \mathrm{Out}(\mathrm{Soc}(G))$ in the definition of OUTER ACTION COMPATIBILITY, and the choice of $t\colon Q \to N$ is handled by considering codes whose codewords correspond to elements of G rather than just elements of Q.

p-groups of class 2 and exponent p, especially quotients of generalized Heisenberg groups. For p-groups of class 2 and exponent p with odd p, Baer's correspondence [17] suggests considering the alternating bilinear maps defined by the commutator bracket: Isomorphism of p-groups corresponds to pseudo-isometry of these bilinear maps. These bilinear maps are 2-cocycles, and two such cocycles are isomorphic as cohomology classes if and only if the bilinear maps are pseudo-isometric, so we see that this is a particular instance of COHOMOLOGY CLASS ISOMORPHISM and Baer's correspondence can be viewed as a special case of the abelian version of the Main Lemma (Lemma 2.3).

The bilinear map viewpoint has been the main stage for the recent progress on testing isomorphism of such p-groups [60, 23]. For example, Lewis and Wilson [60] studied a decently large class of p-groups—quotients of generalized Heisenberg groups—which are indistinguishable to the eyes of classical invariants but for which they nonetheless present a polynomial-time isomorphism test. Such groups admit

---

[6]The *socle* of a group G is the subgroup generated by the union of the minimal normal subgroups of G.

a nice characterization from the bilinear map viewpoint [60, Thm. 3.13], and the polynomial-time isomorphism test for these groups takes advantage of the special structure of the bilinear maps corresponding to these groups [60, Thm. 4.1]. We remark that the algorithm in [60] works with much more succinct models for representing groups, including permutation groups and matrix groups, and runs in time polynomial in the input size, which can be as small as $\log |G|$ instead of $|G|$.

Groups with abelian Sylow towers. Though solving GpI for the obscure-sounding group class "groups with abelian Sylow towers," the core of [16] (following [58, 70]) deals with the case of *coprime extensions*, namely extension of an abelian A by Q where $(|A|, |Q|) = 1$. The Schur–Zassenhaus Theorem guarantees that coprime extensions split, thus reducing EXTENSION DATA PSEUDO-CONGRUENCE to ACTION COMPATIBILITY for such groups. Assuming $\mathrm{Aut}(Q)$ is known (via recursive divide-and-conquer), [16] views the actions of Q on A as linear representations and utilizes the complete reducibility of these representations by Maschke's Theorem, which requires coprimality. The tactic they use is to view the induced action of $\mathrm{Aut}(Q)$ on the irreducible constituents as a permutation group action, and then to develop a parameterized permutation group algorithm to finally solve ACTION COMPATIBILITY in this case [16].

Tame extensions. In a follow-up work [42], the current authors used the viewpoint of this paper to generalize the preceding from coprime extensions to so-called "tame" extensions. These are extensions of $\mathbb{Z}_p^k$ by Q (assuming $\mathrm{Aut}(Q)$ is known, e.g., by recursive divide-and-conquer) where the Sylow p-subgroups of Q are cyclic, or $p = 2$ and the Sylow 2-subgroups are dihedral, semidihedral, or generalized quaternion. (The case above is when the Sylow p-subgroups of Q are trivial.)

In fact, the story behind [42] is a perfect example of the utility of explicitly splitting GpI into ACTION COMPATIBILITY and COHOMOLOGY CLASS ISOMORPHISM. Namely, independently, one of the current authors had solved ACTION COMPATIBILITY for the tame case, and the other had solved COHOMOLOGY CLASS ISOMORPHISM for the tame case under the assumption that ACTION COMPATIBILITY could be solved; when they met in Chicago each was eager to tell the other of their result, asking if the other "half" of the problem could be solved. The result was nearly immediate from there.

Central-radical groups. Similarly, by considering cohomology rather than actions, we will see in the following how to handle central-radical groups. An elementary way of manipulating the 2-cohomology classes yields an $n^{O(\log \log n)}$-time algorithm for groups with central radicals. For a subclass of groups with central radicals, a more detailed understanding of 2-cohomology classes (Lemma 7.4) helps establish the polynomial-time algorithms in Theorem C=Theorem 7.1. In particular, singly exponential algorithms for LINEAR CODE EQUIVALENCE and for COSET INTERSECTION inevitably enter into the algorithm for this theorem.

Summary. In summary, although the strategy we propose will rarely solve the problem completely, even for a restricted class of groups, it has the virtue of quickly dispensing with structural issues to highlight the needed algorithmic tactics.

4.3. Necessity of pseudo-congruence and cohomology. The Main Lemma (Lemma 2.3; Lemma 3.7) suggests studying EXTENSION DATA PSEUDO-CONGRUENCE to make progress toward GpI for groups with abelian normal subgroups. In this section, we shall see that pseudo-congruence tests for certain classes of extension data are exactly isomorphism tests for certain interesting group classes. While the Main Lemma 3.7 almost implies so, a pitfall is that in the reconstruction procedure we need

the normal copy of A in H to be the image of a characteristic subgroup function. (In the setting of Lemmas 2.3 and 3.7 the standard reconstruction procedure does return groups H_i with the copy of $A = \mathcal{S}(H_i)$, but this is because of the assumptions of that lemma.) This leads us to look at some concrete classes of extension data, for which this property holds.

For split extensions, a well-known example is the case when $|A|$ and $|Q|$ are coprime, as ensured by the Schur–Zassenhaus Theorem. In this case G is said to be a coprime extension of A by Q, and A is a normal Hall subgroup in G. Noting that taking a normal Hall subgroup of a specific order is a characteristic subgroup function, with the standard reconstruction procedure we have the following fact.

FACT 4.1. *There is a polynomial-time function r which takes any group action $\theta \colon G \to \mathrm{Aut}(B)$ (for any groups B, G) to a group $r(\theta)$ with the following property. When A is abelian, Q is a group of order coprime to $|A|$, and $\theta_i \colon Q \to \mathrm{Aut}(A)$ ($i = 1, 2$) are group actions, then $(\theta_1, \theta_2) \mapsto (r(\theta_1), r(\theta_2))$ is a Karp reduction from these instances of* ACTION COMPATIBILITY *to* GPI.

A polynomial-time algorithm for ACTION COMPATIBILITY for the case in Fact 4.1 was given by Babai and Qiao [16], yielding a polynomial-time time algorithm for "groups with abelian Sylow towers."

Remark 4.2. Despite the polynomial-time algorithm for ACTION COMPATIBILITY for coprime case—which trivially implies there is a Karp reduction to *any* problem, including GPI—Fact 4.1 remains nontrivial. In particular, recall that given equivalence relations $\sim_1$ on X_1 and $\sim_2$ on X_2, a *kernel reduction* between them is a function $f \colon X_1 \to X_2$ such that $x \sim_1 y$ if and only if $f(x) \sim_2 f(y)$ [36]. Fact 4.1 states that there is a polynomial-time kernel reduction, which does *not* follow automatically from the polynomial-time algorithm for coprime ACTION COMPATIBILITY.

For central extensions, let $p \neq 2$ be a prime. If A and Q are both abelian, then a $\mathbb{Z}$-bilinear map $f \colon Q \times Q \to A$ is a 2-cocycle, as the cocycle identity follows directly from bilinearity. Note that the action of $\mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ preserves bilinearity. The following proposition is known (see [17, 57]; see also [87, sect. 5] and [89]); the standard reconstruction procedure is altered to make the image of A the commutator subgroup.

FACT 4.3. *Given a prime $p \neq 2$ and finite abelian p-groups A and Q, let $f_i \colon Q \times Q \to A$ be an alternating bilinear map. Then* COHOMOLOGY CLASS ISOMORPHISM *for f_1, f_2 Karp-reduces to* GPI.

Proof. Given f_i for $i = 1, 2$, alter the standard construction as follows. For $a, b \in A$ and $q, q' \in Q$, we define the group G_i with operation $\circ$ over the set $A \times Q$ as $(a, q) \circ (b, q') = (a + b + \frac{1}{2} f_i(q, q'), q + q')$. It is known that G_i's are p-groups of class 2, the copy of A in G_i is the commutator subgroup, and $f_1 \cong f_2$ if and only if $G_1 \cong G_2$ (see, e.g., [87, sect. 5]). □

Finally, let us examine groups whose solvable radicals are abelian, a superclass of central-radical groups. When Q is semisimple, the standard reconstruction procedure sends A to the solvable radical. (A solvable normal subgroup N is the solvable radical if and only if G/N is semisimple.) This hints at the fact that for central-radical groups, group isomorphism is equivalent to COHOMOLOGY CLASS ISOMORPHISM.

FACT 4.4. *Let A be abelian and Q semisimple. For $i = 1, 2$, let $\theta_i \colon Q \to \mathrm{Aut}(A)$ be a homomorphism and $f_i \colon Q \times Q \to A$ be a 2-cocycle in $Z^2(Q, A, \theta_i)$. Then* EXTENSION DATA PSEUDO-CONGRUENCE *for (θ_1, f_1) and (θ_2, f_2) Karp-reduces to*

GpI.

5. Preliminaries for the algorithms. Some general notation is described at the beginning of section 2.

Further notation and some group-theoretic facts. Given a finite set Ω, $\mathrm{Sym}(\Omega)$ denotes the symmetric group consisting of all permutations of Ω. A permutation group acting on Ω is a subgroup of $\mathrm{Sym}(\Omega)$. Given $\pi \in \mathrm{Sym}(\Omega)$ and $a \in \Omega$, the image of a under π is denoted by a^π. If $\Omega = [n]$, $n \in \mathbb{N}$, we use S_n to denote $\mathrm{Sym}(\Omega)$, and $A_n \leq S_n$ consists of permutations of even signs. For a vector space V over a field $\mathbb{F}$, the general linear group $\mathrm{GL}(V)$ consists of all nonsingular linear transformations of V. If $V = \mathbb{F}_q^n$, q is a prime power, we may write $\mathrm{GL}(n, q)$ for $\mathrm{GL}(V)$.

By the Fundamental Theorem of Finite Abelian Groups, a finite abelian group is isomorphic to a direct product of cyclic groups of prime power orders. Formally, let A be an abelian group; then there exists a direct product decomposition of A as $A = \langle e_1 \rangle \times \langle e_2 \rangle \times \cdots \times \langle e_n \rangle$, where $e_i \in A$ has order $p_i^{k_i}$ such that $p_1 \leq p_2 \leq \cdots \leq p_n$, and if $p_i = p_{i+1}$, then $k_i \leq k_{i+1}$ for all i. This decomposition is called the primary decomposition of A, and the tuple $(e_1, \ldots, e_n)$ forms a basis of A. The elementary abelian groups are those groups of the form $\mathbb{Z}_p^n$ for some prime p and any n. Note that $\mathrm{Aut}(\mathbb{Z}_p^n) \cong \mathrm{GL}(n, p)$.

A group G is *simple* if $|G| > 1$ and G has no proper nontrivial normal subgroups. The celebrated Classification of Finite Simple Groups lists all finite simple groups explicitly [30]. The only abelian simple groups are the cyclic groups of prime order. We use the following fact, which (currently) depends on the Classification for its proof.

FACT 5.1 (see [81, 3]). *Every nonabelian simple group can be generated by two elements.*

Let T be a nonabelian simple group; it is easily shown that $\mathrm{Aut}(T^k) \cong \mathrm{Aut}(T) \wr S_k$ where $\wr$ denotes the wreath product. If a group G is a direct product of nonabelian simple groups, then this direct product decomposition is unique, not just up to isomorphism: If $G = T_1 \times \cdots \times T_k = S_1 \times \cdots \times S_\ell$, T_i, S_j nonabelian simple, then $k = \ell$ and there exists $\sigma \in S_k$ for all $i \in [k]$, $T_i = S_{i^\sigma}$ as subsets of G.

A group G is *perfect* if $G = [G, G]$, *centerless* if $Z(G) = 1$, and *(directly) indecomposable* if it cannot be written as a nontrivial direct product. Nonabelian simple groups are perfect, centerless, and indecomposable, and several properties of nonabelian simple groups generalize to this class. In particular, if T is a perfect, centerless, indecomposable group, then $\mathrm{Aut}(T^k) \cong \mathrm{Aut}(T) \wr S_k$. Furthermore, if $G = T_1 \times \cdots \times T_k$ where each T_i is perfect, centerless, and indecomposable, then the direct factors of G are uniquely determined as subsets of G (not just up to isomorphism), as in the case of nonabelian simple groups.

(However, perfect, centerless, indecomposable groups can require many more than two generators and correspondingly can have quite large automorphism groups—for example, the group $\{(\pi_1, \ldots, \pi_k) | \pi_i \in S_5, \prod_{i \in [k]} \pi_i \in A_5\}$.)

The following observation is often useful to upgrade algorithms from merely deciding isomorphism to computing the entire coset of isomorphisms. The reader should have in mind that this observation will typically be applied in settings where the groups X, Y, Z are automorphism groups of other groups.

OBSERVATION 5.2. *Let X be a subgroup of an extension $Y \overset{\iota}{\hookrightarrow} \overline{X} \overset{\pi}{\twoheadrightarrow} Z$. If $\mathcal{Y} \subseteq \iota(Y)$ generates $\iota(Y) \cap X$, and $\mathcal{Z} \subseteq Z$ generates $\pi(X)$, and for each $z \in \mathcal{Z}$, x_z is such that $x_z \in X$ and $\pi(x_z) = z$, then $\mathcal{Y} \cup \{x_z \in X : z \in \mathcal{Z}\}$ generates X.*

Proof. Given $x \in X$, first write $z = \pi(x)$ as a word in the generators $\mathcal{Z}$, say

$z = z_1 \cdots z_\ell$, with each $z_i \in \mathcal{Z}$. Then $y = x \cdot (x_{z_1} \cdots x_{z_\ell})^{-1}$ is in $X \cap \mathrm{Ker}(\pi) = X \cap \iota(Y)$. Now write y as a word in $\mathcal{Y}$. □

In Theorem 7.1(1), we will need the following deep result bounding the size of the second cohomology of simple groups:

THEOREM 5.3 (see [43, Thm. B]). *For any finite (quasi-[7])simple group G, any field $\mathbb{F}$, and any $\mathbb{F}$-representation $\theta\colon G \to \mathrm{GL}(V)$ (V an $\mathbb{F}$-vector space), we have* $\dim_{\mathbb{F}} H^2(G, V, \theta) \leq 17.5 \dim_{\mathbb{F}} V$.

Useful algorithms. We shall need some known algorithmic results. Recall that in permutation group algorithms (see [62, 79]), a coset $P\sigma \subseteq S_n$ is represented by a set of generators for $P \leq S_n$ and a coset representative σ. A particularly relevant problem on permutation groups is the COSET INTERSECTION problem: Given two cosets of subgroups of $\mathrm{Sym}(A)$, find their intersection. GRAPHI can be Karp-reduced to COSET INTERSECTION [61]. The COSET INTERSECTION problem for permutation groups of degree n can be solved in quasi-polynomial ($\exp((\log n)^{O(1)})$) time [8], while a relatively easier singly exponential ($\exp(O(n))$) algorithm has been obtained by Luks [63]. Algorithms over finite-dimensional algebras have been considered in, e. g., [28, 22, 52]. In particular, over finite fields, polynomial-time algorithms for MODULE ISOMORPHISM and MODULE CYCLICITY were first shown in [28].

A linear code of length n is a linear subspace $V \leq \mathbb{F}^n$, represented by a $d \times n$ matrix where $d = \dim(V)$, and the rows form a linear basis of V. S_n acts on a linear code by permuting the coordinates (that is, the columns of the matrices). Two linear codes $V, U \leq \mathbb{F}^n$ are equivalent if there exists a permutation $\sigma \in S_n$ such that $V^\sigma = U$ as linear subspaces. Such a σ is called an equivalence between V and U, and the set of all equivalences, denoted by $\mathrm{CodeEq}(V, U)$ is either empty or a coset in S_n. This problem is GRAPHI-hard in general [69], while Babai presents a singly exponential time algorithm.

THEOREM 5.4 (see Babai [7]; cf. [11]). *The set of equivalences of two linear codes of length n (over any field) given by generator matrices can be found in $(2 + o(1))^n$ time, assuming field operations at unit cost.*

We will also need the following results of Babai et al. [11].

THEOREM 5.5 (see [11, Thm. 1.1]). *All isomorphisms between two semisimple groups Q_1 and Q_2 of order n can be listed in time $n^{c \log \log n + O(1)}$, where $c = 1/\log(60) \approx 0.16929$.*

It is also noted in [11] that there exist semisimple groups G of order n with $|\mathrm{Aut}(G)| \geq n^{c \log \log n}$, namely $G = A_5^k$. Hence, for listing all isomorphisms this result is essentially optimal.

The number of minimal normal subgroups of any semisimple group of order n is at most $O(\log n)$. If it happens to be $O(\log n / \log \log n)$, they show the following.

THEOREM 5.6 (see [11, Cor. 4.4]). *Suppose Q_1 and Q_2 are semisimple groups of order n with at most $O(\log n / \log \log n)$ minimal normal subgroups. Then all isomorphisms between Q_1 and Q_2 can be listed in polynomial time.*

In section 6 we extend both of these results to algorithms for isomorphism of groups with central radicals, with the same time bounds as above.

---

[7]G is *quasi-simple* if it is perfect, and $G/Z(G)$ is simple. Examples include $\mathrm{SL}_n(\mathbb{F}_q)$ and $\mathrm{GL}_n(\mathbb{F}_q)$. We will only need the case of nonabelian simple groups, but the theorem holds in the generality of quasi-simple groups, which might be useful in the future.

To find the coset of isomorphisms efficiently in Theorem 7.1, beyond just deciding isomorphism or finding a single isomorphism, we need the following result. Recall that a $d \times d$ *matrix algebra* over a field $\mathbb{F}$ is a linear subspace of $d \times d$ matrices over $\mathbb{F}$ that is also closed under multiplication of matrices. The *unit group* of a (matrix) algebra A is the group consisting of all invertible elements of A (the group operation being the same as the product in the algebra).

THEOREM 5.7 (see [24, Thm. 2.3] and [74]). *Given a linear basis of a $d \times d$ matrix algebra A over a finite field $\mathbb{F}_q$, generators of the unit group of A can be computed in time polynomial in d and q.*

Although the preceding theorem from the original papers is stated as a Las Vegas randomized algorithm that runs in time polynomial in d and $\log q$, they only need randomization to get an algorithm whose running time depends polynomially on $\log q$ rather than q. Essentially the same algorithm works deterministically in time polynomial in d and q.

We also mention a useful result for groups in the Cayley table model by Kayal and Nezhmetdinov [55], though it is not strictly required in the following. They show that decomposing a group G into indecomposable direct factors can be done in polynomial time. Even in the stronger setting of permutation groups given by generators, Wilson [91] showed that this task can be performed in polynomial time.

6. When enumerating $\mathrm{Aut}(Q)$ is allowed. Our main results in this section are $n^{O(\log \log n)}$-time algorithms to test isomorphism of (1) groups with central radicals (Corollary 6.2) and (2) groups with elementary abelian radicals that need not be central (Corollary 6.12). In case (1), if the radical is furthermore *elementary* abelian, we can compute the coset of isomorphisms in the same time bound. These results follow from our more general Theorems 6.1 and 6.11, respectively, and a theorem on semisimple groups from [11] (reproduced above as Theorem 5.5).

6.1. For central extensions of general abelian groups (Theorem A). We first consider the case when both extensions G and H are central.

THEOREM 6.1. *Let $\mathcal{S}$ be a polynomial-time-computable characteristic subgroup function. For two groups G, H of order n, if $\mathcal{S}(G) \leq Z(G)$ and $\mathrm{Aut}(G/\mathcal{S}(G))$ can be listed in time $t(n)$, then isomorphism of G and H can be decided in time $t(n)n^{O(1)}$.*

If, furthermore, $\mathcal{S}(G)$ is elementary abelian, then the coset of isomorphisms can be found in the same time bound.

Before proving Theorem 6.1, let us see how it is applied to groups with central radicals. Combining Theorem 6.1 with Theorem 5.5 and, respectively, Theorem 5.6, we have our first two main results.

COROLLARY 6.2 (Theorem A). *Isomorphism of central-radical groups of order n can be decided in time $n^{c \log \log n + O(1)}$ for $c = 1/\log_2(60) \approx 0.169$. If, furthermore, the center is elementary abelian, the coset of isomorphisms can be found in the same time bound.*

COROLLARY 6.3. *Let G and H be central-radical groups of order n. If $G/\mathrm{Rad}(G)$ has $O(\log n/\log \log n)$ minimal normal subgroups, isomorphism between G and H can be decided in $\mathrm{poly}(n)$ time. If, furthermore, the center is elementary abelian, the coset of isomorphisms can be found in the same time bound.*

We give two different proofs of Theorem 6.1. The first proof has the advantage of dealing with cohomology classes in a very direct manner and working for arbitrary

central $\mathcal{S}(G)$ but has the disadvantage of not making it obvious how to compute the full coset of isomorphisms. For clarity, we first give this "direct" proof in the elementary abelian case, then in subsection 6.1.1 we give it for the general abelian case, and finally in subsection 6.1.2 we go back to the elementary abelian case to show how to compute the coset of isomorphisms.

Let us consider how to work with 2-cohomology classes in algorithms. Let G be a central extension of $A = \mathbb{Z}_p^k$ by Q (thinking of $A = \mathcal{S}(G)$ and $Q = G/\mathcal{S}(G)$). As the action is trivial in central extensions, we drop it from the notation, as in $Z^2(Q, A)$, $B^2(Q, A)$, and $H^2(Q, A)$. By choosing an arbitrary section, we get a 2-cocycle $f : Q \times Q \to A$. Let $e_1, \ldots, e_k$ be the standard basis of $\mathbb{Z}_p^k$. We may view f as a $k \times |Q|^2$-size $\mathbb{Z}_p$-matrix, which we denote by M_f. The rows are indexed by the set $[k]$, and the columns are indexed by $Q \times Q$. For $i \in [k]$ and $(q, q') \in Q \times Q$, $M_f[i, (q, q')]$ is the ith coordinate of $f(q, q')$ relative to the basis $\{e_1, \ldots, e_k\}$. Note that the actions of $\mathrm{Aut}(A)$ and $\mathrm{Aut}(Q)$ commute.

