Journal of Algebra 353 (2012) 286–297

Contents lists available at SciVerse ScienceDirect

Journal of Algebra

www.elsevier.com/locate/jalgebra

Intersecting two classical groups ☆

Peter A. Brooksbank [a,*], James B. Wilson [b]

[a] Department of Mathematics, Bucknell University, Lewisburg, PA 17837, United States
[b] Department of Mathematics, Colorado State University, Fort Collins, CO 80523, United States

ARTICLE INFO

Article history:
Received 18 September 2011
Available online 4 January 2012
Communicated by Derek Holt

Keywords:
$*$-Algebra
Bilinear map
Isometry group
Polynomial-time algorithm

ABSTRACT

A new algorithm is presented to compute the algebra of adjoints associated to a pair of forms on a common finite vector space. This algebra is used in several recent and ongoing projects to study central products, intersections of classical groups, and automorphism groups. The implementation of the new algorithms in Magma greatly outperforms its predecessor and hence, in restricted yet important settings, increases the practicality of the various algorithms that use adjoint algebras.

1. Introduction

This paper considers a classical problem, apparently studied first by Kronecker, concerning pairs of reflexive forms on a common vector space. Such pairs arise, for example, when intersecting two classical groups [BF,GG1,BO,GG2], and also when constructing automorphism groups of groups and rings [BW2]. The general objective is to understand the geometry common to both reflexive forms. The specific goal of this paper is to introduce new algebraic ideas and algorithmic techniques that enable us to determine the necessary geometric information very efficiently.

In earlier work [BW1], the authors solved the general problem of describing the intersection of an arbitrary set of classical subgroups defined on a common vector space of odd order (that is, without restriction on numbers of forms, their geometric types, degeneracy, or their specific natural field of definition). In addition, a polynomial-time algorithm was presented to construct this intersection and to describe its structure. The algorithm uses $O^{\sim}(\mathrm{MM}(d^2))$ field operations,

☆ Project sponsored by the National Security Agency under Grant Number H98230-11-1-0146. The United States Government is authorized to reproduce and distribute reprints not-withstanding any copyright notation herein.

* Corresponding author.
E-mail addresses: pbrooksb@bucknell.edu (P.A. Brooksbank), jwilson@math.colostate.edu (J.B. Wilson).

0021-8693/$ – see front matter
doi:10.1016/j.jalgebra.2011.12.004

where $\mathrm{MM}(n)$ is the number of field operations required to multiply $(n \times n)$-matrices. (Recall that $O^{\sim}(f(n)) = O(f(n)\log^c n)$, where c is a constant.)

The key algorithmic and theoretical tool in [BW1] is the *ring of adjoints* (Section 2) and the construction of this ring is also the principal algorithmic bottleneck. Thus, improving the practical reach of our methods demands a faster technique to compute the ring of adjoints. The main result of this paper shows that a significant improvement is indeed possible if we restrict to a pair of classical groups.

Theorem 1.1. *There is a Las Vegas algorithm which, given classical subgroups G, H of $\mathrm{GL}(V)$, where V is a finite d-dimensional vector space over a finite field k of odd characteristic, returns a generating set $\mathcal{X}$ for the group $G \cap H$. The algorithm uses $O^{\sim}(|\mathcal{X}| \cdot \mathrm{MM}(d) + \mathrm{RND}(d) + \mathrm{FAC}(d, e))$ operations in k, where $e = 2$ if precisely one of G or H is unitary, and $e = 1$ otherwise.*

In the complexity statement, $\mathrm{RND}(d)$ denotes the number of operations needed to select independent, (nearly) uniformly distributed random elements from an algebra of $(d \times d)$-matrices, and $\mathrm{FAC}(d, e)$ denotes the number of field operations needed to factorize a polynomial of degree d in a skew-polynomial ring $k[x; \theta]$, where θ has order e. These complexity parameters are discussed in more detail in the remarks below.

The main tool in Theorem 1.1 is the introduction of a single semilinear transformation, which we call the *slope* (Section 3), whose centralizer is the ring of adjoints. To take advantage of this observation we devise a polynomial-time algorithm to construct the centralizer of a semilinear transformation (Theorem 4.1) which we believe will have uses independent of our present application. Once we have constructed the ring of adjoints as a centralizer, our strategy follows [BW1].

A notion analogous to the slope is used in both [GG1] and [BO]. We abstract that notion here to admit cases where one or both of the forms is Hermitian. As far as we are aware no attempt has previously been made to study such pairs of forms directly on their natural domain. They can, for example, be handled using the (slower) methods of [BW1], but there one must consider instead the isometries common to a set of four forms defined on a subfield.

In [BW1] we reported on an implementation of our algorithm in Magma [BCP]. We have since incorporated the methods used to prove Theorem 1.1 and produced a Magma package to compute effectively with reflexive forms and their associated algebraic structures. In Section 6 we report on this package, noting the dramatic improvement in performance one sees by using the newer "slope methods" for pairs of forms instead of the earlier generic algorithm.

Remarks on complexity. There are several aspects of the complexity stated in Theorem 1.1 that require further comment.

(1) The complexity of $\mathrm{MM}(d)$ is $O(d^{\omega})$, where $2 \leqslant \omega \leqslant 3$. In practice the exponent ω is closer to 3 than to 2 [vzGG].
(2) The complexity of $\mathrm{FAC}(d, e)$ is $O(\mathrm{MM}(d) + e(e-1)d^4)$ [Ge]. Hence, in our setting the complexity is $O(d^4)$ if precisely one of the input classical groups is unitary, and $O(\mathrm{MM}(d))$ for all other pairs of input groups.
(3) Despite much effort, the complexity of $\mathrm{RND}(d)$ is not fully understood (cf. [Iv, Section 2]). It is no worse than $O(\mathrm{MM}(d^2))$, the number of field operations needed to construct a basis of the given algebra. Fortunately, since our randomized algorithm is Las Vegas (any result is guaranteed to be correct), we may use standard $O(\mathrm{MM}(d))$-time random generation heuristics with greater confidence.
(4) The crude bound of $O(d^2)$ on the cardinality of the output generating set $\mathcal{X}$ is already enough to establish the improved complexity of the new algorithm over its predecessor in [BW1]. However, based both on intuition and experimental evidence, we believe that $|\mathcal{X}|$ is actually $O(d)$. Moreover, the run-times of our implementation support the view that the complexity of the entire algorithm is $O^{\sim}(d \cdot \mathrm{MM}(d))$ (see Section 6).

2. Background

To understand the new ideas we first need a brief description of the main concepts in [BW1] along with various standard definitions. Throughout this section, k will denote a field, and V and W will denote finite-dimensional k-vector spaces.

Bimaps and reflexive forms. A *k-bimap* on V is a function $b : V \times V \to W$ that is k-semilinear in the first variable and k-linear in the second variable. We associate to a k-bimap b its group of isometries

$$\operatorname{Isom}(b) = \big\{ f \in \mathrm{GL}(V) \colon b(uf, vf) = b(u, v) \text{ for all } u, v \in V \big\}. \tag{2.1}$$

In the special case that W is 1-dimensional, a k-bimap is usually referred to as a *k-form*. A k-form φ is *reflexive* if $\varphi(u, v) = 0$ whenever $\varphi(v, u) = 0$ for $u, v \in V$, meaning that the usual perpendicularity relation is symmetric. A reflexive k-form φ is *nondegenerate* if $\operatorname{rad}(\varphi) = \{v \in V \colon \varphi(v, V) = 0\} = 0$.

The (misattributed) Birkhoff–von Neumann Theorem [Tay, Theorem 7.1] asserts that every nondegenerate reflexive form on V is a scalar multiple of one of the following types of forms:

(i) *alternating* if $\varphi(v, v) = 0$ for all $v \in V$;
(ii) *symmetric* if $\varphi(u, v) = \varphi(v, u)$ for all $u, v \in V$; or
(iii) *Hermitian* if there is an automorphism α of k such that $\varphi(u, v) = \varphi(v, u)^\alpha$ for all $u, v \in V$.

These definitions apply also to degenerate reflexive k-forms. As the group of isometries is not changed by scaling a form, we regard all reflexive forms as alternating, symmetric, or Hermitian.

Classical groups. We say that $G \leqslant \mathrm{GL}(V)$ is *classical* if $\operatorname{Isom}(\varphi)' \leqslant G \leqslant \operatorname{Isom}(\varphi)$ for some nondegenerate reflexive k-form φ on V. Thus, the intersection of two classical groups contains isometries of a pair $[\varphi, \gamma]$ of distinct reflexive forms. Following [BW1], we associate to such a pair the bimap $(\varphi \cap \gamma) : V \times V \to k^2$ defined naturally as

$$(\varphi \cap \gamma)(u, v) = \big(\varphi(u, v), \gamma(u, v)\big) \quad (\forall u, v \in V) \tag{2.2}$$

and work instead with $b = \varphi \cap \gamma$, observing that $\operatorname{Isom}(\varphi \cap \gamma) = \operatorname{Isom}(\varphi) \cap \operatorname{Isom}(\gamma)$. This provides some flexibility. For, if both φ and γ are degenerate (making them unsuitable for Theorem 1.1), it may be the case that $\varphi \cap \gamma = \varphi' \cap \gamma'$ with one of φ' or γ' nondegenerate; one may then replace $[\varphi, \gamma]$ by $[\varphi', \gamma']$. An example of such a pair is given in Section 7.

Algebra of adjoints. Next we associate to a k-bimap $b : V \times V \to W$ a k-algebra of *adjoints*,

$$\operatorname{Adj}(b) = \big\{ (f, g) \in \operatorname{End} V \times (\operatorname{End} V)^{\mathrm{op}} \colon \forall u, v \in V,\ b(uf, v) = b(u, gv) \big\}, \tag{2.3}$$

where $\operatorname{End} V$ is the ring of endomorphisms of V, and $(\operatorname{End} V)^{\mathrm{op}}$ is its opposite ring. We assume that elements of $\operatorname{End} V$ and $(\operatorname{End} V)^{\mathrm{op}}$ act on V on the right and on the left, respectively. Then, for $f \in \operatorname{End} V$ and $g \in (\operatorname{End} V)^{\mathrm{op}}$, the assignments $f^{\mathrm{op}} v := vf$ and $v g^{\mathrm{op}} := gv$ for all $v \in V$ define inverse isomorphisms between $\operatorname{End} V$ and $(\operatorname{End} V)^{\mathrm{op}}$. (Writing endomorphisms as matrices relative to a fixed basis of V, the "op" map is simply transposition.)