Under the above identification, the set $C^2(Q, A)$ of 2-cochains is identified with the set of all $k \times |Q|^2$ matrices over $\mathbb{Z}_p$. Then $Z^2(Q, A)$ is not just a subgroup (under matrix addition) but also a $\mathbb{Z}_p$-linear subspace of $C^2(Q, A)$, and similarly $B^2(Q, A)$ is a $\mathbb{Z}_p$-linear subspace of $Z^2(Q, A)$. $\mathrm{Aut}(A) \cong \mathrm{GL}(k, p)$ acts on $C^2(Q, A)$ by left multiplication, and $\mathrm{Aut}(Q)$ acts on $C^2(Q, A)$ by permuting the columns according to the diagonal action of $\mathrm{Aut}(Q)$ on $Q \times Q$.

PROPOSITION 6.4. *A basis of $B^2(Q, \mathbb{Z}_p)$ can be computed in time $O(|Q|^3(\log|Q| + \log p))$.*

In the larger context of GPI, note that the running time here is $O(|G|^3 \log |G|)$.

Proof. For $q \in Q$, $q \neq \mathrm{id}$, let $u_q : Q \to \mathbb{Z}_p$ be $u_q(q') = \delta(q, q')$, where δ is the Kronecker delta. Let $f_q : Q \times Q \to \mathbb{Z}_p$ be the 2-coboundary based on u_q. $V := \{f_q \mid q \in Q\}$ then forms a basis of $B^2(Q, \mathbb{Z}_p)$. There are $|Q|$ basis elements, each of which is constructed by computing its $|Q|^2$ values; each value can be computed by a constant number of additions in $\mathbb{Z}_p$ (taking $O(\log p)$ steps) and one table lookup to compute a single product in Q (taking $O(\log|Q|)$ steps). $\square$

As we identified $C^2(Q, A)$ as $k \times |Q|^2$ matrices over $\mathbb{Z}_p$, let $E_{i,j}$, $i \in [k]$, $j \in Q \times Q$, be the $k \times |Q|^2$ matrix with a 1 in the (i, j) position and 0 everywhere else. Then $\{E_{i,j} \mid i \in [k], j \in Q \times Q\}$ is a basis of $C^2(Q, A)$. Let U_i be the subspace of $C^2(Q, A)$, spanned by $\{E_{i,j} \mid j \in Q \times Q\}$, corresponding to matrices whose only nonzero entries are in the ith row. Then $C^2(Q, A) = \oplus_{i \in [k]} U_i$. The following proposition says that not only does $C^2(Q, A)$ split as a direct sum over the rows, but $B^2(Q, A)$ does as well. It follows directly from the fact that the condition to be a 2-coboundary in $B^2(Q, A)$ only depends on the columns $(Q \times Q)$ and not on the rows $([k])$.

PROPOSITION 6.5. *Let V be the basis of $B^2(Q, \mathbb{Z}_p)$ constructed in the proof of Proposition 6.4, and let $V_i \leq C^2(Q, A)$ be a copy of V in U_i. Then $\bigsqcup_{i \in [k]} V_i$ (disjoint union) is a basis of $B^2(Q, \mathbb{Z}_p^k)$.*

Given two 2-cocycles f_1 and f_2, let M_i be the matrix representation of f_i, $i = 1, 2$, and $R_i \subseteq \mathbb{Z}_p^{|Q|^2}$ be the set of rows in M_i. Recall that $\alpha \in \mathrm{GL}(k, p)$ acts on the rows of M_i.

PROPOSITION 6.6. *With notation as above, there exists $\alpha \in \mathrm{GL}(k, p)$ such that f_1 and f_2^α are cohomologous if and only if $\langle R_1, B^2(Q, \mathbb{Z}_p)\rangle = \langle R_2, B^2(Q, \mathbb{Z}_p)\rangle$, where $\langle \cdot \rangle$ denotes the $\mathbb{Z}_p$-linear span.*

Proof. Let $r_{i,j} \in \mathbb{Z}_p^{|Q|^2}$ be the jth row in M_i, $j \in [k]$, $i = 1, 2$. Let B denote $B^2(Q, \mathbb{Z}_p)$. Note that Proposition 6.5 says that $B^2(Q, \mathbb{Z}_p^k) = B \oplus B \oplus \cdots \oplus B$ (k summands).

($\Rightarrow$) f_1 and f_2^α are cohomologous if and only if $f_1 - f_2^\alpha \in B^2(Q, \mathbb{Z}_p^k)$. Let $r_{2,j}^\alpha$ be the jth row in the matrix representation of f_2^α. By Proposition 6.5, for every $i \in [k]$, $r_{1,i} - r_{2,i}^\alpha \in \langle V_i \rangle = B^2(Q, \mathbb{Z}_p) = B$. That is, $r_{1,i} \in \langle R_2, B \rangle$ as $r_{2,j}^\alpha \in \langle R_2 \rangle$ (note that the linear span of R_2, i.e., the rowspan of M_2, is left unchanged by the action of α). Similarly we have $r_{2,i} \in \langle R_1, B \rangle$ for all $i \in [k]$. This shows $\langle R_1, B \rangle = \langle R_2, B \rangle$.

($\Leftarrow$) For $\alpha \in \mathrm{GL}(k, p)$, again let $r_{2,j}^\alpha$ be the jth row of f_2^α. Given $\langle R_1, B \rangle = \langle R_2, B \rangle$, we have that $\langle R_1, B \rangle/B$ and $\langle R_2, B \rangle/B$ are the same as subspaces of $\mathbb{Z}_p^{|Q|^2}/B$. That means that we can choose $\alpha \in \mathrm{GL}(k, p)$ such that $r_{1,i} + B = r_{2,i}^\alpha + B$ for all $i \in [k]$. This gives $f_1 - f_2^\alpha \in B^2(Q, A)$. □

Proof of Theorem 6.1 *(decision version only) when* $\mathcal{S}(G)$ *is elementary abelian.* We list $\mathrm{Aut}(Q)$ in time $t(n)$. For $i = 1, 2$, choose an arbitrary section of Q in G_i to get a 2-cocycle f_i. By the abelian version of the main lemma (Lemma 2.3), it is necessary and sufficient to test whether there exists an $(\alpha, \beta) \in \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ such that f_1 and $f_2^{(\alpha,\beta)}$ are cohomologous.

For each $\beta \in \mathrm{Aut}(Q)$ we get $f_2' = f_2^{(\mathrm{id},\beta)}$. We first use Proposition 6.4 to get a basis V of $B^2(Q, \mathbb{Z}_p)$. Let M_1 be the matrix representation of f_1, and M_2 for f_2'. We now need to determine whether there exists α such that f_1 and f_2' are cohomologous. By Proposition 6.6 it is enough to decide whether the linear span of the rows of f_1 with V and the linear span of the rows of f_2' with V, are the same. This is a standard task in linear algebra and can be determined in time polynomial in $|Q|$ and $\dim_{\mathbb{Z}_p} |A| = k$. Lemma 2.3 implies that $G_1 \cong G_2$ if and only if the preceding test succeeds for some $\beta \in \mathrm{Aut}(Q)$. □

6.1.1. From elementary abelian to general abelian.

The proof here follows the same steps as in the elementary abelian case. As each abelian group A is the direct product of its Sylow p-subgroups A_p, we essentially treat the case of a single Sylow p-subgroup, that is, when A is an abelian p-group $\mathbb{Z}_{p^{\mu_1}} \times \cdots \times \mathbb{Z}_{p^{\mu_k}}$ (not necessarily elementary). We begin by extending Propositions 6.4 to 6.6 to the case of abelian p-groups.

As such groups are no longer just vector spaces over $\mathbb{Z}_p$, we must speak of subgroups of A rather than subspaces, and generating sets rather than $\mathbb{Z}_p$-bases. To emphasize the similarities, we use the terminology "$\mathbb{Z}$-basis" for "irredundant generating set." We use similar terminology for $C^2(Q, A)$, $Z^2(Q, A)$, and $B^2(Q, A)$. We represent a 2-cochain $f: Q \times Q \to A$ by a $k \times |Q|^2$ integer matrix, where we consider the entries in the ith row modulo p^{μ_i}, that is, as elements of $\mathbb{Z}_{p^{\mu_i}}$. As before, we use U_i to denote the subgroup of $C^2(Q, A)$ consisting of matrices whose only nonzero entries are in the ith row (in particular, $U_i \cong \mathbb{Z}_{p^{\mu_i}}^{|Q|^2}$).

For these first two propositions, the proofs are the same as the analogous propositions above for elementary abelian A.

PROPOSITION 6.7. *For any* $\mu \geq 1$, *a* $\mathbb{Z}$-*basis of* $B^2(Q, \mathbb{Z}_{p^\mu})$ *can be computed in time* $O(|Q|^3(\log |Q| + \mu \log p))$.

PROPOSITION 6.8. *Let* $V^{(\mu)}$ *denote the* $\mathbb{Z}$-*basis of* $B^2(Q, \mathbb{Z}_{p^\mu})$ *constructed in the proof of Proposition* 6.7, *and let* $V_i^{(\mu_i)} \leq C^2(Q, A)$ *be a copy of* $V^{(\mu_i)}$ *in* U_i. *Then* $\sqcup_{i \in [k]} V_i^{(\mu_i)}$ *(disjoint union) is a* $\mathbb{Z}$-*basis of* $B^2(Q, \mathbb{Z}_{p^{\mu_1}} \times \cdots \times \mathbb{Z}_{p^{\mu_k}})$.

Before giving the analogue of Proposition 6.6 for general abelian A, we recall the structure of $\mathrm{Aut}(A)$ (see, e.g., the exposition in [46]); it is only slightly more complicated than the fact that $\mathrm{Aut}(\mathbb{Z}_p^k) = \mathrm{GL}(k,p)$. First, if A_p is the p-Sylow subgroup of A, then $\mathrm{Aut}(A) = \mathrm{Aut}(A_{p_1} \times \cdots \times A_{p_d}) = \mathrm{Aut}(A_{p_1}) \times \cdots \times \mathrm{Aut}(A_{p_d})$ where $p_1, \ldots, p_d$ are the distinct primes dividing $|A|$. So we reduce to the case where A is an abelian p-group $\mathbb{Z}_{p^{\mu_1}} \times \cdots \times \mathbb{Z}_{p^{\mu_k}}$ with $1 \le \mu_1 \le \mu_2 \le \cdots \le \mu_k$. Think of elements of A as integer column vectors of length k, where the ith entry is considered modulo p^{μ_i}. As in the elementary abelian case (where $\mu_1 = \cdots = \mu_k = 1$), an automorphism may replace each entry with a $\mathbb{Z}$-linear combination of the entries, as follows. For $i < j$, the ith coordinate can contribute to the jth coordinate by multiplying by $p^{\mu_j - \mu_i}$—in other words, by using the unique inclusion $\mathbb{Z}_{p^{\mu_i}} \hookrightarrow \mathbb{Z}_{p^{\mu_j}}$. In the opposite direction, the jth coordinate can contribute to the ith coordinate by taking the jth coordinate modulo p^{μ_i}—in other words, using the natural surjection $\mathbb{Z}_{p^{\mu_j}} \twoheadrightarrow \mathbb{Z}_{p^{\mu_i}}$. (Note that when $\mu_i = \mu_j$, these two operations are the same, corresponding to the identity map on $\mathbb{Z}_{p^{\mu_i}}$.)

More symbolically, we may consider each element of $\mathrm{Aut}(A)$ as an integer $k \times k$ matrix α such that (1) for $i > j$, $p^{\mu_j - \mu_i}$ divides the (i,j) entry, (2) the entries in row i are considered modulo p^{μ_i}, and (3) α is invertible when taken modulo p.

Finally, consider (row) subgroups $R \le \mathbb{Z}_{p^{\mu_i}}^{|Q|^2}$. In accord with the above description of the automorphisms of A, for $\mu < \mu_i$ let $R^{(\mu)}$ denote the subgroup of $\mathbb{Z}_{p^\mu}^{|Q|^2}$ that is given by taking R modulo p^μ; for $\mu > \mu_i$, let $R^{(\mu)}$ denote the subgroup of $\mathbb{Z}_{p^\mu}^{|Q|^2}$ that is given by multiplying every element of R by $p^{\mu - \mu_i}$. For any prime q, let $R^{(q,\mu)}$ denote $R^{(\mu)}$ if $q = p$ and the trivial subgroup 0 otherwise.

Now, return to A being an arbitrary abelian group. For a 2-cochain $f_1 \in C^2(Q, A)$ with corresponding $k \times |Q|^2$ matrix M with ith row $R_{1,i} \le \mathbb{Z}_{p_i^{\mu_i}}^{|Q|^2}$, let $R_1^{(p,\mu)}$ denote the subgroup of $\mathbb{Z}_{p^\mu}^{|Q|^2}$ generated by all the $R_{1,i}^{(p,\mu)}$; we write $R_1^{(p,\mu)} = \langle R_{1,1}^{(p,\mu)}, \ldots, R_{1,k}^{(p,\mu)} \rangle$.

PROPOSITION 6.9. *Let $A = \mathbb{Z}_{p_1^{\mu_1}} \times \cdots \times \mathbb{Z}_{p_k^{\mu_k}}$ be an arbitrary abelian group (the p_i are primes, not necessarily distinct). With other notation as above, there exists $\alpha \in \mathrm{Aut}(A)$ such that f_1 and f_2^α are cohomologous if and only if $\langle R_1^{(p_i,\mu_i)}, B^2(Q, \mathbb{Z}_{p_i^{\mu_i}}) \rangle = \langle R_2^{(p_i,\mu_i)}, B^2(Q, \mathbb{Z}_{p_i^{\mu_i}}) \rangle$ for each $1 \le i \le k$, where $\langle \cdot \rangle$ denotes the $\mathbb{Z}$-span (=group generated by).*

Proof. Let $r_{i,j} \in \mathbb{Z}_{p_j^{\mu_j}}^{|Q|^2}$ be the jth row in M_i, $j \in [k]$, $i = 1,2$. Let $B^{(p,\mu)}$ denote $B^2(Q, \mathbb{Z}_{p^\mu})$. Note that Proposition 6.5 says that $B^2(Q, \mathbb{Z}_{p_1^{\mu_1}} \times \cdots \times \mathbb{Z}_{p_k^{\mu_k}}) = B^{(p_1,\mu_1)} \oplus \cdots \oplus B^{(p_k,\mu_k)}$.

($\Rightarrow$) f_1 and f_2^α are cohomologous if and only if $f_1 - f_2^\alpha \in B^2(Q, A)$. Let $r_{2,j}^\alpha$ be the jth row in the matrix representation of f_2^α. By Proposition 6.7, for every $i \in [k]$, $r_{1,i} - r_{2,i}^\alpha \in \langle V_i^{(p,\mu_i)} \rangle = B^2(Q, \mathbb{Z}_{p_i^{\mu_i}}) = B^{(p_i,\mu_i)}$. That is, $r_{1,i} \in \langle R_2^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle$ as $r_{2,j}^\alpha \in \langle R_2^{(p_i,\mu_i)} \rangle$ (note that the subgroup generated by of $R_2^{(p_i,\mu_i)}$ is, by definition, left unchanged by the action of α). Similarly we have $r_{2,i} \in \langle R_1^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle$ for all $i \in [k]$. This shows $\langle R_1^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle = \langle R_2^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle$ for each i.

($\Leftarrow$) For $\alpha \in \mathrm{Aut}(A)$, again let $r_{2,j}^\alpha$ be the jth row of f_2^α. Given $\langle R_1^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle = \langle R_2^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle$ for each i, we have

$$\langle R_1^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle / B^{(p_i,\mu_i)} = \langle R_2^{(p_i,\mu_i)}, B^{(p_i,\mu_i)} \rangle / B^{(p_i,\mu_i)}$$

as subgroups of $\mathbb{Z}_{p_i^{\mu_i}}^{|Q|^2}/B^{(p_i,\mu_i)}$. That means that we can choose $\alpha \in \mathrm{Aut}(A)$ such that $r_{1,i} + B^{(p_i,\mu_i)} = r_{2,i}^{\alpha} + B^{(p_i,\mu_i)}$ for all $i \in [k]$. This gives $f_1 - f_2^{\alpha} \in B^2(Q,A)$.   □

Finally, we come to the proof of Theorem 6.1 for general abelian A.

Proof of Theorem 6.1 *(decision version only) for general abelian* $\mathcal{S}(G)$. Same as for the elementary abelian case, but using Propositions 6.8 and 6.9 instead of Propositions 6.4 and 6.6, respectively. Checking the condition of Proposition 6.9 amounts to solving a system of equations over the abelian group A ("linear algebra over A"), which can be done in polynomial time (see, e.g., [39]).   □

6.1.2. Computing the coset of isomorphisms in the elementary abelian case. Here we give an alternative proof of Theorem 6.1 in the elementary abelian case, which also provides the structure needed to compute the full coset of isomorphisms. The proof follows the same general lines as the proof above, except instead of including a basis of $B^2(Q,A)$ among the rows of an extended matrix $\widetilde{M_{f_j}}$, it uses the following $\mathrm{Aut}(A)$-invariant projection from $C^2(Q,A)$ to $C^2(Q,A)/B^2(Q,A)$. One might think that this approach could be generalized to the general abelian case following the idea in the previous section, thereby computing the coset of isomorphisms in that case as well; see the end of this section for a discussion of the difficulties in using that approach.

PROPOSITION 6.10. *For* $A = \mathbb{Z}_p^k$ *and a group* Q, *let* $n = |A| \cdot |Q|$. *In time* $O(n^2 \log n)$ *one can compute an* $\mathrm{Aut}(A)$-*invariant complement* W *of* $B^2(Q,A)$ *in* $C^2(Q,A)$, *and a* $\mathbb{Z}_p$-*linear projection* π *from* $C^2(Q,A)$ *to* W *that commutes with every* $\alpha \in \mathrm{Aut}(A)$.

Proof. Let V_0 be the basis of $B^2(Q,\mathbb{Z}_p)$ from Proposition 6.4, and let W_0 be a linear complement to V_0 in $C^2(Q,\mathbb{Z}_p)$ (which we can think of as the space of row vectors of length $|Q|^2$). For each row index $i \in [k]$, let R_i denote the subgroup of $C^2(Q,A)$ consisting of those matrices whose only nonzero entries are in the ith row, let B_i be the copy of $B^2(Q,\mathbb{Z}_p)$ in R_i, let W_i be the copy of W in R_i, and let $\pi_i\colon R_i \to W_i$ be the projection of R_i to W_i along B_i. (Here we are identifying each R_i with $C^2(Q,\mathbb{Z}_p)$ as in Proposition 6.5.) Collectively, the π_i's define a projection π to $\oplus_i W_i$ along $\oplus_i B_i$ in $\oplus_i R_i = C^2(Q,A)$.

This π makes it easy to identify the 2-cohomology class of $f \in Z^2(Q,A) \leq C^2(Q,A)$: $f,g \in Z^2(Q,A)$ are cohomologous if and only if $\pi(f) = \pi(g)$. Furthermore, since we have chosen π to be "the same" on each row, it follows that for any $\alpha \in \mathrm{Aut}(A)$, $\alpha\pi = \pi\alpha$—an important fact we shall need later. Also, the above procedure involves only standard linear algebra tasks, and such π can be constructed in time $O(n^2 \log n)$.   □

To summarize, each $\alpha \in \mathrm{Aut}(A)$, $\beta \in \mathrm{Aut}(Q)$, and the map $\pi\colon C^2(Q,A) \to C^2(Q,A)/B^2(Q,A)$ just introduced can be viewed as a linear map on $C^2(Q,A)$, with α and β nonsingular. We then note that (1) α and β commute, and (2) α and π commute.

Proof of Theorem 6.1 *for elementary abelian* $\mathcal{S}(G)$, *coset of isomorphisms.* We list $\mathrm{Aut}(Q)$ in time $t(n)$. For $i = 1, 2$, choose an arbitrary section of Q in G_i to get a 2-cocycle f_i. Use Proposition 6.10 to get the projection $\pi : C^2(Q,A) \to W$ for some complement W of $B^2(Q,A)$ such that W is invariant under α and $\alpha\pi = \pi\alpha$ for every $\alpha \in \mathrm{Aut}(A)$. By Lemma 2.3, it is necessary and sufficient to test whether there exists an $(\alpha,\beta) \in \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ such that $\pi(f_1) = \pi(f_2^{(\alpha,\beta)})$. As every α

commutes with both π and β, this condition is equivalent to $\pi(f_1) = (\pi(f_2^{(\mathrm{id},\beta)}))^{(\alpha,\mathrm{id})}$. In other words, we may leave α unspecified until the final step.

For each $\beta \in \mathrm{Aut}(Q)$, compute $f_1' = \pi(f_1)$ and $f_2' = \pi(f_2^{(\mathrm{id},\beta)})$. Note that f_i' are in $W \leq C^2(Q, A)$. The goal is then to find $\alpha \in \mathrm{Aut}(A)$ such that $f_1' = \alpha^{-1}(f_2')$. $f_1' = \alpha^{-1}(f_2')$ for some $\alpha \in \mathrm{Aut}(A)$ if and only if the row spans of $M_{f_1'}$ and $M_{f_2'}$ are the same in $\mathbb{Z}_p^{|Q|^2}$. The latter task is standard in linear algebra and can be checked in time $O(|Q|^6 \log p)$. Lemma 2.3 implies that $G_1 \cong G_2$ if and only if the preceding test succeeds for some $\beta \in \mathrm{Aut}(Q)$.

Computing the coset of isomorphisms. (This is essentially the same procedure as in Smith's thesis [80]; we review the procedure here in detail, in order to give a rigorous analysis of its running time.) Note that $\mathrm{Aut}(G)$ maps into $\mathrm{Aut}(A) \times \mathrm{Aut}(Q)$; let ρ denote this homomorphism. We apply Observation 5.2 twice: once to split the computation into computing generators of $\mathrm{Im}(\rho) \leq \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ and $\mathrm{Ker}(\rho)$, and then a second time to split the computation of $\mathrm{Im}(\rho)$ into computing its projections onto $\mathrm{Aut}(A)$ and $\mathrm{Aut}(Q)$, respectively. We handle this latter step first.

Let $\mathrm{Aut}(A)$ play the role of Y in Observation 5.2 and $\mathrm{Aut}(Q)$ play the role of Z. In our algorithm, we may in fact take $\mathcal{Z}$ to be the entirety of the projection of $\mathrm{Aut}(G)$ into $\mathrm{Aut}(Q)$, since we enumerate over all automorphisms of Q. To use Observation 5.2, for each automorphism of Q that admits a compatible automorphism of A, we need one such automorphism of A, which are easily found as solutions to the linear algebra, as described above.

We then also need a generating set of the automorphisms of A that preserve the cohomology class $f: Q \times Q \to A$, that is, those elements of $\mathrm{Aut}(A)$ that send f to itself modulo $B^2(Q, A)$. Equivalently, we want the stabilizer of the matrix $M_{\pi(f)}$ in $\mathrm{Aut}(A) \cong \mathrm{GL}(k, p)$. Since $\mathrm{Aut}(A)$ acts on each column of a 2-cochain independently, this stabilizer is the same as the pointwise stabilizer of the *set* of columns of $M_{\pi(f)}$. If the columns of $M_{\pi(f)}$ span A, then this stabilizer is trivial, and we are done. More generally, let $B \leq A$ be the $\mathbb{Z}_p$-linear span of the columns of $M_{\pi(f)}$ (although we will not need it here, Lemma 7.5 may help the reader's intuition, which says that in this case, there is a complementary subgroup $B' \leq A$ such that $A \cong B \oplus B'$ and $G \cong (G/B') \times B'$). Choose a basis $\{e_1, \ldots, e_k\}$ of A such that a prefix of this basis, say, $\{e_1, \ldots, e_{\dim B}\}$, is a basis of B. In this basis, the stabilizer of $M_{\pi(f)}$ in $\mathrm{GL}(k, p)$ then consists of all block matrices of the form

$$\begin{pmatrix} \mathrm{id}_B & * \\ 0 & \eta \end{pmatrix},$$

where $\eta \in \mathrm{GL}(k - \dim B, p)$ and "$*$" indicates any $(\dim B) \times (\dim A - \dim B)$ matrix. A generating set for this subgroup is easily written down directly, taking advantage of a standard generating set for GL (e.g., elementary matrices). Observation 5.2 then gives us a generating set for $\mathrm{Im}(\rho) \leq \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$.

Finally, we apply Observation 5.2 again, now with $\mathrm{Ker}(\rho)$ in the role of Y and $\mathrm{Im}(\rho)$ in the role of Z. We have already computed a generating set $\mathcal{Z}$. For each such generator, we must compute a lift of it in $\mathrm{Aut}(G)$, and then we also need to compute generators for $\mathrm{Ker}(\rho)$. The lifts of $\mathrm{Im}(\rho)$ to $\mathrm{Aut}(G)$ are easily computed once we have chosen our section $s: Q \to G$: Given $(\alpha, \beta) \in \mathrm{Im}(\rho) \leq \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$, every element of G is uniquely represented as $as(q)$ for some $a \in A, q \in Q$, and then the corresponding automorphism of G acts by sending $as(q)$ to $\alpha(a)s(\beta(q))$. It is readily verified that this is an automorphism of G.

The kernel of ρ consists of those automorphisms of G that fix Q pointwise and fix

A pointwise. Thus, all they can do is move elements around within their respective cosets of A. That is, any automorphism in $\mathrm{Ker}(\rho)$ is determined by a map $\delta\colon Q \to A$ such that the automorphism is given by sending $as(q)$ to $a\delta(q)s(q)$. Since the automorphism fixes A pointwise, we must have $\delta(\mathrm{id}_Q) = \mathrm{id}_A$. To see the condition as a homomorphism, recall that $(as(q))(a's(q')) = a(s(q)a's(q)^{-1})f_s(q,q')s(qq')$, and thus we must have $\delta(q)\delta(q') = \delta(qq')$. These $|Q|^2$ equations are homogeneous linear equations in $|Q|\dim A$ variables over the field $\mathbb{Z}_p$; such linear equations can be solved in polynomial time, yielding a generating set for $\mathrm{Ker}(\rho)$. Putting it all together, this yields a generating set for $\mathrm{Aut}(G)$. $\qquad\square$

To see the difficulty in extending this proof to the general abelian case, we first mention an alternative viewpoint on the preceding proof: The set $C^2(Q, \mathbb{Z}_p^k)$, viewed as $k \times |Q|^2$ matrices, is acted on by $\mathrm{Aut}(A) \cong \mathrm{GL}(k, p)$ by left multiplication. Let $V \cong \mathbb{Z}_p^k$ be the defining representation of $\mathrm{GL}(k, p)$; then $C^2(Q, A)$ is isomorphic, as an $\mathrm{Aut}(A)$-module, to $V^{\oplus |Q|^2}$. Since V is irreducible, this module is semisimple, so every submodule is a direct summand, and thus every submodule is of the form $V^{\oplus \ell}$ for some ℓ. So $H^2(Q, A)$ is isomorphic to $V^{\oplus \ell}$ as an $\mathrm{Aut}(A)$-module, and finding the $\mathrm{Aut}(A)$-stabilizer of a point in $V^{\oplus \ell}$ amounts to finding the stabilizer of a point in V several times, which is easily done since V is the defining representation of $\mathrm{Aut}(A) \cong \mathrm{GL}(V)$.

The difficulty in extending this proof to the general abelian case is that $B^2(Q, A)$ need not, in general, be a direct summand of $Z^2(Q, A)$, nor of $C^2(Q, A)$; in the general case $C^2(Q, A) \cong A^{\oplus |Q|^2}$ as $\mathrm{Aut}(A)$-modules, but this module is no longer semisimple in general. For example, $C^2(A_5, \mathbb{Z}_4) \cong \mathbb{Z}_4^{60^2}$, which contains many copies of $\mathbb{Z}_2$ that are submodules but not direct summands. Note that $H^2(A_5, \mathbb{Z}_4) \cong \mathbb{Z}_2$ (use the well-known fact that $H_2(A_5, \mathbb{Z}) \cong \mathbb{Z}_2$, together with the Universal Coefficient Theorem, e.g., [45, Thm. 3.2], and the fact that $H_1(A_5, \mathbb{Z}) = 0$, since A_5 is perfect). So either $Z^2(A_5, \mathbb{Z}_4)$ is not a direct summand of $C^2(A_5, \mathbb{Z}_4)$ or $B^2(A_5, \mathbb{Z}_4)$ is not a direct summand of $Z^2(A_5, \mathbb{Z}_4)$.

Without a direct summand decomposition, the straightforward way to obtain an action of $\mathrm{Aut}(A)$ on $H^2(Q, A)$ is to realize $H^2(Q, A)$ explicitly as a quotient module. This can be done quite easily, using the fact that $H^2(Q, \mathbb{Z}_{p_1^{\mu_1}} \times \cdots \times \mathbb{Z}_{p_k^{\mu_k}}) \cong \bigoplus_{i \in [k]} H^2(Q, \mathbb{Z}_{p_i^{\mu_i}})$. This even yields an $\mathrm{Aut}(A)$-invariant projection $Z^2(Q, A) \to H^2(Q, A)$ (the best we can hope for in the absence of a direct sum decomposition). But then the difficulty becomes that we do not know the structure of this quotient module well enough. If it were just the action of $\mathrm{Aut}(A)$ on $A^{\oplus \ell}$ for some ℓ, then we would be very close to being done, as one can find the linear subspace of $\mathrm{Aut}(A)$ that fixes a point in $A^{\oplus \ell}$ using Smith normal form, and then all that would be needed is the general abelian analogue of Theorem 5.7 (see Open Problem 8.2). However, while the group action of $\mathrm{Aut}(A)$ on $H^2(Q, A)$ looks very similar to the action of $\mathrm{Aut}(A)$ on $A^{\oplus \ell}$ for some ℓ, it differs in an important way: The elements in $\mathrm{Aut}(A)$ that would normally induce surjections $\mathbb{Z}_{p^\mu}^d \twoheadrightarrow \mathbb{Z}_{p^{\mu'}}^d$ with $\mu' < \mu$ induce maps $H^2(Q, \mathbb{Z}_{p^\mu}) \to H^2(Q, \mathbb{Z}_{p^{\mu'}})$ which need not be surjections. (The failure to be a surjection is measured precisely by $H^3(Q, \mathbb{Z}_{p^{\mu - \mu'}})$.) So in this case, the best we can do naively is to apply generic algorithms for finding stabilizers in matrix groups, but the current state of the art is not fast enough for this to run in time $\mathrm{poly}(|Q|, |A|)$.

6.2. For general extensions of elementary abelian groups (Theorem B).

THEOREM 6.11. *Let $\mathcal{S}$ be a polynomial-time-computable characteristic subgroup*

function. For two groups G, H of order n, if $\mathcal{S}(G) \cong \mathbb{Z}_p^k$ and $\mathrm{Aut}(G/\mathcal{S}(G))$ can be listed in time $t(n)$, then isomorphism of G and H can be decided, and the coset of isomorphisms found, in time $t(n)n^{O(1)}$.

As before, let us first see how this applies to groups with elementary abelian radicals. Combining Theorem 6.11 with Theorem 5.5 and, respectively, Theorem 5.6, we have the following.

COROLLARY 6.12 (Theorem B). *Isomorphism of groups of order n with elementary abelian radicals can be decided, and the coset of isomorphisms found, in time $n^{c \log \log n + O(1)}$, for $c = 1/\log_2(60) \approx 0.169$.*

COROLLARY 6.13. *Let G and H be groups of order n with elementary abelian radicals. If $G/\mathrm{Rad}(G)$ has $O(\log n/\log \log n)$ minimal normal subgroups, isomorphism between G and H can be decided, and the coset of isomorphisms found, in $\mathrm{poly}(n)$ time.*

The proof of Theorem 6.11 is a reduction to MODULE CYCLICITY, for which a deterministic polynomial-time algorithm over finite fields is provided by Chistov, Ivanyos, and Karpinski [28]. This was recently generalized to finite modules (not necessarily over a field) [29], which we use to decide the case of general abelian radicals, but where the extension by the radical is split (see Remark 6.14), and we expect to have further uses. Before the reduction it might be helpful to see this problem in a special case, when the extensions are split.

Proof of Theorem 6.11 *for split extensions, a.k.a.* MODULE ISOMORPHISM. Recall that G_1 and G_2 are extensions of $A = \mathbb{Z}_p^k$ by Q. Furthermore, suppose both extensions split. Then to test isomorphism we are left with the ACTION COMPATIBILITY; that is, we extract the actions of Q on A in G_i as $\theta_i : Q \to \mathrm{Aut}(A) = \mathrm{GL}(k, p)$, and the goal is to find $(\alpha, \beta) \in \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ such that $\theta_1 = \theta_2^{(\alpha, \beta)}$. As $\mathrm{Aut}(Q)$ is enumerable, we fix a β and all that remains is to test whether there exists $\alpha \in \mathrm{GL}(k, p)$ such that for all $q \in Q$, $\theta_1(q) = \alpha^{-1}\theta_2(q)\alpha$. In other words, viewing θ_i as linear representations of Q over the field $\mathbb{F}_p$, the problem is to test whether these two representations are equivalent. This can also be formulated as finding a nonsingular matrix α such that $\alpha\theta_1(q) = \theta_2(q)\alpha$ for all $q \in Q$, namely MODULE ISOMORPHISM. Over finite fields this problem admits deterministic polynomial-time algorithms [28, 22].

Computing the coset of isomorphisms. When A is elementary abelian, we can also find the full coset of isomorphisms. For this, we need a generating set of the group of units of the matrix algebra $\{\alpha \in M(k, p) : \alpha\theta_1(q) = \theta_2(q)\alpha \forall q \in Q\}$. First, compute an $\mathbb{F}_p$-linear basis for this algebra by solving the given equations, which are linear in the entries of α. Next, from this linear basis for this matrix algebra, use Theorem 5.7 to compute a generating set of the group of units. □

Remark 6.14. In the case of split extensions, we can also decide isomorphism even when A is a general abelian group, not necessarily elementary. For this, we cannot use $\mathrm{GL}(k, p)$ but must stick with the more general group $\mathrm{Aut}(A)$; the rest of the above proof for deciding isomorphism goes through mutatis mutandis, reducing now to MODULE ISOMORPHISM for finite $\mathbb{Z}G$-modules, rather than finite(-dimensional) $\mathbb{F}_pG$-modules. MODULE ISOMORPHISM for finite $\mathbb{Z}G$-modules can also be solved deterministically in polynomial time [29]. For computing the coset of isomorphisms, we do not yet know how to compute the unit group of a finite $\mathbb{Z}$-algebra (rather than one over a field). See Open Problem 8.2.

Now we present the reduction for the general case.

Proof of Theorem 6.11: *Reduction to* MODULE CYCLICITY. As always, we start with the elementary abelian case. Let G_1 and G_2 be extensions of $A = \mathbb{Z}_p^k$ by Q, and (θ_i, f_i) the extension data of $A \hookrightarrow G_i \twoheadrightarrow Q$. It can be verified that if (α, β) satisfies $\theta_1 = \theta_2^{(\alpha,\beta)}$, then (α, β) sends $Z^2(Q, A, \theta_2)$ to $Z^2(Q, A, \theta_1)$ and sends $B^2(Q, A, \theta_2)$ to $B^2(Q, A, \theta_1)$. As $\mathrm{Aut}(Q)$ is enumerable, the problem is to find $\alpha \in \mathrm{GL}(k, p)$ such that (1) for all $q \in Q$, $\alpha\theta_1(q) = \theta_2(q)\alpha$, and (2) $\alpha f_1 = f_2$ as cohomology classes in $Z^2(Q, A, \theta_2)$ (that is, $[\alpha f_1] = [f_2]$).

This task can be reduced to MODULE CYCLICITY over finite-dimensional algebras, in almost the same way as the reduction from MODULE ISOMORPHISM to MODULE CYCLICITY [28]. We include a full proof here for completeness. The basic idea is as follows: Just as $\mathrm{Iso}(G_1, G_2)$ is a coset of $\mathrm{Aut}(G_1)$, $\mathrm{Hom}(G_1, G_2)$ is a "coset of"—more precisely, is acted on by—$\mathrm{End}(G_1) := \mathrm{Hom}(G_1, G_1)$. If one could find $\mathrm{Hom}(G_1, G_2)$, the problem would be reduced to finding an invertible element there. The following is the linear-algebraic version of these ideas, in which finding an invertible element inside a set of not-necessarily-invertible homomorphisms is reduced to MODULE CYCLICITY, for the analogue of $\mathrm{Hom}(G_1, G_2)$ considered as a module over the analogue of $\mathrm{End}(G_1)$.

Let $M(k, p)$ be the linear space of $k \times k$ matrices over $\mathbb{Z}_p$. Consider a linear subspace of $M(k,p)$, $V = \{\alpha \in M(k,p) \mid \forall q \in Q, \alpha\theta_1(q) = \theta_2(q)\alpha,$ and $\exists a \in \mathbb{Z}_p, [\alpha f_1] = [af_2]\}$. Also consider $U = \{\gamma \in M(k,p) \mid \forall q \in Q, \gamma\theta_2(q) = \theta_2(q)\gamma,$ and $\exists a \in \mathbb{Z}_p, [\gamma f_2] = [af_2]\}$. It can be verified that U is an associative algebra over $\mathbb{Z}_p$ with identity. Note that we need to allow the possible scalar a in the definition of U in order for U to be closed under addition. This does not hurt the multiplicative condition, since $\alpha\theta_2(q) = \theta_2(q)\alpha$ is not changed if we multiply both sides by a scalar. Once the scalar is allowed in the definition of U, we must also allow it in the definition of V, in order for V to be a U-module (see the next paragraph). However, if $[\gamma f_2] = [af_2]$, then $a^{-1}\gamma \in U$ and $a^{-1}\gamma)[f_2] = [f_2]$, so U essentially consists of endomorphisms of G "up to scale."

Then V is a left U-module: for $\alpha \in V$, $\gamma \in U$, and $q \in Q$, $\gamma\alpha\theta_1(q) = \gamma\theta_2(q)\alpha = \theta_2(q)\gamma\alpha$. To show that $[\gamma\alpha f_1] = [af_2]$ is a little subtle, and for this we need to recall the fact that if $\gamma\theta_2(q) = \theta_2(q)\gamma$ for every $q \in Q$, then γ preserves $B^2(Q, A, \theta_2)$. That is, $\alpha f_1 = af_2 + g$ for some $a \in \mathbb{Z}_p$ and $g \in B^2(Q, A, \theta_2)$, and $\gamma\alpha f_1 = \gamma(af_2 + g) = a\gamma f_2 + \gamma g = a'f_2 + g' + g''$, where $a' \in \mathbb{Z}_p$, $a\gamma f_2 = a'f_2 + g'$, and $\gamma g = g''$. Now we claim that if V contains invertible elements, then (1) it is cyclic, and (2) every generator is invertible. To show (1), let $\alpha' \in V$ be invertible, and form $\phi : U \to V$ by sending $\gamma \to \gamma\alpha'$. Then ϕ is a U-module isomorphism between U and V, whose inverse is $V \to U$ by $\alpha \to \alpha\alpha'^{-1}$; $\alpha\alpha'^{-1} \in U$ again follows from the fact that α and α' can be shown to send $B^2(Q, A, \theta_2)$ to $B^2(Q, A, \theta_1)$ as a consequence of $\alpha\theta_1(q) = \theta_2(q)\alpha$. For (2), if α'' generates V, then $\alpha''\alpha'^{-1}$ generates U as a left U-module, and thus $\alpha''\alpha'^{-1}$ is invertible, showing that α'' is invertible. Finally, we note that if some invertible $\alpha' \in V$ sends $[f_1]$ to $[af_2]$ for some $a \in \mathbb{Z}_p$, then $a^{-1}\alpha' \in V$ is also invertible and sends $[f_1]$ to $[f_2]$.

Given the above reduction, here is an algorithm for the general case; we still represent 2-cocycles by $k \times |Q|^2$ matrices over $\mathbb{Z}_p$. We first compute a $\mathbb{Z}_p$-basis of $B^2(Q, A, \theta_2)$ as in Proposition 6.5. Using these 2-cocycles, we can represent V and U as solution spaces of homogeneous linear equations. Finally, we apply the polynomial-time MODULE CYCLICITY algorithm [28], either to get that V is not cyclic, and thus does not contain invertible elements, or to get a generator $\alpha' \in V$. In the latter case we conclude based on whether α' is invertible or not.

Computing the coset of isomorphisms. Any single invertible element of V yields an isomorphism. As in Theorem 6.1, a generating set of automorphisms is determined by (1) its image in $\mathrm{Aut}(A) \times \mathrm{Aut}(Q)$, and (2) the kernel of the natural map $\mathrm{Aut}(G) \to \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$. In this case, (1) is easily determined by a generating set of the group of units of U, and (2) is a generalization of a homomorphism from $Q \to A$ called a "derivation," which we describe below (see, e.g., [80, sect. 3.3]). In the elementary abelian case, the unit group of U can be computed from a linear basis for U by Theorem 5.7. All that remains is to compute the derivations $Q \to A$.

Recall that, for a section $s\colon Q \to G$, every element of G is uniquely represented as $as(q)$ for some $a \in A$, $q \in Q$. Then we have that $(as(q))(a's(q')) = a(s(q)a's(q)^{-1})f_s(q,q')s(qq')$. An automorphism that fixes both Q and A pointwise is fully determined by a map $\delta\colon Q \to A$, as in Theorem 6.1, but now satisfying the more general condition:

$$\delta(q)(s(q)\delta(q')s(q)^1) = \delta(qq').$$

Such maps δ are called *derivations*. As conjugation of Q on A is an action by automorphisms, this amounts to $|Q|^2$ homogeneous linear equations in $|Q| \dim A$ variables over the field $\mathbb{Z}_p$; such linear equations can be solved in polynomial time, yielding a generating set for the kernel of the map $\mathrm{Aut}(G) \to \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$. Putting this all together yields a generating set for $\mathrm{Aut}(G)$. □

7. When Aut(Q) is too big. In this section we present polynomial-time algorithms for certain central-radical groups even when $\mathrm{Aut}(Q)$ cannot be enumerated in polynomial time (unconditionally, simply because its size is superpolynomial). In particular, we present two fixed-parameter polynomial-time algorithms for central-radical groups with $G/\mathrm{Rad}(G)$ a direct product nonabelian simple groups, or a direct product of small perfect groups.

THEOREM 7.1 (Theorem C). *Isomorphism between two groups G_1, G_2 with central, elementary abelian radicals can be decided, and the coset of isomorphisms found, in* $\mathrm{poly}(|G_i|)$ *time if either*
1. $G_1/\mathrm{Rad}(G_1)$ *is a direct product of simple groups; or*
2. $G_1/\mathrm{Rad}(G_1)$ *is a direct product of perfect groups, each of order $O(1)$.*

Theorem 7.1 yields polynomial-time algorithms for the following concrete cases: For example, (1) covers groups with $\mathrm{Rad}(G) = Z(G) = \mathbb{Z}_p^k$ and $G/Z(G) \cong A_m^k$; (2) covers the case when $\mathrm{Rad}(G) = Z(G) = \mathbb{Z}_p^k$ and $G/Z(G) \cong (A_5 \wr \mathrm{PSL}_{20}(\mathbb{F}_7))^k$, where the wreath product is taken with respect to any permutation representation of $\mathrm{PSL}_{20}(\mathbb{F}_7)$.

Remark 7.2. In Theorem 7.1(2) we can relax the $O(1)$ size bound to $O(f(n))$ to get an algorithm that runs in time $n^{O(f(n)^2 \log f(n))}$. For $f(n) = (\log \log n)^{1/2-\varepsilon}$ for some fixed $\varepsilon > 0$, Theorem 7.1(2) becomes $n^{o(\log \log n)}$-time (really-very-nearly-polynomial time).

We remark that while it is possible that some of the algorithms in this paper could yield practically efficient implementations by using suitable subroutines (with or without worst-case guarantees), Theorem 7.1(2) is admittedly not practical: The square of the $O(1)$ bound appears in the exponent of the runtime of the algorithm, and the smallest centerless, indecomposable perfect group that is not simple—since simple groups are handled by the much more efficient algorithm from the first part of the theorem—has order 960 (the group is a semidirect product of A_5 acting on $\mathbb{F}_2^4$; see [48, sect. 5.3, Group (4,1)]).

For Theorem 7.1, currently we can only work with elementary abelian groups; an open problem posed in [11, sect. 7.7], namely the group code equivalence problem over cyclic groups, seems to be the only current obstacle for part (2); see Observation 7.8. For part (1), we show in Lemma 7.9 that the cohomological bound *can* be extended to the general abelian case; however, essentially the same obstacle as in subsection 6.1.2 still needs to be overcome in this approach.

As mentioned before, for Theorem 7.1 we need a more detailed understanding of central extensions in this special group class, including the deep Theorem 5.3 [43]. We also remind the reader that singly exponential algorithms for LINEAR CODE EQUIVALENCE and COSET INTERSECTION play an important role in Theorem 7.1(1).