If b is nondegenerate then $(f, g), (f, g') \in \operatorname{Adj}(b)$ implies that $g = g'$; in that case we usually write f^* for g. Furthermore, since φ and γ are reflexive, if $(f, f^*) \in \operatorname{Adj}(b)$, then also $((f^*)^{\mathrm{op}}, f^{\mathrm{op}}) \in \operatorname{Adj}(b)$. Thus $\operatorname{Adj}(b)$ is equipped with a natural *involution* (anti-automorphism of order at most 2) defined by

$$(f, g)^* = \big(g^{\mathrm{op}}, f^{\mathrm{op}}\big) \quad \big(\forall (f, g) \in \operatorname{Adj}(b)\big). \tag{2.4}$$

As with isometry groups, observe that $\mathrm{Adj}(\varphi \cap \gamma) = \mathrm{Adj}(\varphi) \cap \mathrm{Adj}(\gamma)$. Observe also that V is both a left and a right $\mathrm{Adj}(b)$-module and so we can form $V \otimes_{\mathrm{Adj}(b)} V$. Then Eq. (2.3) may be interpreted as saying that b factors through the k-bimap $\otimes_{\mathrm{Adj}(b)} : V \times V \to V \otimes_{\mathrm{Adj}(b)} V$; indeed $\mathrm{Adj}(b)$ is universal with that property.

Isometry groups. In Eq. (2.1) notice that isometries of $b : V \times V \to W$ are applied only to V and not to W. Hence, we might hope to replace b with another k-bimap on V possessing the same isometries as b, but having a more natural codomain. Indeed, observe that $\mathrm{Isom}(b) = \mathrm{Isom}(\otimes_{\mathrm{Adj}(b)})$. Expressed another way, we see that

$$\mathrm{Isom}(b) = \left\{ f \in \mathrm{GL}(V) \colon \left(f, \left(f^{-1}\right)^{\mathrm{op}}\right) \in \mathrm{Adj}(b) \right\} \tag{2.5}$$

$$\cong \left\{ (f, g) \in \mathrm{Adj}(b) \colon (f, g)(f, g)^* = (1, 1) \right\}. \tag{2.6}$$

The group in Eq. (2.6) is referred to as the group of *unitary* elements (or the *norm* 1 *group*) of the $*$-algebra $\mathrm{Adj}(b)$ and we denote it $\mathrm{Adj}(b)^\sharp$.

Remark 2.1. When defining $\mathrm{Isom}(b)$ and $\mathrm{Adj}(b)$ it is not necessary to specify k. Indeed we were careful not do so explicitly in their definitions in Eqs. (2.1) and (2.3). We need only define elements of $\mathrm{Adj}(b)$ and $\mathrm{Isom}(b)$ as endomorphisms of V *as an abelian group*. For, if $(f, g) \in \mathrm{Adj}(\varphi)$, then, for all $u, v \in V$ and $a \in k$, we have

$$b\big((au)f, v\big) = b(au, gv) = a^\alpha b(u, gv) = a^\alpha b(uf, v) = b\big(a(uf), v\big).$$

Since b is nondegenerate, f is k-linear. Likewise, g is k-linear. Using the observation that $\mathrm{Isom}(b)$ embeds in $\mathrm{Adj}(b)$ as $\mathrm{Adj}(b)^\sharp$ we see that elements of $\mathrm{Isom}(b)$ are also k-linear. This subtle observation is crucial to our proof of Lemma 3.2.

From Adj(b) to Isom(b). Eq. (2.6) translates the problem of computing isometries of a k-bimap into the problem of constructing unitary elements in a $*$-algebra. The following result, which is extracted from [BW1], shows that the latter can be solved very efficiently over finite fields k of odd characteristic.

Theorem 2.1. *There is a Las Vegas algorithm which, given $A \leqslant \mathrm{End}_k V \times (\mathrm{End}_k V)^{\mathrm{op}}$, where k is a finite field of odd characteristic, and an involution $(f, g)^* = (g^{\mathrm{op}}, f^{\mathrm{op}})$ on A, returns a generating set, $\mathcal{X}$, for $A^\sharp$ and describes the structure of the group. The algorithm uses $O(|\mathcal{X}| \cdot \mathrm{MM}(d) + \mathrm{RND}(d))$ operations in k.*

Proof. We use the algorithm IsometryGroup [BW1, Section 5], but omit Line 1 (since we make the assumption that the $*$-ring is given). Its complexity is then $O^\sim(|\mathcal{X}| \cdot \mathrm{MM}(d) + \mathrm{RND}(d))$ field operations [BW1, Section 5.5]. Note that, in [BW1], complexity estimates are stated assuming standard cubic algorithms for matrix operations, so that $\mathrm{MM}(d) = \Theta(d^3)$. □

The hypothesis that k has odd characteristic is needed only if A has a nontrivial Jacobson radical. In particular, the same result holds for semisimple $*$-algebras over any finite field.

3. The slope of a pair of reflexive forms

Throughout this section φ and γ denote reflexive forms on a finite-dimensional k-vector space V, and we shall further assume that φ is nondegenerate. We intend to prove Theorem 1.1 using Theorem 2.1 together with an efficient method to construct the adjoint algebra $\mathrm{Adj}(\varphi \cap \gamma) = \mathrm{Adj}(\varphi) \cap \mathrm{Adj}(\gamma)$. In this section we make the crucial observation that $\mathrm{Adj}(\varphi \cap \gamma)$ may be realized as the centralizer of a certain isometry invariant that we call the *slope* of $[\varphi, \gamma]$.

This section applies to all fields in all characteristics.

As mentioned in the introduction, the concept of slope has appeared in various guises in the literature. However, it has always been expressed with matrices, which naturally resulted in the exclusion of Hermitian forms. The following general definition captures the essential properties needed for our purpose, and allows us to apply the slope to a wider range of geometric configurations.