7.1. Preparations from cohomology. Let A be an abelian group, and T_1, ..., T_ℓ be perfect, centerless, directly indecomposable groups (this includes the case in which the T_i are nonabelian simple; in fact, the reader unfamiliar with perfect groups will not lose too much by just considering the nonabelian simple case instead). For an extension $A \hookrightarrow G \twoheadrightarrow Q$ with $Q = \prod_{i \in [\ell]} T_i$, let U_i be the inverse image of T_i in G under the natural projection from G to Q. The following proposition adapted from Suzuki [83] is crucial (see Appendix A.5 for a proof).

PROPOSITION 7.3 (cf. [83, Chap. 6, Prop. 6.5]). *Let the notation be as above. For $i, j \in [\ell]$, $i \neq j$, $[U_i, U_j] = \mathrm{id}$. That is, for all $x \in U_i$, for all $y \in U_j$, $xy = yx$.*

We now consider the U_i not just as subgroups of G but as extensions $A \hookrightarrow U_i \twoheadrightarrow T_i$. As Proposition 7.3 shows that $[U_i, U_j] = \mathrm{id}$ for $i \neq j$, these extensions determine the extension $A \hookrightarrow G \twoheadrightarrow Q$ as follows.

LEMMA 7.4. *Given two central extensions $A \hookrightarrow G_j \twoheadrightarrow Q$ ($j = 1, 2$) with $A = Z(G_j)$ and $Q = \prod_{i=1}^{\ell} T_i$, with the T_i perfect, centerless, and indecomposable, let $U_{j,i}$ be the inverse image of T_i under the natural projection $G_j \to G_j/A$. The extensions $A \hookrightarrow G_j \twoheadrightarrow Q$ ($j = 1, 2$) are equivalent if and only if for each $i \in [\ell]$ the extensions $A \hookrightarrow U_{j,i} \twoheadrightarrow T_i$ ($j = 1, 2$) are equivalent.*

Proof. The "only if" direction is trivial. For the other direction, for $j = 1, 2$ and $i \in [\ell]$, let $f_{j,i}$ be the 2-cocycle of the extension $A \hookrightarrow U_{j,i} \twoheadrightarrow T_i$ induced by some section $s_{j,i} : T_i \to U_i$. By standard cohomology (see, e.g., Appendix A), as $U_{1,i}$ and $U_{2,i}$ are equivalent extensions for each i, $f_{1,i} - f_{2,i}$ is some 2-coboundary $b_i \in B^2(T_i, A)$. To show the equivalence of $A \hookrightarrow G_j \twoheadrightarrow \prod_{i \in [\ell]} T_i$, we need only exhibit two 2-cocycles f_j for $A \hookrightarrow G_j \twoheadrightarrow Q$ that differ by a 2-coboundary.

As Q is decomposed uniquely as $\prod_i T_i$ (that is, uniquely as subsets of Q, not just up to isomorphism), we can identify elements in Q as from $\prod_i T_i$ without ambiguity. Let $(p_1, \ldots, p_\ell)$ and $(q_1, \ldots, q_\ell)$ be two elements in Q, $p_i, q_i \in T_i$ for $i \in [\ell]$. Then define $b : Q \times Q \to A$ as

$$\tag{7} b((p_1, \ldots, p_\ell), (q_1, \ldots, q_\ell)) = \sum_{i \in [\ell]} b_i(p_i, q_i).$$

Using the $\mathbb{Z}$-linearity of the coboundary condition, one can verify that b is a 2-coboundary in $B^2(Q, A)$.

Recall that the 2-cocycle $f_{j,i}$ is induced by the section $s_{j,i} : T_i \to U_i$. We define a section $s_j : Q \to G_j$, as $s_j((p_1, \ldots, p_\ell)) = s_{j,1}(p_1) \ldots s_{j,\ell}(p_\ell)$. Let f_j be the 2-cocycle induced by s_j, then—noting that for $i_1 \neq i_2$, $s_{j,i_1}(p_{i_1})$ and $s_{j,i_2}(p_{i_2})$ commute by

Proposition 7.3—it can be verified that

$$(8) \qquad f_j((p_1, \ldots, p_\ell), (q_1, \ldots, q_\ell)) = \sum_{i \in [\ell]} f_{j,i}(p_i, q_i).$$

Thus $f_1 - f_2 = b \in B^2(Q, A)$, finishing the proof. $\qquad\square$

For convenience, in the following we shall call $U_{j,i}$ the *restriction of G_j to T_i* and use $G_j|_{T_i}$ to denote it. The next lemma concerns the direct product structure of the normal part; its proof is included in Appendix A.5 for completeness.

LEMMA 7.5. *Let $A' \times A'' \hookrightarrow G \twoheadrightarrow Q$ be a central extension of $A' \times A''$ by Q. Let $p_{A'} \colon A' \times A'' \to A'$ be the projection onto A' along A''. If there is a 2-cocycle $f \colon Q \times Q \to A' \times A''$ such that $p_{A'} \circ f \colon Q \times Q \to A'$ is a 2-coboundary, then G is isomorphic (even equivalent) to the direct product $A' \times (G/A')$.*

Furthermore, A' can be computed in $\mathrm{poly}(|G|)$ time using linear algebra over abelian groups.

Using general algorithms for decomposing direct products [55, 91], we could compute A' in polynomial time without the "furthermore." However, in the setting of Lemma 7.5, we give a much simpler algorithm to compute A' using linear algebra over abelian groups.

7.2. Warm-up result. The full proof of Theorem 7.1 requires several ideas. To highlight these ideas separately, we start by proving a warm-up result that only needs some of these ideas. The following result is the same as Theorem 7.1, except that instead of requiring that the central radical be elementary abelian, it requires that $|\mathrm{Aut}(\mathrm{Rad}(G))| \leq \mathrm{poly}(|G|)$. (Formally, these two conditions are incomparable, but the condition used here "feels" more stringent.)

PROPOSITION 7.6. *Isomorphism between two groups G_1, G_2 with central radicals satisfying $|\mathrm{Aut}(\mathrm{Rad}(G_1))| \leq \mathrm{poly}(|G_1|)$ can be decided, and the coset of isomorphisms found, in $\mathrm{poly}(|G_1|)$ time if either*
1. *$G_1/\mathrm{Rad}(G_1)$ is a direct product of simple groups; or*
2. *$G_1/\mathrm{Rad}(G_1)$ is a direct product of perfect groups, each of order $O(1)$.*

As a few examples of when the condition on $\mathrm{Aut}(\mathrm{Rad}(G))$ holds, note that $\mathrm{Aut}(R) \leq |R|^{\log |R|}$ for all groups R, and $|\mathrm{Aut}(\mathbb{Z}_{p^e}^\ell)| \sim p^{e\ell^2}$, so the condition holds if $|\mathrm{Rad}(G)| \leq 2^{O(\sqrt{\log |G|})}$ or if $\mathrm{Rad}(G) = \mathbb{Z}_{p^e}^\ell$ with $e\ell^2 \log p \leq O(\log |G|)$.

Proof. Decompose each G_j ($j = 1, 2$) as an extension of A by $Q = \prod_{i \in [\ell]} T_i$, where each T_i is perfect and directly indecomposable, and each T_i is necessarily centerless because $A = \mathrm{Rad}(G)$. Regardless of whether we are in case (1) or (2), this decomposition is algorithmically straightforward by [11, Prop. 2.1]; although it was stated there only for case (1), the same algorithm works in both cases, because the direct factors are uniquely determined not only up to isomorphism, but as subsets of G_j, even in the perfect case.

Now we group the T_i according to their isomorphism types, identifying $Q = \prod_{i \in [r]} Q_i^{\ell_i}$, where r is the number of isomorphism types among the T_i, each Q_i is isomorphic to some T_i, and the Q_i's are pairwise nonisomorphic. In case (1), this can be done as each T_i is generated by two elements, and in case (2) it can be done as each T_i has size $O(1)$ by assumption.

Then $\mathrm{Aut}(Q) \cong \prod_{i \in [r]} \mathrm{Aut}(Q_i) \wr S_{\ell_i} \cong \prod_{i \in [r]} (\mathrm{Aut}(Q_i)^{\ell_i} \rtimes S_{\ell_i})$. A *diagonal* of $\mathrm{Aut}(Q)$ is an element in $\prod_{i \in [r]} \mathrm{Aut}(Q_i)^{\ell_i}$. All diagonals are enumerable in polynomial

time by the $n^{\# \text{ generators}}$ technique: In case (1), by Fact 5.1, $|\prod_{i \in [r]} \operatorname{Aut}(Q_i)^{\ell_i}| \leq$ $(\prod_{i \in [r]} |Q_i|^2)^{\ell_i} \leq |G_j|^2$, and in case (2), $|\prod_{i \in [r]} \operatorname{Aut}(Q_i)^{\ell_i}| \leq (\prod_{i \in [r]} |Q_i|^{\log |Q_i|})^{\ell_i} \leq$ $O(1)^{O(\log |G_j|)} = \operatorname{poly}(|G_j|)$.

By Lemma 2.3, G_1 and G_2 are isomorphic if and only if they are pseudo-congruent extensions of A by Q. The extensions are pseudo-congruent if and only if there is an element of $\operatorname{Aut}(A) \times \operatorname{Aut}(Q)$ such that, after twisting by this element, the resulting extensions are equivalent. Once an element of $\operatorname{Aut}(A) \times \operatorname{Aut}(Q)$ is fixed, by Lemma 7.4, the latter problem is reduced to determining the equivalence of $G_1|_{T_i}$ and $G_2|_{T_i}$ for each $i \in [\ell]$.

Note that the equivalence type of the extension $G_j|_{T_i}$ can be computed in polynomial time by Theorem 6.1, as $\operatorname{Aut}(T_i)$ can be listed in polynomial time: In case (1) each T_i is generated by two elements, and in case (2) each T_i has size $O(1)$.

Having decomposed G_j as $A \hookrightarrow G_j \twoheadrightarrow \prod_{i \in [r]} Q_i^{\ell_i}$, the algorithm then proceeds as follows. For every $\alpha \in \operatorname{Aut}(A)$, and every diagonal $\prod_{i \in [\ell]} \delta_i$ of $\prod_{i \in [\ell]} \operatorname{Aut}(T_i)$, do the following. Apply α^{-1} and δ_i to each restricted extension $G_2|_{T_i}$. Now compute the equivalence types of $G_2|_{T_i}$. If the multiset of equivalence types coming from $G_1|_{T_i}$ is equal to the multiset of equivalence types from the $(\alpha, \prod_i \delta_i)$-twisted $G_2|_{T_i}$, then G_1 and G_2 are pseudo-congruent as extensions, and the algorithm reports "isomorphic." On the other hand, if the equivalence of multisets is not detected for any $(\alpha, \prod_i \delta_i)$, then the algorithm returns "not isomorphic."

It is obvious that the above procedure runs in polynomial time in $|G_1|$ and $|\operatorname{Aut}(A)|$. We remark that if $T_{i_1} \not\cong T_{i_2}$, then $G_1|_{T_{i_1}}$, and the (α, δ_{i_2})-twisted $G_2|_{T_{i_2}}$ cannot be equivalent. Thus the multiset of equivalence types distinguishes the isomorphism types of T_i's automatically. Finally, it is enough to compare the multisets because we have full symmetric groups S_{ℓ_i} acting on the isomorphic factors.

To find the coset of isomorphisms, note that the preceding algorithm enumerates over every element of $\operatorname{Aut}(A)$ and every diagonal of $\prod_{i \in [\ell]} \operatorname{Aut}(T_i)$. For each such pair, the algorithm determines a multiset of equivalence types, so the relevant subgroup of S_{ℓ_i} is just a set transporter, which is easily calculated. Compute the derivations and apply Observation 5.2 twice, as in the proof of Theorem 6.11. □

7.3. Proof of Theorem C. Unlike the warm-up result from the previous section, for Theorem 7.1 we can no longer afford to enumerate $\operatorname{Aut}(\operatorname{Rad}(G))$. Avoiding this enumeration leads us to two distinct, more advanced methods, each of which may have further applications.

Proof of Theorem 7.1 *(Theorem C).* We start as in the proof of Proposition 7.6. For G_j, $j = 1, 2$, decompose it as an extension of $A = \mathbb{Z}_p^k$ by $Q = \prod_{i \in [\ell]} T_i$. Classify T_i's according to their isomorphism types and group them together, identifying $Q = \prod_{i \in [r]} Q_i^{\ell_i}$. Then $\operatorname{Aut}(Q) \cong \prod_{i \in [r]} \operatorname{Aut}(Q_i) \wr S_{\ell_i} \cong \prod_{i \in [r]} (\operatorname{Aut}(Q_i)^{\ell_i} \rtimes S_{\ell_i})$. A *diagonal* of $\operatorname{Aut}(Q)$ is an element in $\prod_{i \in [r]} \operatorname{Aut}(Q_i)^{\ell_i}$. All diagonals are enumerable in polynomial time by the $n^{\# \text{ generators}}$ technique. By Lemma 2.3, $G_1 \cong G_2$ if and only if they are pseudo-congruent extensions of A by Q. The extensions are pseudo-congruent if and only if there is an element of $\operatorname{Aut}(A) \times \operatorname{Aut}(Q)$ such that, after twisting by this element, the resulting extensions are equivalent. Once an element of $\operatorname{Aut}(A) \times \operatorname{Aut}(Q)$ is fixed, by Lemma 7.4, the latter problem is reduced to determining the equivalence of $G_1|_{T_i}$ and $G_2|_{T_i}$ for each $i \in [\ell]$.

We will also need an additional consequence of Lemma 7.4. For G_j, $j = 1, 2$, we say a 2-cocycle $f_j : Q \times Q \to A$ respects the direct factors if there exist $\tilde{f}_{j,i} : T_i \times T_i \to A$, $i \in [\ell]$ such that (8) holds. Let $Z^2_{\text{prod}}(Q, A)$ denote the set of 2-cocycles respecting

the direct factors. The proof of Lemma 7.4 shows that $Z^2_{\text{prod}}(Q, A) \neq \emptyset$. Similarly we can define 2-coboundaries that respect the direct factors $B^2_{\text{prod}}(Q, A)$ using (7). The difference of two cohomologous 2-cocycles in $Z^2_{\text{prod}}(Q, A)$ is automatically in $B^2_{\text{prod}}(Q, A)$. We define $C^2_{\text{prod}}(Q, A)$ as the 2-cochains respecting the direct factors. We may view elements of $C^2_{\text{prod}}(Q, A)$ as $k \times \sum_{i \in [\ell]} |T_i|^2$ matrices, whose rows are indexed by $[k]$ and whose columns are indexed by triples $(i; p, q)$ with $p, q \in T_i$. That is, $C^2_{\text{prod}}(Q, A) = \bigoplus_{i \in [\ell]} C^2(T_i, A)$.

(1). Note that the equivalence type of the extension $G_j|_{T_i}$ can be computed in polynomial time by Theorem 6.1, as each T_i is generated by two elements, and thus $\text{Aut}(T_i)$ can be listed in polynomial time. The algorithm then proceeds as follows.

As every $\beta \in \text{Aut}(Q)$ can be represented as $(\delta, \sigma) \in (\prod_i \text{Aut}(Q_i)^{\ell_i}) \rtimes (\prod_i S_{\ell_i})$, it follows that $Z^2_{\text{prod}}(Q, A)$ (resp., $B^2_{\text{prod}}(Q, A)$) is an invariant subset in $Z^2(Q, A)$ (resp., $B^2(Q, A)$) under the actions of both $\text{Aut}(A)$ and $\text{Aut}(Q)$. Using Proposition 6.10, for each $i \in [\ell]$ we get projection $\pi_i \colon C^2(T_i, A) \to W_i$ for some $\text{Aut}(A)$-invariant $W_i \leq C^2(T_i, A)$ such that $C^2(T_i, A) = B^2(T_i, A) \oplus W_i$ and π_i commutes with the action of $\text{Aut}(A)$.

Theorem 5.3 tells us that $\dim W_i = \dim H^2(T_i, A) \leq 17.5 \dim A$, and the "alternative viewpoint" at the end of subsection 6.1.2 implies that W_i is a direct sum of copies of A. By a judicious choice of basis, we may thus write elements of W_i as $k \times 17$ matrices, in a way that is still $\text{Aut}(A)$-equivariant. (This does not imply that the 17.5 in Theorem 5.3 can be replaced by 17 in general, but rather only in the case when Q acts trivially on A.) From now on, let π_i denote the composition of the previous π_i, followed by this mapping onto $k \times 17$ matrices.

For each choice of diagonal $\delta \in \prod_{i \in [\ell]} \text{Aut}(T_i)$, we may choose the complements W_i so that whenever i, j are such that $T_i \cong T_j$, that W_i and W_j are identified by δ. With this choice, we may direct sum these π_i together to get a homomorphism $\pi = \pi_\delta \colon C^2_{\text{prod}}(Q, A) \to W \leq M(k, 17\ell)$ for some $\text{Aut}(A)$-invariant $W \leq C^2_{\text{prod}}(Q, A)$ such that $C^2_{\text{prod}}(Q, A) = B^2_{\text{prod}}(Q, A) \oplus W$ and π commutes with the action of $\text{Aut}(A) \times \prod_i S_{\ell_i}$. For each diagonal δ, the question then is to decide the existence of $(\alpha, \sigma) \in \text{Aut}(A) \times (\prod_i S_{\ell_i})$ such that $\pi(f_1) = \pi(f_2^{(\text{id}, \delta, \text{id})})^{(\alpha, \text{id}, \sigma)}$. (Note that although $(\text{id}, \text{id}, \sigma)$ does not commute with $(\text{id}, \delta, \text{id})$, we do have that $(\alpha, \delta, \sigma) = (\text{id}, \delta, \text{id})(\alpha, \text{id}, \sigma)$, since σ is already "on the right" of δ, and α and δ do commute.)

Let $M_1 = M_{\pi(f_1)}$. By Lemma 7.5, without loss of generality we may assume M_1 is of rank k. Otherwise Lemma 7.5 splits a direct factor out of the center as $A' \times G_j/A'$ (although Lemma 7.5 is concerned with $Z^2(Q, A)$ and $B^2(Q, A)$, it is readily adapted to $Z^2_{\text{prod}}(Q, A)$ and $B^2_{\text{prod}}(Q, A)$). By the Remak–Krull–Schmidt theorem, we then reduce to testing isomorphism between G_1/A' and G_2/A', where the desired rank condition holds. Lemma 7.5 allows us to compute such $A' \leq Z(G_j)$.

The algorithm thus proceeds as follows. For every diagonal $\delta = \prod_{j \in [\ell]} \delta_j$ of $\prod_j T_j$, compute $\pi(f_1)$ and $\pi(f_2^{(\text{id}, \delta, \text{id})})$, and let $M_1 = M_{\pi(f_1)}$ and $M_2 = M_{\pi(f_2^{(\text{id}, \delta, \text{id})})}$. Thus M_j is the matrix of size $k \times (17\ell)$ corresponding to f_j.

As the action of $\text{Aut}(A)$ on the M_i is by left multiplication, and the action of $\prod_{i \in [\ell]} S_{\ell_i}$ is by permuting blocks of columns, we treat the M_i as generator matrices of two $\mathbb{Z}_p$-linear codes of dimension k and length 17ℓ. Compute the coset of equivalences $\text{CodeEq}(M_1, M_2) \subseteq S_{17\ell}$ using Theorem 5.4. On the other hand, $\prod_i S_{\ell_i}$ induces an action on $[17\ell]$, permuting blocks of size 17, and these are the only permutations we want. Thus we intersect $\text{CodeEq}(M_1, M_2)$ with $\prod_i S_{\ell_i}$. If the intersection is nonempty, the algorithm reports "isomorphic." On the other hand, if for all diagonals we get

empty intersection, then the algorithm returns "not isomorphic."

To analyze the running time, the outer loop depending on the diagonals is polynomially related to n. Both the applications of the LINEAR CODE EQUIVALENCE algorithm (Theorem 5.4) and the singly exponential time algorithm for COSET INTERSECTION [63] take time $c^{17\ell}$ for some absolute constant c. Since $\ell \leq \log |G|$, this is poly($|G|$) time.

(2). This case proceeds as in the previous case, with minor modifications. Now, to decide isomorphism of the centerless, indecomposable, perfect groups T_i we rely on the assumption that they have size $O(1)$ (since we do not, otherwise, know good bounds on testing them for isomorphism). Also, in the absence of an analogue of Theorem 5.3 for this class of groups, rather than 17ℓ, we use the naive bound that $\dim H^2(T_i, A) \leq \dim C^2(T_i, A) = |T_i|^2 \dim A$. Let $D \leq O(1)$ be the maximum of the $|T_i|$; then the singly exponential algorithms for LINEAR CODE EQUIVALENCE and COSET INTERSECTION run in time $c^{\ell D^2} \leq c^{\ell \cdot O(1)} \leq$ poly($|G|$).

Computing the coset of isomorphisms. The algorithm is essentially the same in both cases (1) and (2). To find a single isomorphism, whenever δ is a diagonal such that the coset intersection CodeEq(M_1, M_2) $\cap \prod_i S_{\ell_i}$ is nonempty, we pick a single element σ of this coset intersection. Apply $(\delta, \sigma) \in \text{Aut}(Q)$ to the columns of M_2, and then compute using linear algebra a matrix in $\text{Aut}(A) \cong \text{GL}(k, p)$ that makes M_1 equal to the column-permuted M_2.

To find a generating set for the automorphisms of G_1, we proceed as follows. As in Theorems 6.1 and 6.11, we can compute the homomorphisms $Q \to A$ using linear algebra; these constitute the kernel of the map $\text{Aut}(G_1) \to \text{Aut}(A) \times \text{Aut}(Q)$. Next, we compute generators of the image of $\text{Aut}(G_1)$ in $\text{Aut}(A) \times \text{Aut}(Q)$. We apply Observation 5.2 again, viewing $\text{Aut}(G_1)$ as an extension of $\text{Aut}(A) \times \prod_i \text{Aut}(Q_i)^{\ell_i}$ by $\prod_i S_{\ell_i}$. For each diagonal $\delta \in \prod_i \text{Aut}(Q_i)^{\ell_i}$, we compute the coset intersection CodeEq(M_1, M_1) $\cap \prod_i S_{\ell_i}$. The union of these coset intersections generates the image of $\text{Aut}(G_1)$ in $\prod_i S_{\ell_i}$. For each of these generators σ, we know which diagonal δ it came from, and we can compute an $\alpha \in \text{Aut}(A)$ such that $(\alpha, \delta, \sigma) \in \text{Aut}(G_1)$, as in the preceding paragraph. This gives us lifts of all the generators of the image of $\text{Aut}(G_1)$ in $\prod_i S_{\ell_i}$. Finally, we need to compute generators of the kernel of the map $\text{Aut}(G_1) \to \prod_i S_{\ell_i}$, as a subgroup of $\text{Aut}(A) \times \prod_i \text{Aut}(Q_i)^{\ell_i}$.

For this, we will apply Observation 5.2 one last time. Let $\text{Aut}_0(G_1)$ denote the subgroup of $\text{Aut}(G_1)$ whose projection to $\prod_i S_{\ell_i}$ is trivial. We can determine the projection of $\text{Aut}_0(G_1)$ to $\prod_i \text{Aut}(Q_i)^{\ell_i}$ by enumerating over all diagonals, and including only those diagonals δ for which the matrix associated to $\pi(f_1)$ and the matrix associated to $\pi(f_1^{(\text{id}, \delta, \text{id})})$ have the same rowspan. For each one, we lift it to $\text{Aut}_0(G_1)$ by using linear algebra to find an element of $\text{Aut}(A) \cong \text{GL}(k, p)$ that makes the two matrices equal. Finally, we need to compute generators of the subgroup of $\text{Aut}(A)$ that fix $\pi(f_1)$. This proceeds exactly as in Theorem 6.1.   □

7.4. Toward the general abelian case. In this section we give two results headed in the direction of extending Theorem 7.1 from elementary abelian central radicals to general central radicals. First, we give an alternative proof of Theorem 7.1(2), without using Proposition 6.10. Avoiding Proposition 6.10 is key, because it may not have an analogue for the general abelian case (see the remarks at the end of subsection 6.1.2). This alternative proof also has the advantage that it makes clear how to extend to the general abelian case by using group code equivalence over cyclic p-groups (Observation 7.8). Then we observe that Theorem 5.3 can be extended from modules over fields to general finite $\mathbb{Z}G$-modules, which we expect to be useful for

proving the general abelian case of Theorem 7.1(1).

Alternative proof of Theorem 7.1(2). As in section 6, let M_{f_j} denote the matrix representation of f_j with row index set $[k]$ and column index set $Q \times Q$. As f_j is completely determined by the direct factors, we can focus on M_{f_j} with column indices from $\cup_{i \in [\ell]} T_i \times T_i$. Thus for $f_j \in Z^2_{\text{prod}}(Q, A)$ the size of M_{f_j} becomes $k \times (\sum_{i \in [\ell]} |T_i|^2)$.

We will need the following analogue of Proposition 6.6. Let $\widetilde{M_{f_j}}$ denote the matrix with $\sum_{i \in [\ell]} |T_i|^2$ columns, and whose first rows are just M_{f_j}. The remaining rows will be the union of bases for $B^2(T_i, \mathbb{Z}_p)$ for each i (see Proposition 6.4).

PROPOSITION 7.7. *Suppose that $\widetilde{M_{f_j}}$ has full rank for $j = 1, 2$. Fix some diagonal $\delta \in \prod_i \text{Aut}(T_i)$ and let $\widetilde{M_{f_1}}' = \widetilde{M_{f_1}}^{(\text{id},\delta,\text{id})}$. Then the intersection $\text{CodeEq}(\widetilde{M_{f_1}}', \widetilde{M_{f_2}}) \cap \prod_i S_{\ell_i}$ is nonempty if and only if there exists $(\alpha, \sigma) \in \text{Aut}(A) \times \prod_i S_{\ell_i}$ such that (α, δ, σ) is an isomorphism of f_1 and f_2 as cohomology classes.*

Proof. Let N be the number of rows of $\widetilde{M_{f_j}}$, and let $b = N - k$ be the number of rows of $\widetilde{M_{f_j}}$ that were added to $\widetilde{M_{f_j}}$ compared to M_{f_j}.

($\Rightarrow$) Suppose $\text{CodeEq}(\widetilde{M_{f_1}}', \widetilde{M_{f_2}}) \cap \prod_i S_{\ell_i}$ is nonempty. Then there is some permutation $\sigma \in \prod_i S_{\ell_i}$ and some $\Lambda \in \text{GL}(N, p)$ such that $\Lambda \widetilde{M_{f_1}}^{(\delta,\sigma)} = \widetilde{M_{f_2}}$. As the last b rows of $\widetilde{M_{f_1}}$ and $\widetilde{M_{f_2}}$ have the same rowspan (namely, $B^2_{prod}(Q, A)$) and this rowspan is preserved by all permutations of the columns (i.e., elements of $\prod_i S_{\ell_i}$), we may assume without loss of generality that Λ has the following block form:

$$\Lambda = \begin{pmatrix} \alpha & \gamma \\ 0 & \eta \end{pmatrix}.$$

In other words, to make the row spans of the bottom b rows equal, it is never necessary to add any multiples of the top k rows to the bottom b rows, as the bottom b rows already have equal row spans that are preserved by all relevant permutations, in particular by σ.

The submatrix γ contributes by adding elements of $B^2_{prod}(Q, A)$ to M_{f_1} so that

$$\begin{pmatrix} \text{id} & \gamma \\ 0 & \eta \end{pmatrix} M_{f_1}$$

corresponds to a cocycle that is cohomologous to f_1. Finally, the contribution of the submatrix α is to send f_1 to a pseudo-congruent cocycle, since $\alpha \in \text{Aut}(A)$. Therefore, we have shown that (α, δ, σ) is an isomorphism of the cohomology classes f_1, f_2.

($\Leftarrow$) Suppose that $f_1^{(\alpha,\delta,\sigma)}$ is cohomologous to f_2 for some $\alpha \in \text{Aut}(A)$ and $\sigma \in \prod S_{\ell_i}$. Then some matrix

$$\Lambda' = \begin{pmatrix} \alpha & \gamma \\ 0 & \text{id} \end{pmatrix}$$

will make the first k rows of $\Lambda' \widetilde{M_{f_1}}^{(\delta,\sigma)}$ equal to the first k rows of $\widetilde{M_{f_2}}$. The last b rows of $\Lambda' \widetilde{M_{f_1}}^{(\delta,\sigma)}$ and $\widetilde{M_{f_2}}$ have the same row span, so there is some η such that

$$\Lambda = \begin{pmatrix} \alpha & \gamma \\ 0 & \eta \end{pmatrix}$$

makes the two matrices equal. (In fact, η will be a block-permutation matrix, which permutes the blocks in the same way that δ permutes the factors.) In particular, this

shows that σ is a code equivalence, and hence that $\mathrm{CodeEq}(\widetilde{M_{f_1}}', \widetilde{M_{f_2}}) \cap \prod_i S_{\ell_i}$ is nonempty. $\qquad\square$

To finish the proof of Theorem 7.1(2), we proceed as follows. First, find a direct decomposition of G_j [55]. If any direct factor is contained in $Z(G_j)$, set this direct factor aside. Then by Lemma 7.5, the matrices M_{f_j} will be full rank, *even after modding out by coboundaries.* The latter fact implies that the rank of $\widetilde{M_{f_j}}$ is then the rank of M_{f_j} plus the dimension of $B_{prod}^2(Q, A)$; in other words, the matrices $\widetilde{M_{f_j}}$ are also of full rank.

Finally, we compute the coset of equivalences $\mathrm{CodeEq}(\widetilde{M_1}, \widetilde{M_2}) \subseteq S_m$ using Theorem 5.4. On the other hand, $\prod_i S_{\ell_i}$ induces an action on $[m]$, which contains the permutations we want. Thus we need to intersect $\mathrm{CodeEq}(M_1, M_2)$ with $\prod_i S_{\ell_i}$. By Proposition 7.7, if this intersection is nonempty, the algorithm returns "isomorphic." If the intersection is empty for all diagonals δ, the algorithm returns "nonisomorphic."

To analyze the running time, the outer loop depending on the diagonals is polynomially related to n. Both the applications of the LINEAR CODE EQUIVALENCE algorithm (Theorem 5.4), and the singly exponential time algorithm for COSET INTERSECTION [63] take time $c^m \le c^{\ell D^2}$ for some absolute constant c, where D is the maximum size of any of the T_i.

Computing the coset of isomorphisms. As in the other proof, the key is to compute the image of $\mathrm{Aut}(G_1)$ in $\mathrm{Aut}(A) \times \prod_{i \in [r]} \left(\mathrm{Aut}(Q_i)^{\ell_i} \rtimes S_{\ell_i} \right)$. For each diagonal $\delta \in \prod_{i \in [\ell]} \mathrm{Aut}(T_i)$, we compute the coset intersection of Proposition 7.7, which is a subcoset of $\prod_{i \in [r]} S_{\ell_i}$. This subcoset is given by a single element of $\prod_{i \in [r]} S_{\ell_i}$ together with a generating set for the image of $\mathrm{Aut}(G_1)$ in $\prod_{i \in [r]} S_{\ell_i}$. Given the diagonal δ and one of the generators σ of this subcoset, we compute an element $\alpha \in \mathrm{Aut}(A)$ such that $(\alpha, \delta, \sigma) \in \mathrm{Aut}(G_1)$ following the first half of the proof of Proposition 7.7. The only part unspecified is how to find Λ (in the notation of that proof); given δ and σ, this Λ is just an arbitrary element of a linear transporter, which is easily computed.

Finally, we need to compute a generating set for the $\mathrm{Aut}_0(G_1)$ (notation as in the first proof of this theorem). We compute such a generating set by enumerating all diagonals of $\prod \mathrm{Aut}(T_i)$ and computing generators for the linear stabilizers of the rowspan of the matrix $\widetilde{M_{f_1}}^{(\mathrm{id}, \delta, \mathrm{id})}$. $\qquad\square$

Given a group A, a *group code* of length n over A is a subgroup of A^n. Note that when $A = \mathbb{Z}_p$ is a prime field, group codes are the same as linear codes.

OBSERVATION 7.8. *If the coset of group code equivalences of length n over cyclic p-groups $\mathbb{Z}_{p^k}$ can be found in time $2^{O(n)}$ with arithmetic operations in $\mathbb{Z}_{p^k}$ at unit cost, then Theorem 7.1(2) also holds when $\mathrm{Rad}(G) = Z(G)$ is an arbitrary abelian group.*

Proof. Following the definition of $R^{(\mu)}$ in subsection 6.1.1, for a given $d \times n$ matrix M representing a 2-cochain over an abelian group $A = \prod_{i=1}^k \prod_{j=1}^{\ell_i} \mathbb{Z}_{p_i^{\mu_{i,j}}}$, let $M^{(p,\mu)}$ denote the subgroup of $\mathbb{Z}_{p^\mu}^n$ that is generated by the corresponding (possibly scaled/quotiented) rows $R^{(p,\mu)}$. That is, $M^{(p,\mu)}$ is the subgroup that is generated by the following: (a) the rows of M that correspond to $\mathbb{Z}_{p^\mu}$ factors of A; (b) $p^{\mu-\mu'}$ times the rows of M that correspond to $\mathbb{Z}_{p^{\mu'}}$ with $\mu' < \mu$; and (c) the rows of M that correspond to $\mathbb{Z}_{p^{\mu'}}$ with $\mu' > \mu$, taken modulo p^μ. (Rows corresponding to primes $q \ne p$ do not contribute to $M^{(p,\mu)}$.) In extending the argument of Theorem 7.1(2) to arbitrary central radicals, it is not difficult to show that the corresponding cocycles

are isomorphic if and only if there is a product-respecting automorphism of Q that is simultaneously a code equivalence for each $M^{(p,\mu)}$. Thus, to solve this we compute, for each p^μ that appears as a factor of $\mathrm{Rad}(G)$, the coset of code equivalences of the $M^{(p,\mu)}$. Intersecting these cosets then yields the desired permutations, and the rest of the proof proceeds as in the theorem. $\square$

Finally, we observe that the cohomological bound of Theorem 7.1 extends from elementary abelian G-modules to general abelian G-modules. Although we do not yet know how to prove Theorem 7.1 in the general abelian case, because we do not have the general abelian analogue of Proposition 6.10, we expect the following to be useful.

There are two possible ways to generalize the notion of "dimension" from G-modules over $\mathbb{F}_p$ to general finite G-modules, both of which are potentially useful for complexity analysis of algorithms for GPI; our result will cover both simultaneously. The first is the log of the cardinality of a finite G-module M, which we denote $\ell(M) := \log|M|$ (with the base of the logarithm unspecified, as the result will hold in any base); this generalizes dimension in that for any $\mathbb{F}_pG$-module M, $|M| = p^{\dim M}$. The second measure is the minimal number of generators, $d(M)$ (which equals dimension for $\mathbb{F}_p$-modules), which is also equal to the smallest d such that M is a quotient of the free abelian group $\mathbb{Z}^d$. The measure $d(M)$ is also closely related to the number of direct factors appearing in the primary decomposition of M; more precisely, $d(M)$ is at most the number of direct factors in the primary decomposition of M, and at least the number of p-primary direct factors in the primary decomposition of M, for any prime p. The ratio $d(H^2(G,M))/d(M)$ governs the number of columns that appear in the matrices representing 2-cohomology classes as in the first proof of Theorem 7.1(1).

LEMMA 7.9. *For any quasi-simple finite group G, and any finite $\mathbb{Z}G$-module M, $\ell(H^2(G,M)) \leq 17.5\ell(M)$ and $d(H^2(G,M)) \leq 17.5d(M)$.*

If the value 17.5 in [43, Thm. B] gets improved to a smaller constant c, then the 17.5 in this lemma also immediately can be replaced by c. (The authors of [43] state that they believe that with additional work the 17.5 could be replaced by 2.)

Proof. We give the proof for $\ell(\cdot)$; the same proof works mutatis mutandis for $d(\cdot)$.

M is a finite abelian group, so the Sylow p-subgroups of M are characteristic subgroups, and therefore are G-invariant. Let M_p denote the Sylow p-subgroup of M. Then $M \cong \bigoplus_{p||M|} M_p$ is a direct sum decomposition of $\mathbb{Z}G$-modules. As H^2 is additive in its second factor, we have $H^2(G,M) \cong \bigoplus_{p||M|} H^2(G, M_p)$. Thus, if we can show the result when M is an abelian p-group, we are done, for then we have $\ell(H^2(G,M)) = \sum_{p||M|} \ell(H^2(G, M_p)) \leq \sum_p 17.5\ell(M_p) = 17.5\ell(M)$.

So we now assume that M is a finite abelian p-group acted on by G. In any abelian p-group, the pth powers form a characteristic subgroup, which we denote pM since we are using additive notation. So pM is a G-submodule, albeit not in general a direct summand. Since H^2 is a left-exact covariant functor in its second argument, from the exact sequence $0 \to pM \hookrightarrow M \twoheadrightarrow M/pM \to 0$, we get a left exact sequence

$$0 \to H^2(G, pM) \hookrightarrow H^2(G,M) \to H^2(G, M/pM),$$

in which the final map need not be a surjection. However, even if it is not a surjection, by bounding the measures of the first and last term of this sequence, we still get a bound on the measure of the middle term $H^2(G,M)$.

Now, let e be the exponent of M (that is, $p^e M = 0$ but $p^{e-1}M \neq 0$). We proceed by induction on e. For $e = 1$, M is in fact an $\mathbb{F}_pG$-module, and we have that $\ell(H^2(G,M)) \leq 17.5\ell(M)$ by [43, Thm. B].

For $e > 1$, we have that $\ell(H^2(G, M/pM)) \leq 17.5\ell(M/pM)$ by [43, Thm. B], since M/pM is an $\mathbb{F}_pG$-module. Since the exponent of pM is strictly less than the exponent of M, we have, by induction, that $\ell(H^2(G, pM)) \leq 17.5\ell(pM)$. Thus we have that $\ell(H^2(G, M)) \leq \ell(H^2(G, pM)) + \ell(H^2(G, M/pM)) \leq 17.5(\ell(pM) + \ell(M/pM)) = 17.5\ell(M)$. □

8. Future directions. In this paper we developed an $n^{O(\log \log n)}$-time algorithm to test isomorphism of groups with central radicals, extending the results of [12] and beginning to resolve an open problem from [11]. We also developed an $n^{O(\log \log n)}$-time algorithm for groups with elementary abelian radical (not necessarily central), and polynomial-time algorithms for several prominent subclasses of central radical groups. The "difficult" cases—those where we do not yet know how to improve beyond $n^{O(\log \log n)}$—seem to be when the radical $\mathrm{Rad}(G)$ and the semisimple quotient $G/\mathrm{Rad}(G)$ are roughly of the same size—say both are of order $\sqrt{n}$—and $G/\mathrm{Rad}(G)$ is complicated (without this last condition, we handle such groups in Theorem 7.1; see Remark 7.2). Although a polynomial-time algorithm for the general case of central radicals remains open, we propose several directions for extending our work which we believe may now be within reach.

8.1. Abelian radical. A nearby next step is to extend our results to groups with general abelian radicals (not necessarily central nor elementary abelian).

Open Problem 8.1. Extend Theorem 6.11 to groups whose solvable radicals are general abelian. Ultimately, decide isomorphism of groups with abelian radicals in polynomial time.

We note that even with the recent algorithm for MODULE CYCLICITY for arbitrary finite modules (not necessarily over a field) [29], it is not immediately clear how to generalize the proof of Theorem 6.11 to the general abelian case. The issue is in the definition of the algebra U: In order to be closed under addition, we must allow the cocycles to get scaled arbitrarily. However, if an element of $\mathrm{End}(A)$ scales a cocycle by a nonunit in $\mathbb{Z}_{p^e}$, it is unclear how to proceed (since then we cannot argue that by rescaling the endomorphism by the inverse of the scalar, we get an automorphism, since the scalar has no inverse).

In both Theorems 6.1 and 6.11, we are currently only able to compute the coset of isomorphisms in the elementary abelian case. This could be generalized to the abelian case of Theorem 6.1 by resolving the following.

Open Problem 8.2. Extend Theorem 5.7 [24, 74] from matrix algebras over a field to finite $\mathbb{Z}$-algebras. That is, given a $\mathbb{Z}$-linear spanning set of a finite algebra over $\mathbb{Z}$, compute a generating set of the group of units in polynomial time.

8.2. The Babai–Beals filtration. The Babai–Beals filtration was defined and used in the context of algorithms for matrix groups [9, 10]—where the groups are given by a generating set of matrices, and the goal is algorithms which run in time polynomial in the input size, which can be polylogarithmic in $|G|$. In the context of GPI, it has also been used successfully in the polynomial-time algorithm for semisimple groups [11, 12].

The Babai–Beals filtration is the following chain of characteristic subgroups:

$$(9) \qquad 1 \leq \mathrm{Rad}(G) \leq \mathrm{Soc}^*(G) \leq \mathrm{Pker}(G) \leq G,$$

where $\mathrm{Rad}(G)$ is the solvable radical of G and $\mathrm{Soc}^*(G)$ is the subgroup such that $\mathrm{Soc}^*(G)/\mathrm{Rad}(G) = \mathrm{Soc}(G/\mathrm{Rad}(G))$. Note that the socle of the semisimple group

$G/\operatorname{Rad}(G)$ is a direct product of nonabelian simple groups. G then acts on this direct product by, among other things, permuting the factors. The final subgroup in the Babai–Beals filtration, $\operatorname{Pker}(G)$, consists of those $g \in G$ which do not permute the direct factors of $\operatorname{Soc}^*(G)/\operatorname{Rad}(G)$.

In Theorem 7.1 we make progress on the case of groups G with central radical which further satisfy $G = \operatorname{Soc}^*(G)$. It is then natural to consider groups with the next step of the Babai–Beals filtration, $G = \operatorname{Pker}(G)$. As a polynomial-time algorithm for isomorphism of *semisimple* groups G satisfying $G = \operatorname{Pker}(G)$ [11] was significantly simpler than the polynomial-time algorithm for general semisimple groups [12], we have hope that the following is achievable.

Open Problem 8.3. Extend Theorem 7.1 to groups with central radical which satisfy $G = \operatorname{Pker}(G)$.

8.3. The Cannon–Holt strategy. Cannon and Holt [26] suggest the following strategy for computing $\operatorname{Aut}(G)$ for a finite group G, as well as for isomorphism testing (see also [49], which contains improvements and new ideas related to this strategy). They consider the following chain of characteristic subgroups:

$$(10) \qquad 1 = N_r \trianglelefteq N_{r-1} \trianglelefteq \cdots \trianglelefteq N_1 = \operatorname{Rad}(G) \trianglelefteq G,$$

where the N_i refine the derived series of $\operatorname{Rad}(G)$ and each N_i/N_{i+1} is elementary abelian. More specifically, let $R^{(i)}$ be the terms of the derived series of $\operatorname{Rad}(G)$; thus $R^{(0)} = \operatorname{Rad}(G)$, $R^{(1)} = [\operatorname{Rad}(G), \operatorname{Rad}(G)]$, and so on. Then the quotients $R^{(i)}/R^{(i+1)}$ are abelian. From an abelian group we get a characteristic series as follows: First break the abelian group into abelian p-groups for each p, taking the primes in ascending order. Within an abelian p-group, we take the subgroup of pth powers as the largest characteristic subgroup in the series, and then recurse. These groups are then lifted from the abelian quotients $R^{(i)}/R^{(i+1)}$ back to subgroups of $\operatorname{Rad}(G)$ in the natural way. The algorithm proceeds by first computing $\operatorname{Aut}(G/N_1) = \operatorname{Aut}(G/\operatorname{Rad}(G))$, and then iteratively computing $\operatorname{Aut}(G/N_{i+1})$ from $\operatorname{Aut}(G/N_i)$.

This chain is convenient for describing known results in the Cayley table model: the case when $\operatorname{Rad}(G) = 1$ (equivalently $r = 1$) corresponds to the semisimple case, which can be solved in polynomial time [12]. When $G = \operatorname{Rad}(G)$ and $r = 2$, the case of $|N_2|$ and $|N_1/N_2|$ being coprime can be solved in polynomial time [16]. When $|N_2|$ and $|N_1/N_2|$ are not coprime, this includes the notorious case of p-groups of class 2. Finally, the present work considers a special case of $r = 2$, namely when $\operatorname{Rad}(G) = Z(G)$.

In light of [12], in the Cayley table model the second step in the Cannon–Holt strategy—to compute $\operatorname{Aut}(G/N_2)$ from $\operatorname{Aut}(G/\operatorname{Rad}(G))$—is equivalent to the special case of Open Problem 8.1 in which $\operatorname{Rad}(G)$ is elementary abelian, which we have solved in $n^{O(\log\log n)}$ time in Theorem 6.11.