Definition 3.1. The (*left*) *slope* of a pair $[\varphi, \gamma]$ of reflexive forms on a k-vector space V is a function $\sigma : V \to V$ which, for all $u, v \in V$, satisfies

$$\varphi(u\sigma, v) = \gamma(u, v). \tag{3.1}$$

We say a pair $[\varphi, \gamma]$ of forms is *sloped* if it has a slope and otherwise we say the pair is *flat*.

Notice we did not demand that σ is k-linear, only that it is a function. The biadditivity of φ and γ implies that σ is additive; however, we will see that σ may only be k-semilinear. It is this generality that enables us to include intersections of unitary groups with non-unitary groups in the present treatment. Those cases had been handled previously only by the very different methods of [BW1].

We give examples of flat pairs $[\varphi, \gamma]$ in Section 7.

Recall that each reflexive form is a scalar multiple of an alternating, symmetric, or Hermitian form. Hence, there exist field automorphisms α and β associated to φ and γ, respectively, such that, for all $u, v \in V$ and all $a \in k$, $\varphi(au, v) = a^\alpha \varphi(u, v)$ and $\gamma(au, v) = a^\beta \gamma(u, v)$.

Our first result establishes the basic properties of slope.

Lemma 3.1. *Let φ and γ be reflexive forms on V, with φ nondegenerate. Then $[\varphi, \gamma]$ determines a unique slope, which is a k-semilinear transformation of V.*

Proof. First we prove uniqueness. Suppose σ and τ are functions that satisfy Eq. (3.1). Then, for all $u, v \in V$, $\varphi(u\sigma, v) = \gamma(u, v) = \varphi(u\tau, v)$, so that $\varphi(u(\sigma - \tau), v) = 0$. As φ is nondegenerate, $\sigma - \tau = 0$.

To establish the existence of a slope, we define $\varphi^\dashv : V \to \mathrm{Hom}_k(V, k)$ such that $u\varphi^\dashv = \varphi(u, -)$. Note that $\varphi^\dashv$ is well-defined by our convention for Hermitian forms, since $u\varphi^\dashv$ is k-linear. Similarly define $\gamma^\dashv : V \to \mathrm{Hom}_k(V, k)$. Both $\varphi^\dashv$ and $\gamma^\dashv$ are semilinear transformations. As φ is nondegenerate and V is finite-dimensional over k, $\varphi^\dashv$ is a k-semilinear isomorphism. Evidently $\sigma := \gamma^\dashv(\varphi^\dashv)^{-1}$ is k-semilinear and satisfies Eq. (3.1). □

The essential condition in this proof that ensures the existence of a slope for a pair $[\varphi, \gamma]$ is that $V\gamma^\dashv \leqslant V\varphi^\dashv$. Thus slopes exist in greater generality than is stated in the lemma. For convenience we write $\sigma = \gamma\varphi^{-1}$ even though that notation is not as precise as $\sigma = \gamma^\dashv(\varphi^\dashv)^{-1}$.

The following result is the key to computing effectively with $\mathrm{Adj}(\varphi \cap \gamma)$. Part (i) shows that $\mathrm{Adj}(\varphi \cap \gamma)$ may be realized as a centralizer ring, and part (ii) may be adapted to construct the involution on this ring efficiently (cf. Section 5).

Lemma 3.2. *Let φ and γ be reflexive k-forms on V, with φ nondegenerate, and let $\sigma = \gamma\varphi^{-1}$. Then the following hold.*

(i) $(x, x^*) \in \mathrm{Adj}(\varphi \cap \gamma) \Leftrightarrow x \in C_{\mathrm{End}_k(V)}(\sigma) = \{f \in \mathrm{End}_k V \colon \sigma f = f\sigma\}$.
(ii) $\mathrm{Adj}(\varphi)|_V = \mathrm{End}_k V$; *moreover, for every $x \in \mathrm{End}_k V$ there is a unique $x^* \in (\mathrm{End}_k V)^{\mathrm{op}}$ with $(x, x^*) \in \mathrm{Adj}(\varphi)$.*

Proof. For part (ii) observe, by Remark 2.1, that $\mathrm{Adj}(\varphi)|_V \subseteq \mathrm{End}_k V$. For the reverse inclusion, we regard $\mathrm{Hom}_k(V, k)$ as a right $(\mathrm{End}_k V)^{\mathrm{op}}$-module where, for each $f \in (\mathrm{End}_k V)^{\mathrm{op}}$, we define $\tau f^\circ := f\tau$ for all $\tau \in \mathrm{Hom}_k(V, K)$. As φ is nondegenerate, $\varphi^\dashv$ is an isomorphism and $x\varphi^\dashv = \varphi^\dashv(x^{*\mathrm{op}})^\circ$ defines x^* uniquely.

Now we prove part (i). Let $(x, x^*) \in \mathrm{Adj}(\varphi \cap \gamma)$. Note that $x \in \mathrm{End}_k V$ by Remark 2.1. Also, for all $u, v \in V$,

$$\varphi(ux\sigma, v) = \gamma(ux, v) = \gamma\big(u, x^* v\big) = \varphi\big(u\sigma, x^* v\big) = \varphi(u\sigma x, v).$$

Since φ is nondegenerate, $x\sigma - \sigma x = 0$. So $x \in C_{\mathrm{End}_k V}(\sigma)$.