However, even before getting polynomial-time algorithms for groups with general abelian radicals (Open Problem 8.1), it may be possible to give a reduction from the third step of the Cannon–Holt strategy to listing isomorphisms of two-step solvable groups. This is headed in the direction of a formal reduction from general group isomorphism to the solvable case. In a related vein, in Proposition 3.13 we showed how isomorphism of groups whose outer action on $\operatorname{Rad}(G)$ is trivial reduces to isomorphism of central-radical groups and isomorphism of solvable groups.

Open Problem 8.4. Extend Theorems 6.1 and 7.1 to groups whose radicals are two-step solvable, allowing access to an oracle for listing $\operatorname{Aut}(\operatorname{Rad}(G))$.

Appendix A. A gentle introduction to group extensions and cohomology. In this appendix we present a gentle introduction to group extensions and cohomology from the viewpoint of isomorphism testing, for readers encountering such notions for the first time. While this content exists in many textbooks, the typical textbook treatments often require reading several chapters first; our goal here is to cover the needed material in a self-contained, leisurely, and expository manner.

A.1. An intuitive starting point. Suppose we have a group G with a normal subgroup N, whose quotient G/N we call Q. We write this as $N \hookrightarrow G \twoheadrightarrow Q$ and refer to G as an *extension* of N by Q (some authors use the opposite convention, and call this as an extension of Q by N). Given N and Q, what additional data determines G up to isomorphism?

If G is the direct product $N \times Q$, then we essentially need no additional information.

A slightly more complicated kind of extension is captured by the semidirect product: In a direct product, G contains subgroups N and $\overline{Q} \cong Q$ such that $G = N\overline{Q} = \{nq : n \in N, q \in \overline{Q}\}$ and $N \cap \overline{Q} = 1$, and both N and $\overline{Q}$ are normal. As a consequence, every element of n commutes with every element of $\overline{Q}$. A semidirect product is defined similarly, except that $\overline{Q}$ need no longer be normal, and consequently the elements of $\overline{Q}$ need no longer commute with the elements of N. However, as N is normal, given any $q \in \overline{Q}$, $qNq^{-1} = N$, and thus the map $\varphi \colon Q \to \mathrm{Aut}(N)$ given by sending q to the map $(n \mapsto qnq^{-1})$ gives an action of $\overline{Q}$ on N. We denote this semidirect product by $N \rtimes_\varphi Q$. (When the action is trivial, we recover the direct product.) So to specify G from N and Q, we also need to specify something about this action.

However, consider the two groups $\mathbb{Z}_2 \times \mathbb{Z}_2$ and $\mathbb{Z}_4$, both as extensions of $\mathbb{Z}_2$ by $\mathbb{Z}_2$. (For $\mathbb{Z}_4$, written additively, we have $\{0, 2\} \cong \mathbb{Z}_2 \trianglelefteq \mathbb{Z}_4$, with the quotient $\mathbb{Z}_4/\mathbb{Z}_2$ again being a copy of $\mathbb{Z}_2$. The quotient map is the parity map.) However, these two extensions are not distinguished by any action of the quotient on the normal subgroup: In $\mathbb{Z}_2 \times \mathbb{Z}_2$ this action is trivial, as always happens for direct products; in $\mathbb{Z}_4$, since $\mathbb{Z}_4$ is abelian, conjugation by any element of $\mathbb{Z}_4$ is trivial, and thus the action of $Q \cong \mathbb{Z}_2$ on $N \cong \mathbb{Z}_2$ is also trivial. If we think of the group elements in both cases as specified by length-2 binary strings, the difference between the groups is that when adding in $\mathbb{Z}_4$, one must perform *carries*, whereas in $\mathbb{Z}_2 \times \mathbb{Z}_2$ one does not. Carrying is a simple example of nontrivial cohomology.

How can we capture this difference formally? Note that in the case of $\mathbb{Z}_4$, there is no choice of coset representatives of $\mathbb{Z}_2$ such that the coset representatives themselves form a subgroup; that is, there is no subgroup $\overline{Q}$ as in the discussion of the semidirect product above. It turns out that measuring the *failure* of the existence of such a subgroup is precisely what we need to capture the difference between $\mathbb{Z}_4$ and $\mathbb{Z}_2 \times \mathbb{Z}_2$; as we will see, in this case, this "measure of failure" will turn out to be described exactly by carrying.

Let $N = \mathbb{Z}_2$, $G = \mathbb{Z}_4$, and $Q = G/N = \mathbb{Z}_2$, and let $\pi \colon G \to Q$ denote the natural quotient map. Let $s \colon Q \to G$ denote any choice of coset representatives; that is, $\pi(s(q)) = q$ for all $q \in Q$. If there *were* a subgroup $\overline{Q}$ as above, then there would be some choice of s so that s is a homomorphism (and, in fact, it is not hard to see that such an s would be an isomorphism $Q \overset{\cong}{\to} \overline{Q}$). So we will measure the failure of s to be a homomorphism, and then consider this measure over all possible s.

For s to be a homomorphism is to say precisely that $s(p)s(q) = s(pq)$ for all $p, q \in Q$, or equivalently that $s(p)s(q)s(pq)^{-1} = 1$ for all $p, q \in Q$. We will use this expression to measure the failure of s to be a homomorphism, and say that the

failure of s to be a homomorphism is exactly captured by the function $f_s(p,q) := s(p)s(q)s(pq)^{-1}$. This function is identically 1 if and only if s is a homomorphism. Any f_s of this form is called a 2-*cocycle*, and all such 2-cocycles arise in this fashion. We then define the following.

DEFINITION A.1 (cohomology class, provisional definition). *The 2-cohomology class of an extension* $N \hookrightarrow G \twoheadrightarrow Q$ *is the collection*

$$[f_s] := \{f_s | s\colon Q \to G \text{ is a choice of coset representatives}\}.$$

Before understanding what this collection is, let us return to our example.

Example A.2. The map $\pi\colon \mathbb{Z}_4 \to \mathbb{Z}_2$ maps $\{0,2\} \mapsto 0$ and $\{1,3\} \mapsto 1$. So we have four choices for $s\colon \mathbb{Z}_2 \to \mathbb{Z}_4$:

$$s_0 \begin{cases} 0 & \mapsto 0, \\ 1 & \mapsto 1, \end{cases} \qquad s_1 \begin{cases} 0 & \mapsto 0, \\ 1 & \mapsto 3, \end{cases} \qquad s_2 \begin{cases} 0 & \mapsto 2, \\ 1 & \mapsto 1, \end{cases} \qquad s_3 \begin{cases} 0 & \mapsto 2, \\ 1 & \mapsto 3. \end{cases}$$

We calculate f_{s_0} for exposition and leave the calculation of the remaining f_{s_i} to the reader. Since we are writing $\mathbb{Z}_4$ and $\mathbb{Z}_2$ additively, our expression for f_s becomes $s(p) + s(q) - s(pq)$.

$$\begin{array}{rclclcl} f_{s_0}(0,0) & = & s_0(0) + s_0(0) - s_0(0+0) & = & s_0(0) & = & 0, \\ f_{s_0}(0,1) & = & s_0(0) + s_0(1) - s_0(0+1) & = & s_0(0) & = & 0, \\ f_{s_0}(1,0) & = & s_0(1) + s_0(0) - s_0(1+0) & = & s_0(0) & = & 0, \\ f_{s_0}(1,1) & = & s_0(1) + s_0(1) - s_0(1+1) & = & 1+1-0 & = & 2. \end{array}$$

There are two things to note. First, every $f_{s_0}(p,q)$ actually landed in the normal subgroup $N = \{0,2\}$; as we will see in a moment, this was no accident. Second, f_{s_0} is exactly the carry function: If we write the values in $\mathbb{Z}_4$ in binary, then $f_{s_0}(p,q)$ was 00 except when $p = q = 1$, when it was 10. This should help give some intuition for what a cohomology class "really is." It is easily verified that none of the f_{s_i} are trivial (with image 0), so this extension is not a semidirect product, as we already knew.

Note that the cohomology class depends not only on G, but also on the choice of N (see Example A.12), *and* on the choice of map $\pi\colon G \to Q$ (see Example A.8). We will give examples of both of these in the next section, after we have developed some more structure on cohomology classes, which will make the examples clearer. (At first, the reader may find the dependence on π surprising, since given N, we think of Q as being the same as G/N, for which there is the natural map $g \mapsto gN$. However, note that for any given map $\pi\colon G \to Q$ with kernel N, and any automorphism $\beta \in \mathrm{Aut}(Q)$, the map $\beta \circ \pi\colon G \to Q$ is also a homomorphism $G \to Q$ with kernel N, and it is this choice that is relevant.)

A.2. The group structure on cohomology. Defining the cohomology class as $\{f_s | s\colon Q \to G\}$ is a perfectly fine definition, but what is this collection? It turns out that, when N is abelian, this collection is a coset of a subgroup of some group, as we now explain. For the remainder of this appendix, N will be abelian, so we will denote N by A. In A we will use additive notation; although $A \leq G$ and G need not be abelian, this will turn out not to cause much confusion, and using this notation helps simplify things greatly.

Let $C^2(Q,A)$ denote the group of all functions (not necessarily homomorphisms) $Q \times Q \to A$; $C^2(Q,A)$ is an abelian group under pointwise addition: $(f + g)(p,q) :=$

$f(p,q) + g(p,q)$. Elements of $C^2(Q,A)$ are called "2-cochains of Q with coefficients in A." First, we show that any f_s as above is in fact a 2-cochain. To show this, we need merely show that $f_s(p,q) \in A$ for all $p,q \in Q$. Recalling that $A = \text{Ker}(\pi)$, we apply $\pi \colon G \to Q$ to $f_s(p,q)$:

$$\pi(f_s(p,q)) = \pi(s(p)s(q)s(pq)^{-1}) = \pi(s(p))\pi(s(q))\pi(s(pq)^{-1}) = pq(pq)^{-1} = \text{id}_Q,$$

and thus $f_s(p,q)$ is indeed in $\text{Ker}(\pi) = A$. So f_s is an element of the abelian group $C^2(Q,A)$.

However, not every 2-cochain arises in this manner; it turns out that those that do are a subgroup of $C^2(Q,A)$. To determine this subgroup, we need to determine the constraint that all f_s as above satisfy. One constraint comes from the fact that multiplication in G is associative; it will turn out that this one constraint suffices (that is, any 2-cochain satisfying this associativity-like constraint will be of the form f_s for some choice of coset representatives $s \colon Q \to G$).

To derive this associativity-like constraint, let us write the elements of G in terms of those of A and Q. Let $s \colon Q \to G$ be a choice of coset representatives; then every element of G can be written uniquely as $as(q)$ for some $a \in A, q \in Q$. Let us determine what the product in G looks like when we represent elements of G in this form:

$$\begin{aligned}
(as(p))(bs(q)) &= as(p)b(s(p)^{-1}s(p))s(q) \\
&= a(s(p)bs(p)^{-1})s(p)s(q) \\
&= a\theta_p(b)s(p)s(q) \\
&= a\theta_p(b)s(p)s(q)s(pq)^{-1}s(pq) \\
&= a\theta_p(b)f_s(p,q)s(pq).
\end{aligned}$$

In other words, if we use the notation $(a,q)_s$ to mean $as(q)$, then we have

(11) $$(a,p)_s \cdot (b,q)_s = (a + \theta_p(b) + f_s(p,q), pq)_s.$$

Exercise. Using (11), expand out the two expressions $((a,p)_s \cdot (b,q)_s) \cdot (c,r)_s$ and $(a,p)_s \cdot ((b,q)_s \cdot (c,r)_s)$ (note that the order of parentheses matters for the syntactic form of the final expressions; do not just assume associativity!).

Since multiplication in G is associative, the two expressions in the preceding exercise must be equal. Expanding them out and setting them equal, we find that f_s must satisfy the condition:

$$f_s(p,q) + f_s(pq,r) = \theta_p(f_s(q,r)) + f_s(p,qr) \qquad \text{(the 2-cocycle identity)}.$$

Reversing the above reasoning, we find that given an action $\theta \colon Q \to \text{Aut}(A)$, any 2-cochain $f \colon Q \times Q \to A$ which satisfies the 2-cocycle identity arises as f_s for some extension $A \hookrightarrow G \twoheadrightarrow Q$ and some choice of coset representatives $s \colon Q \to G$. Any 2-cochain f satisfying the 2-cocycle identity is called a 2-*cocycle* (with respect to the action θ). Since the 2-cocycle identity is easily seen to be $\mathbb{Z}$-linear—if f and g are both 2-cocycles, then so is their pointwise sum $f + g$—the 2-cocycles form a subgroup of the 2-cochains, which we denote $Z^2(Q,A,\theta)$ (unlike 2-cochains, the 2-cocycle identity depends on the action θ). Thus each cohomology class (with respect to θ) is a subset of $Z^2(Q,A,\theta)$.

Finally, let us determine what kind of subset a cohomology class is. Toward this end, suppose $s,t \colon Q \to G$ are two choices of coset representatives of Q in G.

What is the difference between f_s and f_t? As $s(q)$ and $t(q)$ lie in the same coset of A, there is a function $u\colon Q \to A$ such that $s(q) = u(q)t(q)$ for all $q \in Q$. Then $f_s(p,q) = f_t(p,q) + (u(p) + \theta_p(u(q)) - u(pq))$, where $\theta_p\colon A \to A$ is the automorphism given by conjugation by $p \in Q$. A 2-*coboundary* is a function of the form $b_u(p,q) := u(p) + \theta_p(u(q)) - u(pq)$ for any function $u\colon Q \to A$. Hence, if two 2-cocycles come from the same extension, they differ by a 2-coboundary.

Now, for a 2-coboundary of the form $f_s - f_t$, it is clear that it lies in $Z^2(Q, A, \theta)$, since both f_s and f_t are elements of this group. More generally, we have the following exercise.

Exercise. Show that any 2-coboundary satisfies the 2-cocycle identity. Show that if f_s is a 2-cocycle corresponding to the extension $A \hookrightarrow G \twoheadrightarrow Q$, and b_u is any 2-coboundary, then $f_s + b_u$ is another 2-cocycle corresponding to the same extension.

Thus, the 2-coboundaries are elements in $Z^2(Q, A, \theta)$. Finally, note that for any two functions $u, v\colon Q \to A$, we have $b_{u+v} = b_u + b_v$, so the 2-coboundaries in fact form a sub*group* of $Z^2(Q, A, \theta)$, which we denote $B^2(Q, A, \theta)$. Thus, we have finally arrived at the following definition.

DEFINITION A.3. *A 2-cohomology class is an element of the abelian group*

$$H^2(Q, A, \theta) := Z^2(Q, A, \theta)/B^2(Q, A, \theta),$$

called the second cohomology group of Q with coefficients in A, relative to the action θ.

We thus arrive at one of the central notions in this paper.

DEFINITION A.4. *For A an abelian group and Q any group, a pair (θ, f) of an action $\theta\colon Q \to \mathrm{Aut}(A)$ and a 2-cocycle $f\colon Q \times Q \to A$, $f \in Z^2(Q, A, \theta)$ is* extension data. *Two extension data for the pair (Q, A) are* equivalent *if they have the exact same action and if the two 2-cocycles are cohomologous (differ by a coboundary).*

Given an extension $A \hookrightarrow G \twoheadrightarrow Q$, the extension data for this particular extension are the action θ as defined above, and any 2-cocycle f_s for any section $s\colon Q \to G$. Note that extension data are nonunique, as we may choose any representative of the corresponding 2-cohomology class.

Two important special cases of extension data (θ, f) are as follows.

f is trivial (as 2-cohomology class). This implies that there exists $P \leq G$ such that $AP = G$ and $P \cap A = \mathrm{id}$, i.e., that G is the semidirect product $A \rtimes P$. Such P is called a *complement* of A in G, and the extension is called a *split* extension. In this case, an isomorphism test need only focus on one of the two aspects of GpI: EXTENSION DATA PSEUDO-CONGRUENCE simplifies to ACTION COMPATIBILITY.

θ is trivial. This implies that $A \leq Z(G)$, and the extension is called central. In this case, an isomorphism test need only focus on the other aspect of GpI: EXTENSION DATA PSEUDO-CONGRUENCE simplifies to COHOMOLOGY CLASS ISOMORPHISM.

Remark A.5. It is not difficult to test whether an input satisfies one of the above conditions; it is trivial to test whether an extension is central. We leave it as an exercise for the interested reader to devise an algorithm to test whether an extension is split when the normal subgroup is abelian. See [47, sect. 7.6.2] for a practical algorithm for the latter problem.

A.3. Equivalence of extensions and extension data. Now that we see there is a group structure on $H^2(Q, A, \theta)$, we return to give examples of the dependence on A and $\pi \colon G \to Q$. For this, we introduce a standard, slightly different viewpoint on the notion of "equivalence."

DEFINITION A.6. *Two extensions of A by Q, G_1 and G_2, are* equivalent *if there exists an isomorphism $\gamma \colon G_1 \to G_2$ such that in the diagram*

$$
\begin{array}{ccccc}
A & \xrightarrow{\ \iota_1\ } & G_1 & \xrightarrow{\ \pi_1\ } & Q \\
\Big\| & & \cong \Big\downarrow \gamma & & \Big\| \\
A & \xrightarrow{\ \iota_2\ } & G_2 & \xrightarrow{\ \pi_2\ } & Q
\end{array}
$$

we have $\gamma\iota_1 = \iota_2$ and $\pi_1 = \pi_2\gamma$, in which case we say that the diagram commutes. *The vertical double lines are stretched out, rotated equality signs; that is, they denote the identity map.*

It is classical, going back to Hölder and Schreier, that two extensions have equivalent extension data (Definition A.4) if and only if the extensions are equivalent according to Definition A.6. The explicit connection with cohomology was developed by Eilenberg and MacLane [33].

THEOREM A.7 (see [73, Chap. 11] and Eilenberg–MacLane [33]). *There is a bijection between equivalence classes of extensions of A by Q with action θ, and elements of the group $H^2(Q, A, \theta)$.*

Now we give the promised example of the dependence of the cohomology class on the choice of map $\pi \colon G \to Q$. This example also serves a secondary purpose, as an example of two extensions where $G_1 \cong G_2$ but the extensions are not equivalent according to the above definition(s).

Example A.8 (isomorphic groups from nonequivalent extensions). $\mathbb{Z}_9$ can be viewed as an extension of $\mathbb{Z}_3$ by $\mathbb{Z}_3$ in two ways. First, $\mathbb{Z}_3 \hookrightarrow \mathbb{Z}_9$ by sending 1 to 3. Then define $\pi_i(1) = i$ for $i \in [2]$. To see that π_1 and π_2 yield nonequivalent extensions, consider any $\phi \in \mathrm{Aut}(\mathbb{Z}_9)$. Let $k = \phi(1)$; since ϕ is an automorphism, k is one of $\{1, 2, 4, 5, 7, 8\}$. In order for ϕ to be an equivalence of extensions, it must induce the identity on $\langle 3 \rangle$, implying that $3 = \phi(3) = 3\phi(1) = 3k \mod 9$, and thus $k \mod 3$ must be $1 \mod 3$. On the other hand, to be an equivalence of extensions, we must also have $1 = \pi_1(1) = \pi_2(\phi(1)) = \pi_2(k) = 2k \mod 3$, that is $k = 2 \mod 3$, giving a contradiction.

Given the preceding example, the question then becomes how to leverage cohomology for isomorphism testing (since what we will ultimately care about is the isomorphism class of G, and not merely the equivalence class of any given extension with G as its total group). This is what we tackle next.

A.4. Main lemma for abelian characteristic subgroups. Recall that a characteristic subgroup is a subgroup invariant under all automorphisms. The analogous notion for isomorphisms (rather than automorphisms) is a function $\mathcal{S}$ that assigns to each group G a subgroup $\mathcal{S}(G) \leq G$ such that any isomorphism $\varphi \colon G_1 \to G_2$ restricts to an isomorphism $\varphi|_{\mathcal{S}(G_1)} \colon \mathcal{S}(G_1) \to \mathcal{S}(G_2)$. We call such a function a *characteristic subgroup function*. Note that if $G_1 = G_2$, this says that $\mathcal{S}(G_1)$ is sent to itself by every automorphism of G_1, that is, $\mathcal{S}(G_1)$ is a characteristic subgroup of G_1. Most natural characteristic subgroups encountered are characteristic subgroup func-

tions, for example the center $Z(G)$, the commutator subgroup $[G, G]$, or the radical $\mathrm{Rad}(G)$.

Let $\mathcal{S}$ denote a fixed characteristic subgroup function, and suppose we are given two groups G_1, G_2 such that $\mathcal{S}(G_1)$ and $\mathcal{S}(G_2)$ are both abelian. To determine how to use cohomology in isomorphism testing, we first examine the consequences of an isomorphism $G_1 \cong G_2$, as an exercise in reverse engineering. Let $\gamma : G_1 \to G_2$ be an isomorphism. By the definition of characteristic subgroup function, $\gamma(\mathcal{S}(G_1)) = \mathcal{S}(G_2)$, so γ induces $\gamma_1 : \mathcal{S}(G_1) \to \mathcal{S}(G_2)$ and $\gamma_2 : G_1/\mathcal{S}(G_1) \to G_2/\mathcal{S}(G_2)$. Using these isomorphisms, we identify $A = \mathcal{S}(G_2) = \mathcal{S}(G_2)^{\gamma_1}$ and $Q = G_2/\mathcal{S}(G_2) = (G_2/\mathcal{S}(G_2))^{\gamma_2}$. Let (θ_i, f_i) be the extension data of $A \hookrightarrow G_i \twoheadrightarrow Q$, where $\theta_i : Q \to \mathrm{Aut}(A)$ and $f_i \in Z^2(Q, A, \theta_i)$. As we have identified $A = \mathcal{S}(G_1) = \mathcal{S}(G_2)$ and $Q = G_1/\mathcal{S}(G_1) = G_2/\mathcal{S}(G_2)$, γ induces some $\alpha \in \mathrm{Aut}(A)$ and $\beta \in \mathrm{Aut}(Q)$. We write $\theta_{i,q}$ as the shorthand for $\theta_i(q)$ for $i = 1, 2$ and $q \in Q$. It can then be verified that for $q \in Q$ and $a \in A$,

$$\text{(12)} \qquad \theta_{1,q}(a) = \alpha^{-1}(\theta_{2,\beta(q)}(\alpha(a))) =: \theta_2^{(\alpha,\beta)}(q)(a),$$

and we record this as $\theta_1 = \theta_2^{(\alpha,\beta)}$, where $\theta_2^{(\alpha,\beta)}$ is defined as above.

It can be similarly verified that $[f_1] = [f_2^{(\alpha,\beta)}]$ as classes in $H^2(Q, A, \theta_1)$, where $f_2^{(\alpha,\beta)}(p, q) := \alpha^{-1}(f_2(\beta(p), \beta(q)))$ for all $p, q \in Q$. In other words, we have

$$\text{(13)} \qquad f_1(p, q) = \alpha^{-1}(f_2(\beta(p), \beta(q))) + b_u(p, q)$$

for some 2-coboundary $b_u \in B^2(Q, A, \theta_1)$. Note that (12) ensures $f_2^{(\alpha,\beta)}$ is indeed a 2-cocycle relative to θ_1, i.e., in $Z^2(Q, A, \theta_1)$. This discussion leads to the following definition.

DEFINITION A.9. *Let A be an abelian group and Q any group, and let (θ_1, f_1) and (θ_2, f_2) be two extension data for A-by-Q. Then the extension data are* pseudo-congruent[8] *if there exists $(\alpha, \beta) \in \mathrm{Aut}(A) \times \mathrm{Aut}(Q)$ such that $\theta_1 = \theta_2^{(\alpha,\beta)}$ and $[f_1] = [f_2^{(\alpha,\beta)}]$, that is, (12) and (13) hold. In this case we write $(\theta_1, f_1) \cong (\theta_2, f_2)$.*

LEMMA A.10 (Main Lemma, abelian case). *Let $\mathcal{S}$ be a characteristic subgroup function. Given two finite groups G_1 and G_2, suppose $\mathcal{S}(G_1)$ and $\mathcal{S}(G_2)$ are abelian. Then $G_1 \cong G_2$ if and only if both of the following conditions hold:*

1. *$\mathcal{S}(G_1) \cong \mathcal{S}(G_2)$ (which we denote by A) and $G_1/\mathcal{S}(G_1) \cong G_2/\mathcal{S}(G_2)$ (which we denote by Q).*
2. *$(\theta_1, f_1) \cong (\theta_2, f_2)$, where (θ_i, f_i) is the extension data of the extensions $A \hookrightarrow G_i \twoheadrightarrow Q$.*

Proof. The above discussion shows the "only if" direction. For the other direction, suppose we are given an abelian group A, a group Q, an action $\theta : Q \to \mathrm{Aut}(A)$, and a 2-cocycle $f : Q \times Q \to A$, $f \in Z^2(Q, A, \theta)$. We shall need the following procedure of Eilenberg and MacLane [33] that takes A, Q, θ, and f as input, and outputs a group H as an extension of A by Q with extension data (θ, f). We refer to this as the *standard reconstruction procedure*. The set of group elements of H is $A \times Q$. For $(a, p), (b, q) \in A \times Q$, the group operation $\circ_H$ is defined as

$$(a, p) \circ_H (b, q) = (a + \theta_p(b) + f(p, q), pq).$$

---

[8]See Footnote 3 on page 1161.

A simple but tedious calculation verifies that $A \hookrightarrow H \twoheadrightarrow Q$ is an extension with extension data (θ, f).

Getting back to our problem, from $(\theta_1, f_1) \cong (\theta_2, f_2)$, we can choose appropriate sections $s_i \colon Q \to G_i$ such that the corresponding 2-cocycles satisfy $f_1 = f_2^{(\alpha, \beta)}$ in $Z^2(Q, A, \theta_1)$. Note that as $\theta_1 = \theta_2^{(\alpha, \beta)}$, $f_2^{(\alpha, \beta)} \in Z^2(Q, A, \theta_1)$. Now apply the standard reconstruction procedure to (θ_i, f_i) to get $H_i \cong G_i$ (isomorphism follows from Eilenberg and MacLane [32]). It is then straightforward to verify that the bijection $\gamma \colon H_1 \to H_2$ defined by $\gamma((a, p)) = (\alpha(a), \beta(p))$ is in fact an isomorphism. □

A.4.1. Pseudo-congruence of extensions and extension data. As with the standard concept of equivalence, the standard concept of pseudo-congruence applies directly to group extensions themselves, rather than extension data as in our definitions. We use our definitions because the standard definitions seem to presuppose that the total groups are isomorphic, whereas in our setting the whole goal is to determine whether this is the case. However, we show in this section that the definitions are in fact equivalent (which is closely related to the Main Lemma 3.7). We present the standard definition here as it has more intuitive appeal and we believe it makes some discussions in the paper clearer, for example the proof of Theorem 7.1.

Throughout this section, A denotes an abelian (normal sub)group. Nonabelian normal subgroups are handled in section 3.

DEFINITION A.11. *Two extensions $A \hookrightarrow G_i \twoheadrightarrow Q$ $(i = 1, 2)$ of A by Q are* pseudo-congruent *if there is an isomorphism $\gamma \colon G_1 \to G_2$ such that $\gamma(A) = A$. In particular, γ induces automorphisms $\alpha \in \mathrm{Aut}(A)$ and $\beta \in \mathrm{Aut}(Q)$.*

Pictorially, G_1 and G_2 are pseudo-congruent as extensions if there exist $\alpha \in \mathrm{Aut}(A)$, $\beta \in \mathrm{Aut}(Q)$, and $\gamma \in \mathrm{Iso}(G_1, G_2)$ such that the following diagram commutes:

$$
\begin{array}{ccccc}
A & \overset{\iota_1}{\hookrightarrow} & G_1 & \overset{\pi_1}{\twoheadrightarrow} & Q \\
{\cong}\downarrow{\alpha} & & {\cong}\downarrow{\gamma} & & {\cong}\downarrow{\beta} \\
A & \overset{\iota_2}{\hookrightarrow} & G_2 & \overset{\pi_2}{\twoheadrightarrow} & Q
\end{array}
$$

where ι_i is the injective homomorphism from A to G_i and π_i is the surjective homomorphism from G_i to Q with $\mathrm{Ker}(\pi_i) = \mathrm{Im}(\iota_i)$.

It is possible for the total groups G_1 and G_2 to be isomorphic without the extensions being pseudo-congruent. The following example was provided by Vipul Naik [66]. This, finally, also serves as an example of the dependence on the choice of A.

Example A.12 (isomorphic groups from non–pseudo-congruent extensions). $A = \mathbb{Z}_{p^2} \times \mathbb{Z}_p \times \mathbb{Z}_p$, $Q = \mathbb{Z}_{p^2} \times \mathbb{Z}_p$, $G = \mathbb{Z}_{p^3} \times \mathbb{Z}_{p^2} \times \mathbb{Z}_p \times \mathbb{Z}_p$. In one extension, $\iota_1(a, b, c) = (pa, 0, b, c)$ and in the other $\iota_2(a, b, c) = (pa, pb, a \pmod p, c)$. To see that there is no automorphism of G sending $\mathrm{Im}\, \iota_1$ to $\mathrm{Im}\, \iota_2$—and hence that the two extensions are not pseudo-congruent—note that $\mathrm{Im}\, \iota_1$ contains elements that are p times an element of order p^3 in G, but $\mathrm{Im}\, \iota_2$ contains no such elements.

Despite the fact that the usual Definition A.11 presupposes that the total groups are isomorphic, in fact it is equivalent to our Definition A.9. The isomorphism of the total groups follows for free from pseudo-congruence of the extension data.

LEMMA A.13. *Definitions A.9 and A.11 are equivalent. In detail, let $A \hookrightarrow G_i \twoheadrightarrow Q$ $(i = 1, 2)$ be extensions of A by Q, and let (θ_i, f_i) be the corresponding extension data. Then G_1 and G_2 are pseudo-congruent as extensions of A by Q if and only if*

$(\theta_1, f_1) \cong (\theta_2, f_2)$.

Proof. Suppose that the extensions are pseudo-congruent (Definition A.11), and let $\gamma \in \mathrm{Iso}(G_1, G_2)$, $\alpha \in \mathrm{Aut}(A)$, $\beta \in \mathrm{Aut}(Q)$ be as in Definition A.11. It is readily verified that $\theta_1 = \theta_2^{(\alpha,\beta)}$ and $[f_1] = [f_2^{(\alpha,\beta)}]$, that is, that the extension data are pseudo-congruent under Definition A.9.

Conversely, suppose the extension data are pseudo-congruent (Definition A.9). Then the isomorphism γ constructed in the proof of the abelian version of the Main Lemma (Lemma 2.3) satisfies the conditions of Definition A.11.   □

A.4.2. Some algorithmic problems arising from special cases of pseudo-congruence. We describe two special cases of pseudo-congruence of extensions, explain the algorithmic problems corresponding to them, and indicate some of the solutions. The first one was discussed in [72]. Consider the case when a pseudo-congruence γ induces the identity map on Q as follows:

$$
\begin{array}{ccccc}
A & \overset{\iota_1}{\hookrightarrow} & G_1 & \overset{\pi_1}{\twoheadrightarrow} & Q \\
\cong \downarrow \alpha & & \cong \downarrow \gamma & & \| \\
A & \underset{\iota_2}{\hookrightarrow} & G_2 & \underset{\pi_2}{\twoheadrightarrow} & Q
\end{array}
$$

This corresponds to the algorithmic setting when enumerating $\mathrm{Aut}(Q)$ is allowed (section 6), as after fixing some $\beta \in \mathrm{Aut}(Q)$ we are reduced to looking for α such that G_1 and G_2 are pseudo-congruent via (α, β). If the extension is split and $A \cong \mathbb{Z}_p^k$ is elementary abelian, this problem reduces to MODULE ISOMORPHISM: The action of each $q \in Q$ can be expressed as a nonsingular matrix in $\mathrm{GL}(k, p)$. So suppose $Q = \{q_1, \ldots, q_s\}$, and in G_j the conjugation action of Q is written as $\{M(j, i) \mid M(j, i) \in \mathrm{GL}(k, p)\}$, where $M(j, i)$ denotes the action of q_i on A in G_j. The problem of ACTION COMPATIBILITY then reduces to determining whether there exists $T \in \mathrm{GL}(k, p)$ such that $TM(1, i) = M(2, i)T$ for every $i \in [\ell]$. This is a special case of MODULE ISOMORPHISM, which admits deterministic polynomial-time algorithms [28, 22, 52]. At the other extreme, when the extension is central, Theorem 6.1 solves COHOMOLOGY CLASS ISOMORPHISM. At present it is not clear to us how to combine these two procedures to solve EXTENSION DATA PSEUDO-CONGRUENCE as a whole, beyond the elementary abelian case (Theorem 6.11); see subsection 8.1 for a discussion of the difficulties involved.

On the other hand, consider the case in which a pseudo-congruence γ induces the identity map on A, but not necessarily on Q:

$$
\begin{array}{ccccc}
A & \overset{\iota_1}{\hookrightarrow} & G_1 & \overset{\pi_1}{\twoheadrightarrow} & Q \\
\| & & \cong \downarrow \gamma & & \cong \downarrow \beta \\
A & \underset{\iota_2}{\hookrightarrow} & G_2 & \underset{\pi_2}{\twoheadrightarrow} & Q
\end{array}
$$

The isomorphism problem here corresponds to the algorithmic setting when $\mathrm{Aut}(A)$ is enumerable, and our goal is to find β such that G_1 and G_2 are pseudo-congruent. Proposition 7.6 and Remark 3.14 fall into this setting. Recently, Ivanyos and the second author [53] presented a randomized efficient algorithm to decide whether two alternating bilinear maps over a finite field of odd characteristic are isometric or not; given the connections between p-groups of class 2 and exponent p and alternating bilinear maps (cf. [89, sect. 3.4]), this amounts to solving the problem of finding β as

above. That result and the techniques therein are inspired by the work of Brooksbank and Wilson [25], who presented an efficient algorithm to compute the isometry group of a single alternating bilinear map (the automorphism version of the above problem). These two works [25, 53] together suggest that we can compute the coset of isometries between two alternating bilinear maps.

A.5. Some cohomological lemmas. Recall that a group G is perfect if $G = [G, G]$; as $[G, G]$ is always a normal subgroup, all nonabelian simple groups are perfect.

PROPOSITION 7.3 (RESTATED). *Suppose $T_1, \ldots, T_\ell$ are perfect groups. Let A be an abelian group, and $Q = \prod_{i \in [\ell]} T_i$. Let G be a group with $Z(G) = A$ and $G/Z(G) = Q$. Let $U_i = AT_i$ be the full preimage of T_i under the natural map $G \to G/Z(G)$. Then for $i, j \in [\ell]$, $i \neq j$, $[U_i, U_j] = 1$.*

Proof. Let $\pi : G \to G/Z(G)$ be the natural projection. Note that $Q = G/Z(G)$ and T_i's are direct factors of Q. For $i \in [\ell]$, define V_i to be the smallest normal subgroup of G such that $\pi(V_i) = T_i$. Then $U_i = V_i Z(G)$, and for $i \neq j$, $[U_i, U_j] = \mathrm{id}$ if and only if $[V_i, V_j] = \mathrm{id}$.

As T_i is perfect, $\pi([V_i, V_i]) = \pi(V_i)$. Because of minimality of V_i, $V_i = [V_i, V_i]$. For $i \neq j$, $T_i \cap T_j = \mathrm{id}$; thus $[\pi(V_i), \pi(V_j)] = [T_i, T_j] = \mathrm{id}$ in Q, which implies that $[V_i, V_j] \subseteq Z(G)$. Now we have (1) $[[V_i, V_j], V_j] \subseteq [Z(G), V_j] = \mathrm{id}$; (2) $[[V_j, V_i], V_j] = \mathrm{id}$ as $[V_i, V_j] = [V_j, V_i]$. Then Hall's three subgroup lemma [83, Chap. 4, Prop. 1.9] gives that $[[V_j, V_j], V_i] = \mathrm{id}$. Finally noting that $V_j = [V_j, V_j]$, we have $[V_j, V_i] = \mathrm{id}$. $\square$

LEMMA 7.5 (RESTATED). *Let $A' \times A'' \hookrightarrow G \twoheadrightarrow Q$ be a central extension of $A' \times A''$ by Q. Let $p_{A'} : A' \times A'' \to A'$ be the projection onto A' along A''. If there is a 2-cocycle $f : Q \times Q \to A' \times A''$ such that $p_{A'} \circ f : Q \times Q \to A'$ is a 2-coboundary, then G is isomorphic (even equivalent as an extension of $A' \times A''$ by Q) to the direct product $A' \times (G/A')$.*

Furthermore, given the Cayley table of G, A' can be computed in polynomial time using linear algebra over abelian groups.

Proof. We prove directly that $A' \trianglelefteq G$, exhibit a complement of A' in G, and show that this complement is normal. At the end we show how to compute A' using linear algebra.

We may assume without loss of generality that the image of f lies entirely within A''. For if not, then we may add the 2-coboundary $p_{A'} \circ f : Q \times Q \to A' \hookrightarrow A' \times A''$ to f to get an equivalent 2-cocycle satisfying this condition. Similarly, we may assume that f is normalized so that $f(1, q) = f(q, 1) = 0$ for all $q \in Q$.

We construct an extension equivalent to G from the cocycle f in the usual way: the elements are $A' \times A'' \times Q$ as a set, with multiplication given by (writing A' and A'' additively)

$$(a_1, a_1', q_1)(a_2, a_2', q_2) = (a_1 + a_2, a_1' + a_2' + f(q_1, q_2), q_1 q_2)$$

since the image of f lies entirely in A''. We also have $(a, a', q)^{-1} = (-a, -a' - f(q, q^{-1}), q^{-1})$.

A' is normal:

$$
\begin{aligned}
(a, a', q)^{-1}(a_0, 1, 1)(a, a', q) &= (-a, -a' - f(q, q^{-1}), q^{-1})(a_0 + a, a', q) \text{ (since } f(1, q) = 0) \\
&= (-a + a_0 + a, -a' - f(q, q^{-1}) + a' + f(q, q^{-1}), qq^{-1}) \\
&= (a_0, 0, \mathrm{id}_Q).
\end{aligned}
$$

A' has a normal complement: As the image of f lies entirely in A'', it is readily verified that elements of the form $(0, a', q)$ are closed under product and hence form a subgroup of G which is isomorphic to G/A' and intersects A' only in the identity. Moreover, this subgroup is normal. For consider conjugating one of its elements by an arbitrary element of G: $(-a, -a' - f(q, q^{-1}), q^{-1})(0, a'_0, q_0)(a, a', q)$. From the multiplication rule above, it is clear that the first coordinate of this product is just the sum of the first coordinates of the three factors—namely, zero—whatever the second and third coordinates are.

Finally, we show how to compute A' from the Cayley table for G using linear algebra over abelian groups. We give the proof in the case that $Z(G) = \mathbb{Z}_p^k$ is elementary abelian; the general case uses the same ideas as in subsection 6.1.1. First compute $Z(G)$ (which is $A' \times A''$, but we do not yet know this decomposition of $Z(G)$; we are only promised it exists) and $Q = G/Z(G)$. Choose any set-theoretic section $s \colon Q \to G$ and compute the corresponding cocycle $f := f_s$. Let M_f be the $k \times |Q|^2$ $\mathbb{Z}_p$-matrix corresponding to f as in section 6. We may view M_f as a $\mathbb{Z}_p$-linear map from $Z(G) = \mathbb{Z}_p^k$ to $\mathbb{Z}_p^{Q \times Q}$. As in Proposition 6.5, we may compute a basis of $B^2(Q, Z(G))$ that is a direct sum of bases for $B^2(Q, \mathbb{Z}_p)$, one copy for each row of M_f. The maximal A' satisfying the conditions of the theorem is then the inverse image of $B^2(Q, \mathbb{Z}_p)$ under this map. Computing the inverse image of $B^2(Q, \mathbb{Z}_p)$ under the map $M_f^T \colon Z(G) \to \mathbb{Z}_p^{Q \times Q}$ is then just linear algebra over $\mathbb{Z}_p$. $\square$

Appendix B. Generalized Fitting subgroups of groups with central radical.

DEFINITION B.1. *A group G is quasisimple if $G = [G, G]$ and $G/Z(G)$ is a nonabelian simple group. G is m-quasisimple if $G = [G, G]$ and $G/Z(G)$ is a direct product of nonabelian simple groups.*

Our m-quasisimple groups are Suzuki's "semisimple groups" [83, p. 446]; we cannot use Suzuki's terminology as we have used "semisimple" for groups with no abelian normal subgroups. Although central-radical groups G in which $G/Z(G)$ is a direct product of nonabelian simple groups need not be m-quasisimple groups (as G need not be perfect), the difference is small.

PROPOSITION B.2 ([83, Ch. 6, corollary to Thm. 6.4]). *Let G be a group such that $G/Z(G)$ is a direct product of nonabelian simple groups. Then $G = Z(G)[G, G]$ and $[G, G]$ is an m-quasisimple group.*

m-quasisimple groups are crucial for defining the generalized Fitting subgroups.

PROPOSITION B.3 (see [83]). *Let H and K be m-quasisimple normal subgroups of G; then HK is m-quasisimple.*

This motivates the following definition. Recall that the Fitting subgroup $F(G)$ of G is the maximal nilpotent normal subgroup of G.

DEFINITION B.4. *Let G be a group. The layer $E(G)$ of a group G, is the maximal m-quasisimple normal subgroup of G. The generalized Fitting subgroup $F^*(G)$ of G is $E(G)F(G)$.*

PROPOSITION B.5. *For any group G, if $\mathrm{Rad}(G) = Z(G)$, then $\mathrm{Soc}^*(G) = F^*(G)$.*

Proof. As $\mathrm{Rad}(G) = Z(G)$, $F(G) = Z(G)$. Let $D = [\mathrm{Soc}^*(G), \mathrm{Soc}^*(G)]$. Thus D is m-quasisimple and $\mathrm{Soc}^*(G) = Z(G)D = F(G)D$ (see [83, Ch. 6, corollary to Thm. 6.4]). Thus $D \subseteq E(G)$ and $\mathrm{Soc}^*(G) \subseteq F^*(G)$.

To show $\mathrm{Soc}^*(G) \supseteq F^*(G)$, for the purpose of contradiction, suppose $Z(G)D =$

$\mathrm{Soc}^*(G) \subsetneq F^*(G) = Z(G)E(G)$. Consider the decomposition of $E(G)$ into quasisimple groups [83, Ch. 6, Def. 6.8] as $Q_1 \cdot \ldots \cdot Q_d$, where $\cdot$ denotes central product, and Q_i is subnormal in G. Without loss of generality assume $Z(G)Q_1 \nsubseteq Z(G)D$. As Q_1 is subnormal in G, $G/Z(G)$ necessarily has $Q_1 Z(G)/Z(G)$ as a subnormal group, contained in some minimal normal group $N/Z(G) \lhd G/Z(G)$ [51, Lem. 9.17]. By assumption, $Q_1 Z(G)/Z(G)$ is not contained in $DZ(G)/Z(G) = \mathrm{Soc}^*(G)/Z(G) = \mathrm{Soc}(G/Z(G))$, and thus $N/Z(G)$ is a minimal normal subgroup not contained in $\mathrm{Soc}(G/Z(G))$, contradicting the definition of the socle.   □

Appendix C. Relationship with results on practical algorithms.   It is not surprising that when it comes to the group isomorphism problem, theoretical computer scientists and computational group theorists often leverage the same underlying structure of the relevant group classes, even when they consider groups as being input by different kinds of data structures. Here we discuss the relationship between these two sets of results. A general reference for computational group theory (CGT) is the handbook [47]; algorithms in CGT are often implemented in Magma [21] and/or GAP [37].

One major achievement in CGT, as reflected in the Magma and GAP as well as in [47], is to implement a huge collection of practical routines to work with groups, when the groups are represented by generating sets of permutations or matrices, or by polycyclic presentations (for solvable groups). In general, these encodings are of size polylogarithmic in $|G|$, so in that context even a provable worst-case guaranteed running time of $O(|G|)$ is usually impractical. However, in order to achieve better running times, practically fast methods are often used, even without guarantees on their running time. In the general setting, the best-known worst-case guarantee on the running time of these practical algorithms is $|G|^{O(\log|G|)}$, even when the group is given succinctly [88]. Thus, improving the state of the art on large group classes to anything less than $|G|^{\Theta(\log|G|)}$ represents not only an asymptotic, worst-case improvement in the complexity, but potentially a practical improvement in many cases as well. Note that although we sometimes phrase results in terms of the Cayley table model, essentially none of our results depend on this; it just means that running times that depend on $|G|$ are then also being expressed in terms of their dependence on the input size.