Next, let $x \in C_{\mathrm{End}_k V}(\sigma)$. By (ii) there is a unique $x^\bullet \in (\mathrm{End}_k V)^{\mathrm{op}}$ such that $(x, x^\bullet) \in \mathrm{Adj}(\varphi)$. So, for all $u, v \in V$, we have

$$\gamma(ux, v) = \varphi(ux\sigma, v) = \varphi(u\sigma x, v) = \varphi\big(u\sigma, x^\bullet v\big) = \gamma\big(u, x^\bullet v\big).$$

Hence, $(x, x^\bullet) \in \mathrm{Adj}(\gamma)$, so that $(x, x^\bullet) \in \mathrm{Adj}(\varphi \cap \gamma)$. □

Historical comment. Previous uses of "slope" in the literature have focused on the transformation σ itself, rather than on its centralizing ring. While many geometric properties of $\mathrm{Isom}(\varphi) \cap \mathrm{Isom}(\gamma)$ can be understood in terms of this single transformation, there does not seem to be a way to use it directly to describe intersections with unitary groups. We can now see why this is the case.

When one or both of the forms φ and γ is Hermitian, the slope σ can be semilinear, and so σ need not lie in $C_{\mathrm{End}_k V}(\sigma)$. Hence one cannot in general find σ^* such that (σ, σ^*) lies in $\mathrm{Adj}(\varphi \cap \gamma)$. Nevertheless, Lemma 3.2 shows that σ may still be used to recover the entire adjoint ring, rather than just one element of it. That makes all the geometric analysis possible.

4. The centralizer of a semilinear transformation

In this section we describe centralizers of semilinear transformations in a manner that admits an algorithmic treatment. The ideas underlying our exposition are classical [Ja, Chapter 3] and are analogous to the more familiar theory of finitely generated modules over principal ideal domains.

The results in this section hold for all finite fields.

Let K be a finite degree field extension of a finite field k, let U be a finite-dimensional K-vector space, and let $h \in \mathrm{End}_k U$ be K-semilinear. Hence, there exists $\theta \in \mathrm{Gal}(K/k)$ such that

$$(av)h = a^\theta vh \quad (\forall a \in K,\ \forall v \in V). \tag{4.1}$$

We are concerned with the structure and computability of the ring

$$C_{\mathrm{End}_K U}(h) = \{f \in \mathrm{End}_K U \colon hf = fh\}. \tag{4.2}$$

The *skew polynomial ring* $K[x; \theta]$ is the free K-algebra with generator x having relations

$$ax^i = x^i a^{\theta^i} \quad (\forall a \in K,\ \forall i \in \mathbb{N}).$$

Elements of $K[x; \theta]$ can be expressed as right polynomials $\sum_{i \in \mathbb{N}} x^i a_i$, with $a_i \in K$ with only finitely many a_i's nonzero. As θ is invertible and K is a field, $K[x; \theta]$ is a principal right ideal domain (PRI). This ring is noncommutative so when we say one element divides another we mean division on the right.

The ring $K[x; \theta]$ acts on U via $up(x) := up(h)$, for $u \in U$ and $p(x) \in K[x; \theta]$. The image of this representation of $K[x; \theta]$ in $\mathrm{End}_k U$ is the enveloping algebra $K[h]$, namely the ring of right K-polynomials in h. We now use U as a $K[x; \theta]$-module to describe $C_{\mathrm{End}_K U}(h) = \mathrm{End}_{K[h]} U = \mathrm{End}_{K[x;\theta]} U$.

As $K[x;\theta]$ is a PRI, there is a monic polynomial $m_h(x)$ – the *minimum polynomial of* h – generating the annihilator of U in $K[x;\theta]$. Furthermore, there is a decomposition of the form

$$m_h(x) = p_1(x)^{\lambda_1} \cdots p_s(x)^{\lambda_s}, \tag{4.3}$$

where $p_1(x), \ldots, p_s(x) \in K[x;\theta]$ are irreducible, with ordering and choice of p_i unique up to a permutation and unit, respectively [Ja, Chapter 3, Theorem 5].

Fix $i \in \{1, \ldots, s\}$. Define

$$U_i = \ker p_i(h)^{\lambda_i}, \quad \text{and} \tag{4.4}$$

$$q_i(x) = p_1(x)^{\lambda_1} \cdots p_{i-1}(x)^{\lambda_{i-1}} p_{i+1}(x)^{\lambda_{i+1}} \cdots p_s(x)^{\lambda_s} \in K[x;\theta]. \tag{4.5}$$

By the Chinese Remainder Theorem, we have $K[x;\theta]/(m(x)) = (p_i(x)^{\lambda_i}) \oplus (q_i(x))$ so that $U = U(p_i(x)^{\lambda_i}) \oplus U(q_i(x))$. Intentionally, $U(q_i(x)) \leqslant U_i$ and, since the annihilator in $K[x;\theta]$ of U_i is $(p_i(x)^{\lambda_i})$, it follows that $U(q_i(x)) = U_i$. Eq. (4.4) shows that U_i is readily computable, while the second description shows that U_i is the $p_i(x)$-primary component of U, and that $U = U_1 \oplus \cdots \oplus U_s$. Set

$$R_i = K[x;\theta]/\big(p_i(x)^{\lambda_i}\big). \tag{4.6}$$

Then U_i is a right R_i-module. Since R_i is a chain ring (i.e. its ideals are linearly ordered) every directly indecomposable submodule of U_i is isomorphic to $K[x;\theta]/(p_i(x)^d)$ for some $d \leqslant \lambda_i$. Hence,

$$U_i = U_{i1} \oplus \cdots \oplus U_{it_i}, \tag{4.7}$$

where, for $1 \leqslant i \leqslant t_i$, $U_{ij} \cong K[x;\theta]/(p_i(x)^{\lambda_{ij}})$ with $0 < \lambda_{i1} \leqslant \cdots \leqslant \lambda_{it_i} = \lambda_i$.

This information is sufficient to determine the structure of $\mathrm{End}_{K[x;\theta]} U$. As the annihilators of U_i and U_j are relatively prime for distinct $i, j \in \{1, \ldots, s\}$, it follows that $\mathrm{Hom}_{K[x;\theta]}(U_i, U_j) = 0$ if $i \neq j$. Hence,

$$\mathrm{End}_{K[x;\theta]} U = \mathrm{End}_{K[x;\theta]} U_1 \oplus \cdots \oplus \mathrm{End}_{K[x;\theta]} U_s. \tag{4.8}$$

Thus, the problem is reduced to studying the endomorphism rings of the primary components.

Once again, fix $i \in \{1, \ldots, s\}$, and consider the decomposition in Eq. (4.7). For $a, b \in \{1, \ldots, t_i\}$, we see that

$$\mathrm{Hom}_{K[x;\theta]}(U_{ia}, U_{ib}) = \big\{u \mapsto uq(x) \colon\ p_i(x)^{\lambda_{ib}-\lambda_{ia}} q(x)\big\}. \tag{4.9}$$

Note that divisibility is meaningful in $k[x;\theta]$ since this ring possesses a left Euclidean algorithm. Applying the 'checkered matrix theorem' [Pa, p. 42] to Eq. (4.7), we see that $\mathrm{End}_{K[x;\theta]} U_i$ is the epimorphic image of the ring

$$S_i(x) = \big\{\big[[F_{ab}]\big] \in M_{t_i}(R_i) \colon\ p_i(x)^{\lambda_{ib}-\lambda_{ia}} \text{ divides } F_{ab} \text{ for all } 1 \leqslant a < b \leqslant t_i\big\}. \tag{4.10}$$

Evaluating x at h, we see that $S_i(h) = \mathrm{End}_{K[x;\theta]} U_i = C_{\mathrm{End}_K U_i}(h)$.

We can now prove the crucial algorithmic result that we need for our main theorem.

Theorem 4.1. *There is a Las Vegas algorithm which, given a semilinear transformation h on a d-dimensional k-vector space V, where k is a finite field, returns $C_{\mathrm{End}_k V}(h)$ as the enveloping algebra of some set $\mathcal{X}$. The complexity of the algorithm is $O^{\sim}(\mathrm{MM}(d) + e(e-1)d^4)$ field operations, where e is the order of the field automorphism associated with h.*

Proof. The algorithm first computes the characteristic polynomial $\chi_h(x)$ of h and then factors this polynomial in $k[x;\theta]$ using the Las Vegas algorithm of Giesbrecht [Ge]. The total number of field operations needed for this is the complexity stated in the theorem.

Next, the factors of $\chi_h(x)$ are used to produce the minimum polynomial $m_h(x)$ of h factored as in Eq. (4.3). This involves raising h to successive powers, which requires $O(\mathrm{MM}(d)\cdot\log d)$ field operations. Since $V=\bigoplus_i\bigoplus_j U_{ij}$ (as defined above), all the bases for the submodules U_{ij} are computed in aggregate using $O(\mathrm{MM}(d))$ field operations. The algorithm concludes by returning canonical generators for each $S_i(x)$ block of checkered matrices. □

Remark 4.1. The context in which we shall make use of Theorem 4.1 involves a semilinear transformation h associated to a field automorphism of order at most 2. Thus, in our particular application, we construct $C_{\mathrm{End}_k V}(h)$ using $O(\mathrm{MM}(d)+d^4)$ field operations (see comment (2) on complexity in Section 1).

5. Proof of Theorem 1.1

In this section we combine the results of Sections 3 and 4 with Theorem 2.1 to establish our main algorithmic result.

Let k be a finite field of odd characteristic and let V be a k-space of dimension d. Let φ,γ be reflexive forms on V with φ nondegenerate. Take α and β to be the field automorphisms associated to φ and γ, respectively (one or both of these could be the identity automorphism).

By Lemma 3.1, $\sigma:=\gamma\varphi^{-1}$ is the $\alpha\beta$-semilinear slope of $[\varphi,\gamma]$. It follows from Lemma 3.2(ii) that $(x,x^*)\in\mathrm{Adj}(\varphi\cap\gamma)$ if and only if $x\in C_{\mathrm{End}_k V}(\sigma)$. A generating set $\mathcal{X}$ for the latter ring is constructed using Theorem 4.1 in $O(\mathrm{MM}(d)+e(e-1)d^4)$ field operations, where $(\alpha\beta)^e=1$ (so that $e=1$ or 2).

Next, for each $x\in\mathcal{X}$, by Lemma 3.2(i) there is a unique $x^*\in(\mathrm{End}_k V)^{\mathrm{op}}$ such that $\varphi(ux,v)=\varphi(u,x^*v)$, and x^* may be constructed with elementary linear algebra using $O(\mathrm{MM}(d))$ field operations. Hence, using $O(|\mathcal{X}|\cdot\mathrm{MM}(d))$ field operations we construct $\{(x,x^*)\colon\ x\in\mathcal{X}\}$, which generates $\mathrm{Adj}(\varphi\cap\gamma)$ as a $*$-algebra.

Theorem 1.1 now follows by appealing to Theorem 2.1.