Regarding isomorphism testing algorithms in CGT, other than [26] which is mentioned in subsection 8.3, some notable works include [68, 18, 60], and many of the algorithms in this area are summarized in the theses [80, 49]. Very often isomorphism testing arises as a subroutine in the construction of all finite groups up to a certain order (up to isomorphism), as in [84, 18, 19]. Recently, Wilson et al. have produced several results related to isomorphism of p-groups (sometimes reformulated in the context of Hermitian bilinear maps) in [89, 90, 60, 25, 23], including worst-case guarantees (and sometimes even worst-case guarantees that are polynomial in the size of the succinct input!). The structure they are uncovering in p-groups is thus also notable from the complexity perspective, and there is likely more to be discovered in this direction.

Although the two communities often leverage the same structure that is present in various classes of groups, the worst-case guarantees often require further structural results on the group classes considered, in addition to further algorithmic results. For example, Besche and Eick had already considered the group classes in Le Gall's work [58] (the algorithm in [18, Fig. 4]), but Le Gall's work was the first to prove a polynomial-time upper bound on an isomorphism algorithm for groups of the form

$A \rtimes \mathbb{Z}_p$ with A abelian and $p \nmid |A|$. A necessary ingredient in Le Gall's work is a detailed understanding of automorphism groups of abelian groups traced back to Ranum [71] (note that the citation here is to *19*07), which was not needed in the practical setting of Besche and Eick. Another example is the polynomial-time algorithm for semisimple groups [12], where a similar situation is described at the end of that paper, comparing it with the practical work of Cannon and Holt [26]. For example, the algorithm in [12] required bounds on the orders of the transitive permutation groups other than S_n and A_n.

Relations to the present work. As mentioned above, our choice to focus on groups with central radicals is partially motivated by the strategy of Cannon and Holt [26]. Another work of particular relevance is [18]. There Besche and Eick considered construction of finite groups and proposed three heuristics. To support one heuristic, they proposed the concept of "strong isomorphism" of groups, which can be viewed as a special case of our Main Lemma 3.7 in their setting. We use the same structural results to support the approach, but as our goal is worst-case running time upper bounds of $\mathrm{poly}(|G|)$, we have more freedom to handle the 2-cohomology classes directly, as in Theorem 6.1; we also need Lemma 7.4 and Theorem 5.3 from [43], which in turn allow us to apply algorithms for LINEAR CODE EQUIVALENCE and COSET INTERSECTION as in Theorem 7.1. These ingredients are not present in [26] or [18], as they were not needed to get practical algorithms. However, it is nonetheless possible that taking advantage of these ingredients, in combination with the practically fast approaches from [26, 18], could lead to further practical improvements.

REFERENCES

[1] V. ARVIND AND J. TORÁN, *Solvable group isomorphism is (almost) in* NP ∩ coNP, ACM Trans. Comput. Theory, 2 (2011), 4, https://doi.org/10.1145/1944857.1944859.

[2] M. ASCHBACHER, *Finite Group Theory*, 2nd ed., Cambridge Stud. Adv. Math. 10, Cambridge University Press, Cambridge, UK, 2000.

[3] M. ASCHBACHER AND R. GURALNICK, *Some applications of the first cohomology group*, J. Algebra, 90 (1984), pp. 446–460, https://doi.org/10.1016/0021-8693(84)90183-2.

[4] L. BABAI, *Permutation Groups, Coherent Configurations, and Graph Isomorphism*, D.Sc. Thesis, Hungarian Academy of Sciences, Budapest, Hungary, 1983 (in Hungarian).

[5] L. BABAI, *Automorphism groups, isomorphism, reconstruction*, in Handbook of Combinatorics (Vol. 2), R. L. Graham, M. Grötschel, and L. Lovász, eds., MIT Press, Cambridge, MA, 1995, pp. 1447–1540.

[6] L. BABAI, *Coset intersection in moderately exponential time*, Chicago J. Theoret. Comput. Sci., to appear, available onilne at https://people.cs.uchicago.edu/~laci/papers/int.pdf.

[7] L. BABAI, *Equivalence of linear codes*, manuscript, 2010.

[8] L. BABAI, *Graph isomorphism in quasipolynomial time (extended abstract)*, in Proceedings of the 48th Annual ACM Symposium on Theory of Computing, STOC '16, ACM, New York, 2016, pp. 684–697, https://doi.org/10.1145/2897518.2897542.

[9] L. BABAI AND R. BEALS, *A polynomial-time theory of black-box groups I*, in Groups St. Andrews 1997 in Bath, I, C. M. Campbell, E. F. Robertson, N. Ruskuc, and G. C. Smith, eds., London Math. Soc. Lecture Note Ser. 260, Cambridge University Press, Cambridge, UK, 1999, pp. 30–64.

[10] L. BABAI, R. BEALS, AND A. SERESS, *Polynomial-time theory of matrix groups*, in Proceedings of the 41st ACM STOC, ACM, New York, 2009, pp. 55–64, https://doi.org/10.1145/1536414.1536425.

[11] L. Babai, P. Codenotti, J. A. Grochow, and Y. Qiao, *Code equivalence and group isomorphism*, in Proceedings of the 22nd Annual ACM–SIAM Symposium on Discrete Algorithms (SODA'11), SIAM, Philadelphia, 2011, ACM, New York, pp. 1395–1408, https://doi.org/10.1137/1.9781611973082.107.

[12] L. Babai, P. Codenotti, and Y. Qiao, *Polynomial-time isomorphism test for groups with no abelian normal subgroups (extended abstract)*, in International Colloquium on Automata, Languages, and Programming (ICALP), 2012, pp. 51–62, https://doi.org/10.1007/978-3-642-31594-7_5.

[13] L. Babai, W. M. Kantor, and E. M. Luks, *Computational complexity and the classification of finite simple groups*, in Proceedings of the 24th IEEE FOCS, IEEE Computer Society, Los Alamitos, CA, 1983, pp. 162–171, https://doi.org/10.1109/SFCS.1983.10.

[14] L. Babai and E. M. Luks, *Canonical labeling of graphs*, in Proceedings of the 15th ACM STOC, ACM, New York, 1983, pp. 171–183, https://doi.org/10.1145/800061.808746.

[15] L. Babai and S. Moran, *Arthur–Merlin games: A randomized proof system, and a hierarchy of complexity classes*, J. Comput. System Sci., 36 (1988), pp. 254–276, https://doi.org/10.1016/0022-0000(88)90028-1.

[16] L. Babai and Y. Qiao, *Polynomial-time isomorphism test for groups with Abelian Sylow towers*, in Proceedings of the 29th STACS, Lecture Notes in Comput. Sci. 6651, Springer, Berlin, 2012, pp. 453–464, https://doi.org/10.4230/LIPIcs.STACS.2012.453.

[17] R. Baer, *Groups with abelian central quotient group*, Trans. Amer. Math. Soc., 44 (1938), pp. 357–386, https://doi.org/10.2307/1989886.

[18] H. U. Besche and B. Eick, *Construction of finite groups*, J. Symbolic Comput., 27 (1999), pp. 387–404, https://doi.org/10.1006/jsco.1998.0258.

[19] H. U. Besche, B. Eick, and E. O'Brien, *A millennium project: Constructing small groups*, Internat. J. Algebra Comput., 12 (2002), pp. 623–644, https://doi.org/10.1142/S0218196702001115.

[20] R. Boppana, J. Håstad, and S. Zachos, *Does co-NP have short interactive proofs?*, Inform. Process. Lett., 25 (1987), pp. 27–32, https://doi.org/10.1016/0020-0190(87)90232-8.

[21] W. Bosma, J. J. Cannon, and C. Playoust, *The Magma algebra system. I. The user language.*, J. Symbolic Comput., 24 (1997), pp. 235–265.

[22] P. A. Brooksbank and E. M. Luks, *Testing isomorphism of modules*, J. Algebra, 320 (2008), pp. 4020–4029, https://doi.org/10.1016/j.jalgebra.2008.07.014.

[23] P. A. Brooksbank, J. Maglione, and J. B. Wilson, *A fast isomorphism test for groups of genus 2*, J. Algebra, 473 (2017), pp. 545–590, https://doi.org/10.1016/j.jalgebra.2016.12.007.

[24] P. A. Brooksbank and E. A. O'Brien, *Constructing the group preserving a system of forms*, Internat. J. Algebra Comput., 18 (2008), pp. 227–241, https://doi.org/10.1142/S021819670800441X.

[25] P. A. Brooksbank and J. B. Wilson, *Computing isometry groups of Hermitian maps*, Trans. Amer. Math. Soc., 364 (2012), pp. 1975–1996, https://doi.org/10.1090/S0002-9947-2011-05388-2.

[26] J. J. Cannon and D. F. Holt, *Automorphism group computation and isomorphism testing in finite groups*, J. Symbolic Comput., 35 (2003), pp. 241–267, https://doi.org/10.1016/S0747-7171(02)00133-5.

[27] A. Chattopadhyay, J. Torán, and F. Wagner, *Graph isomorphism is not AC^0 reducible to group isomorphism*, ACM Trans. Comput. Theory, 5 (2013), https://doi.org/10.1145/2540088. Originally appeared in FSTTCS 2010, http://doi.org/10.4230/LIPIcs.FSTTCS.2010.317.

[28] A. Chistov, G. Ivanyos, and M. Karpinski, *Polynomial time algorithms for modules over finite dimensional algebras*, in Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation, ISSAC '97, ACM, New York, 1997, pp. 68–74, https://doi.org/10.1145/258726.258751.

[29] I. Ciocănea-Teodorescu, *The Module Isomorphism Problem for Finite Rings and Related Results*, preprint, https://arxiv.org/abs/1512.08365, 2015.

[30] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson, *Atlas of Finite Groups: Maximal Subgroups and Ordinary Characters for Simple Groups*, Oxford University Press, Oxford, UK, 1985.

[31] B. Eick, C. R. Leedham-Green, and E. A. O'Brien, *Constructing automorphism groups of p-groups*, Comm. Algebra, 30 (2002), pp. 2271–2295, https://doi.org/10.1081/AGB-120003468.

[32] S. Eilenberg and S. MacLane, *Cohomology theory in abstract groups. I*, Ann. of Math. (2), 48 (1947), pp. 51–78, https://doi.org/10.2307/1969215.

[33] S. EILENBERG AND S. MACLANE, *Cohomology theory in abstract groups. II: Group extensions with a non-abelian kernel*, Ann. of Math. (2), 48 (1947), pp. 326–341, https://doi.org/10.2307/1969174.

[34] V. FELSCH AND J. NEUBÜSER, *On a programme for the determination of the automorphism group of a finite group.*, in Computational Problems in Abstract Algebra (Proceedings of a Conference on Computational Problems in Algebra, Oxford, 1967), P. J. Leech, ed., Elsevier, Oxford, UK, 1970, pp. 59–60.

[35] H. FITTING, *Die theorie der automorphismenringe abelscher gruppen und ihr analogon bei nicht kommutativen gruppen*, Math. Ann., 107 (1933), pp. 514–542, http://eudml.org/doc/159609.

[36] L. FORTNOW AND J. A. GROCHOW, *Complexity classes of equivalence problems revisited*, Inform. and Comput., 209 (2011), pp. 748–763, https://doi.org/10.1016/j.ic.2011.01.006.

[37] T. GAP GROUP, *GAP—Groups, Algorithms, and Programming, version* 4.6.5, 2013, http://www.gap-system.org/.

[38] M. GARZON AND Y. ZALCSTEIN, *On isomorphism testing of a class of 2-nilpotent groups*, J. Comput. System Sci., 42 (1991), pp. 237–248, https://doi.org/10.1016/0022-0000(91)90012-T.

[39] M. GOLDMANN AND A. RUSSELL, *The complexity of solving equations over finite groups*, Inform. and Comput., 178 (2002), pp. 253–262, https://doi.org/10.1016/S0890-5401(02)93173-1.

[40] J. A. GROCHOW, *Matrix Lie algebra isomorphism*, in Proceedings of the IEEE Conference on Computational Complexity (CCC12), IEEE, Washington, DC, 2012, pp. 203–213, https://doi.org/10.1109/CCC.2012.34.

[41] J. A. GROCHOW AND Y. QIAO, *Algorithms for group isomorphism via group extensions and cohomology*, in Proceedings of the IEEE Conference on Computational Complexity (CCC14), IEEE, Washington, DC, 2014; preprint, https://arxiv.org/abs/1309.1776, 2013.

[42] J. A. GROCHOW AND Y. QIAO, *Polynomial-time isomorphism test of groups that are tame extensions (extended abstract)*, in Proceedings of the 26th International Symposium on Algorithms and Computation (ISAAC 2015), Nagoya, Japan, 2015, pp. 578–589, https://doi.org/10.1007/978-3-662-48971-0_49.

[43] R. GURALNICK, W. M. KANTOR, M. KASSABOV, AND A. LUBOTZKY, *Presentations of finite simple groups: Profinite and cohomological approaches*, Groups Geom. Dyn., 1 (2007), pp. 469–523, https://doi.org/10.4171/GGD/22.

[44] J.-F. HÄMMERLI, *Some remarks on nonconnected compact Lie groups*, Enseign. Math. (2), 49 (2003), pp. 67–84.

[45] A. HATCHER, *Algebraic Topology*, Cambridge University Press, Cambridge, UK, 2002.

[46] C. J. HILLAR AND D. L. RHEA, *Automorphisms of finite abelian groups*, Amer. Math. Monthly, 114 (2007), pp. 917–923; preprint, https://arxiv.org/abs/math/0605185, 2006.

[47] D. F. HOLT, B. EICK, AND E. A. O'BRIEN, *Handbook of Computational Group Theory*, Chapman and Hall/CRC, London, 2005.

[48] D. F. HOLT AND W. PLESKEN, *Perfect Groups*, Oxford Mathematical Monographs, Oxford University Press, New York, 1989.

[49] D. J. A. HOWDEN, *Computing Automorphism Groups and Isomorphism Testing in Finite Groups*, Ph.D. thesis, The University of Warwick, Coventry, UK, 2012; available at http://wrap.warwick.ac.uk/50060/1/WRAP_THESIS_Howden_2012.pdf.

[50] C. S. ILIOPOULOS, *Computing in general abelian groups is hard*, Theoret. Comput. Sci., 41 (1985), pp. 81–93, https://doi.org/10.1016/0304-3975(85)90061-1.

[51] I. M. ISAACS, *Finite Group Theory*, Grad. Stud. Math. 92, AMS, Providence, RI, 2008, https://doi.org/10.1090/gsm/092.

[52] G. IVANYOS, M. KARPINSKI, AND N. SAXENA, *Deterministic polynomial time algorithms for matrix completion problems*, SIAM J. Comput., 39 (2010), pp. 3736–3751, https://doi.org/10.1137/090781231.

[53] G. IVANYOS AND Y. QIAO, *Algorithms Based on *-Algebras, and Their Applications to Isomorphism of Polynomials with One Secret, Group Isomorphism, and Polynomial Identity Testing*, manuscript, 2016.

[54] T. KAVITHA, *Linear time algorithms for abelian group isomorphism and related problems*, J. Comput. System Sci., 73 (2007), pp. 986–996, https://doi.org/10.1016/j.jcss.2007.03.013.

[55] N. KAYAL AND T. NEZHMETDINOV, *Factoring groups efficiently*, in Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP '09), Springer-Verlag, New York, 2009, pp. 585–596, https://doi.org/10.1007/978-3-642-02927-1_49.

[56] J. KÖBLER, U. SCHÖNING, AND J. TORÁN, *The Graph Isomorphism Problem: Its Structural Complexity*, Birkhäuser Verlag, Basel, 1993.

[57] M. LAZARD, *Sur les groupes nilpotents et les anneaux de Lie*, Ann. Sci. Ecole Norm. Sup., 71 (1954), pp. 101–190.

[58] F. LE GALL, *Efficient isomorphism testing for a class of group extensions*, in Proceedings of the 26th STACS, Freiburg, Germany, 2009, pp. 625–636, https://doi.org/10.4230/LIPIcs. STACS.2009.1830.

[59] F. LE GALL AND D. J. ROSENBAUM, *On the Group and Color Isomorphism Problems*, preprint, https://arxiv.org/abs/1609.08253, 2016.

[60] M. L. LEWIS AND J. B. WILSON, *Isomorphism in expanding families of indistinguishable groups*, Groups Complex. Cryptol., 4 (2012), pp. 73–110, https://doi.org/10.1515/gcc-2012-0008.

[61] E. M. LUKS, *Isomorphism of graphs of bounded valence can be tested in polynomial time*, J. Comput. System Sci., 25 (1982), pp. 42–65, https://doi.org/10.1016/0022-0000(82) 90009-5.

[62] E. M. LUKS, *Permutation groups and polynomial-time computation*, in Proceedings of the Workshop on Groups and Computation, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 11, AMS, Providence, RI, 1991, pp. 139–175.

[63] E. M. LUKS, *Hypergraph isomorphism and structural equivalence of boolean functions*, in Proceedings of the 31st ACM STOC, ACM, New York, 1999, pp. 652–658, https://doi.org/10. 1145/301250.301427.

[64] G. L. MILLER, *On the $n \log n$ isomorphism technique (a preliminary report)*, in Proceedings of STOC, ACM, New York, 1978, pp. 51–58, https://doi.org/10.1145/800133.804331.

[65] K. MULMULEY, *On* P *vs.* NP *and geometric complexity theory*, J. ACM, 58 (2011), 5, https: //doi.org/10.1145/1944345.1944346.

[66] V. NAIK, *Isomorphic Extensions that Are Not Pseudo-congruent*, personal communication, 2010.

[67] V. NAIK, *Group extension problem*, http://groupprops.subwiki.org/wiki/Group_extension_ problem, 2012.

[68] E. O'BRIEN, *Isomorphism testing for p-groups*, J. Symbolic Comput., 17 (1994), pp. 133–147, https://doi.org/10.1006/jsco.1994.1007.

[69] E. PETRANK AND R. M. ROTH, *Is code equivalence easy to decide?*, IEEE Trans. Inform. Theory, 43 (1997), pp. 1602–1604, https://doi.org/10.1109/18.623157.

[70] Y. QIAO, J. M. N. SARMA, AND B. TANG, *On isomorphism testing of groups with normal Hall subgroups*, in Proceedings of the 28th STACS, 2011, pp. 567–578, https://doi.org/10.4230/ LIPIcs.STACS.2011.567.

[71] A. RANUM, *The group of classes of congruent matrices with application to the group of isomorphisms of any abelian group*, Trans. Amer. Math. Soc., 8 (1907), pp. 71–91, https://doi.org/10.2307/1986235.

[72] D. ROBINSON, *Applications of cohomology to the theory of groups*, in Groups St. Andrews 1981, London Math. Soc. Lecture Note Ser. 71, Cambridge University Press, Cambridge, UK, 1982, pp. 46–80, https://doi.org/10.1017/CBO9780511661884.005.

[73] D. J. ROBINSON, *A Course in the Theory of Groups*, 2nd ed., Springer, New York, 1996.

[74] L. RÓNYAI, *Computing the order of centralizers in linear groups*, Inform. and Comput., 91 (1991), pp. 172–176, https://doi.org/10.1016/0890-5401(91)90064-9.

[75] D. ROSENBAUM, *Bidirectional Collision Detection and Faster Deterministic Isomorphism Testing*, preprint, https://arxiv.org/abs/1304.3935, 2013.

[76] D. ROSENBAUM, *Breaking the $n^{\log n}$ barrier for solvable-group isomorphism*, in Proceedings of the 24th SODA, SIAM, Philadelphia, ACM, New York, 2013, pp. 1054–1073, https: //doi.org/10.1137/1.9781611973105.76.

[77] J. J. ROTMAN, *An Introduction to the Theory of Groups*, 4th ed., Grad. Texts in Math. 148, Springer-Verlag, New York, 1995, https://doi.org/10.1007/978-1-4612-4176-8.

[78] C. SAVAGE, *An $O(n^2)$ Algorithm for Abelian Group Isomorphism*, Tech. report, North Carolina State University, Raleigh, NC, 1980.

[79] A. SERESS, *Permutation Group Algorithms*, Cambridge University Press, Cambridge, UK, 2003.

[80] M. J. SMITH, *Computing Automorphisms of Finite Soluble Groups*, Ph.D. thesis, Australian National University, Canberra, Australia, 1994; available online at http://www.mathe2.uni-bayreuth.de/axel/papers/smith:computing_automorphisms_of_ finite_soluble_groups.ps.gz, summary in Bull. Austral. Math. Soc., 53 (1996), pp. 169–171.

[81] R. STEINBERG, *Generators for simple groups*, Canad. J. Math., 14 (1962), pp. 277–283, https: //doi.org/10.4153/CJM-1962-018-0.

[82] M. SUZUKI, *Group Theory* I, Springer, Berlin, New York, 1982.

[83] M. SUZUKI, *Group Theory* II, Springer, New York, 1986.

[84] D. R. TAUNT, *Remarks on the isomorphism problem in theories of construction of finite groups*, Proc. Cambridge Philos. Soc., 51 (1955), pp. 16–24.

[85] N. Vikas, *An O(n) algorithm for abelian p-group isomorphism and an O(n log n) algorithm for abelian group isomorphism*, J. Comput. System Sci., 53 (1996), pp. 1–9, https://doi.org/10.1006/jcss.1996.0045.

[86] F. Wagner, *On the complexity of group isomorphism*, Tech. report TR11-052, Electronic Colloquium on Computational Complexity (ECCC), 2011.

[87] R. Warfield, *Nilpotent Groups*, Lecture Notes in Math. 513, Springer-Verlag, Berlin, 1976.

[88] J. Wilson, Personal communication, Conference on Groups, Computation, and Geometry at Colorado State University, co-organized by P. Brooksbank, A. Hulpke, T. Penttila, J. Wilson, and W. Kantor), 2014.

[89] J. B. Wilson, *Decomposing p-groups via Jordan algebras*, J. Algebra, 322 (2009), pp. 2642–2679, https://doi.org/10.1016/j.jalgebra.2009.07.029.

[90] J. B. Wilson, *Finding central decompositions of p-groups*, J. Group Theory, 12 (2009), pp. 813–830, https://doi.org/10.1515/JGT.2009.015.

[91] J. B. Wilson, *Existence, algorithms, and asymptotics of direct product decompositions,* I, Groups Complex. Cryptol., 4 (2012), pp. 33–72, https://doi.org/10.1515/gcc-2012-0007.

[92] J. B. Wilson, *More characteristic subgroups, Lie rings, and isomorphism tests for p-groups*, J. Group Theory, 16 (2013), pp. 875–897, https://doi.org/10.1515/jgt-2013-0026.