6. Implementation and performance

Both the generic algorithm in [BW1], and its variant in Theorem 1.1, have been implemented as a Magma package. In order to illustrate the principal bottlenecks in the original version, and to demonstrate the improvements obtained by its slope variant, we ran more than a thousand random trials across multiple geometric types, fields, and dimensions. In this section we report some of the more significant findings from those trials. All trials were carried out at the University of Auckland on a 2.4 GHz microprocessor with 112 GB RAM.

Generic versus slope. Our trials indicate that any comparison of the two implementations ("generic" versus "slope") will not be unduly influenced by the geometric types of the pair $[\varphi,\gamma]$ of reflexive forms used as input. In this report we take φ to be alternating and γ to be symmetric over the field GF(27) (so we are computing intersections of orthogonal groups with symplectic groups).

There are a couple of reasons for highlighting this particular type of intersection. First, the performance of the generic method was already tested for this type of intersection, and reported in [BW1, Section 6]. Second, it would be natural – given our propaganda concerning our ability to handle unitary groups by the slope method – to make one of the forms Hermitian in our trials. Unfortunately, there are as yet no effective computational tools for skew polynomial rings implemented in Magma. We therefore content ourselves to consider pairs of bilinear forms that cannot be handled by the existing "slope-like" method described in [BO].

Fig. 1 illustrates the overall runtimes of the two methods for 100 random choices of pair $[\varphi,\gamma]$, and random dimension in the range $10\leqslant d\leqslant 70$. As one can readily see from the plot, the timings of the slope method show a larger variance, the likely reason for which is explained below. Nevertheless,

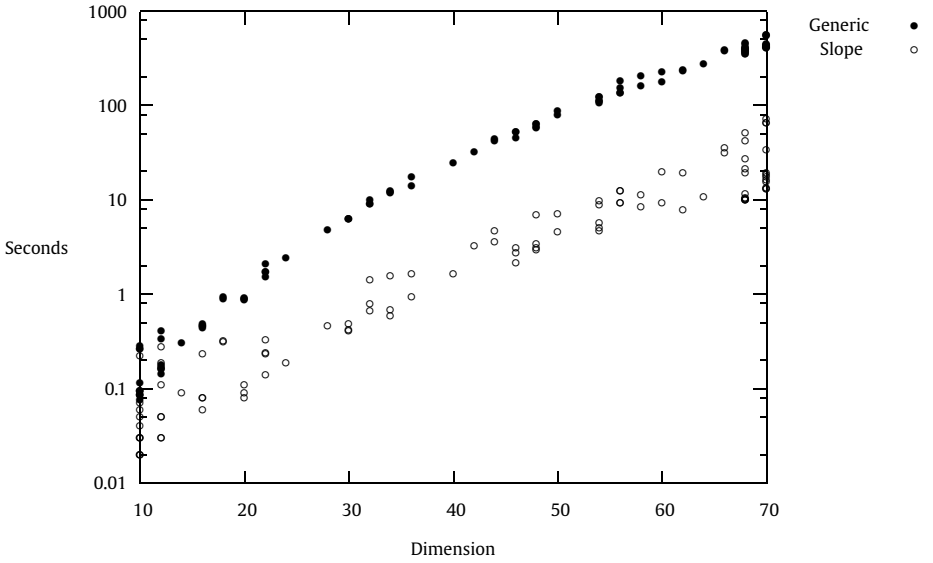

Fig. 1. Total time to compute the intersection of an orthogonal group and a symplectic group over GF(27) with varied dimensions.

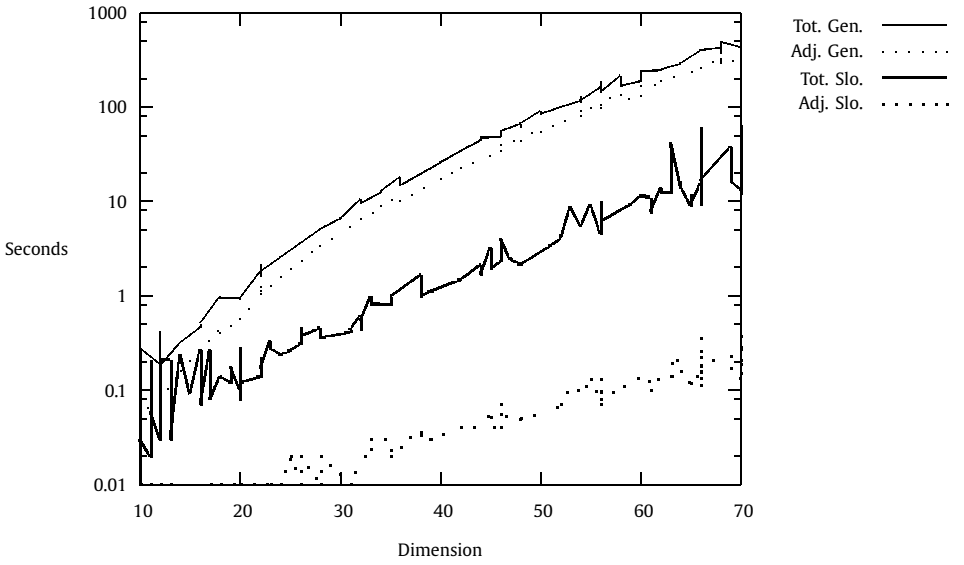

Fig. 2. In the generic method the time spent computing adjoint dominates. In the slope method this is the least costly of the main steps in the algorithm.

the furthest outliers in the slope method (beyond dimension 20) remain an order of magnitude faster than the generic method.

It is instructive in this comparison to profile the amount of time spent computing the adjoint algebra. Fig. 2 shows that, in the generic version of the algorithm, computing the adjoint algebra accounts for most of the runtime. In the slope version, however, the proportion of time spent computing adjoints becomes the least significant of the main components.

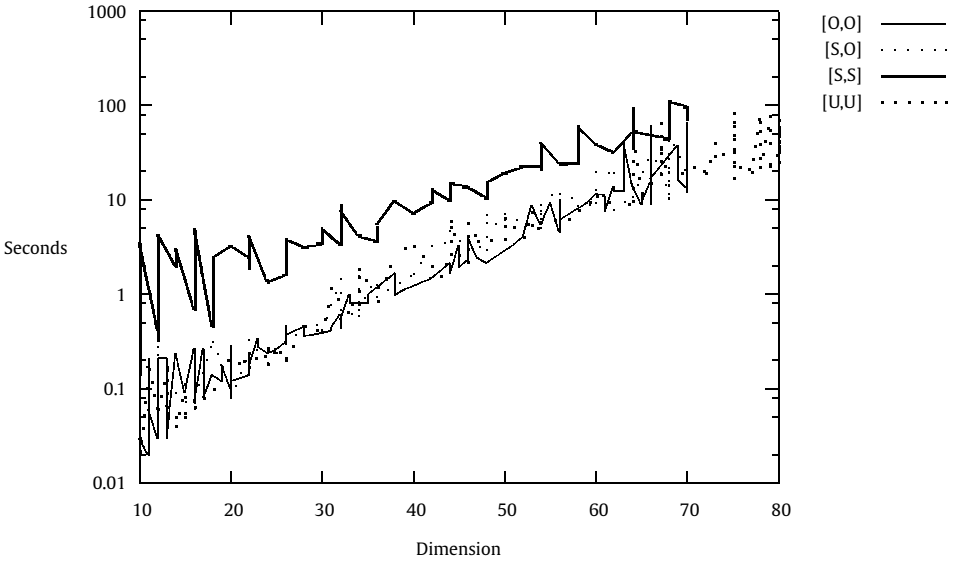

Fig. 3. Comparison of the running times for intersections of classical groups of varied types over GF(27) (or GF(25) in the unitary case).

Dependence on field and geometry. Obviously the time to compute the algebra of adjoints by the slope method (and also by the generic method) increases with the size of the defining field of the input forms. For the fields used in our trials, however, that performance is not significantly influenced by the structure of the field; the runtimes over the cubic field extension GF(27), for example, do not differ significantly from those over the prime field GF(31).

What does seem to influence the performance of the slope method are the geometric types of the input forms φ and γ, and also the geometric types of the simple $*$-ideals that comprise the semisimple quotient of $\mathrm{Adj}(\varphi \cap \gamma)$. The general reason for this is that computing adjoints is no longer a significant contributor to the overall runtime, so variations in the time spent on the other steps are more noticeable.

The specific reason that geometry influences the performance of the constructive recognition seems to be the following. The algebraic structure of the simple $*$-ideals tends generally to be low-degree (typically 1 or 2) matrix algebras over high degree (of the order of the input dimension) field extensions. Although this does not present difficulties from a complexity viewpoint, it does cause practical bottlenecks in the implementation. In short, Magma does not deal as effectively with fields of size 27^{50} as it does prime fields of roughly the same size. Although the identification of the simple $*$-ideals is carried out in identical fashion in both methods, the identification phase itself commands a far greater percentage of the overall runtime in the slope method. Hence fluctuations in the time spent in recognition have a more significant influence on the variance of overall runtimes in the slope method than they do in the generic method.

Furthermore, certain input configurations are more apt than others to produce simple $*$-algebras that require significant computing effort to recognize. For instance, intersections of two symplectic groups always produce $*$-simples of degree at least 2 (never degree 1), and indeed such intersections tend to require more time to compute. Fig. 3 provides an illustration of the variation in runtime over input geometric type.

7. Limitations and examples

Our motivation in formulating the notion of "slope" was to provide, for as broad a range of geometric types of form pair $[\varphi, \gamma]$ as possible, an explicit description of $\mathrm{Adj}(\varphi \cap \gamma)$ that admits a very

effective algorithm to construct this ring. For, it is this ring that contains all of the information needed to construct $\mathrm{Isom}(\varphi) \cap \mathrm{Isom}(\gamma)$, and it is also a crucial component in other ongoing investigations, such as [BW2]. We believe that we succeeded in this goal: for any pair of reflexive forms (including Hermitian forms), one of which is nondegenerate, the target ring $\mathrm{Adj}(\varphi \cap \gamma)$ may be realized as the centralizing ring of the slope transformation that we have defined here.

While we have not been able to find a slope method to incorporate pairs $[\varphi, \gamma]$ where both forms are degenerate, we now see that any attempt to do so using the centralizer of a single (semi)linear transformation is doomed to failure. We conclude by giving two constructions of form pairs $[\varphi, \gamma]$ that are flat (i.e. have no slope) and, more significantly, where $\mathrm{Adj}(\varphi \cap \gamma)$ is not a centralizing ring, even of a set of semilinear transformations.

Degenerate forms. Consider two nonzero alternating forms φ and γ on a 3-dimensional space V over the field $\mathbb{Z}/p$, where p is an odd prime. As V has odd dimension, both forms are degenerate and, unless their radicals coincide, there is a basis for V relative to which

$$\varphi(u, v) = u \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} v^t, \qquad \gamma(u, v) = u \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix} v^t.$$

Observe that $\varphi \cap \gamma$ is nondegenerate and alternating. Also,

$$\mathrm{Adj}(\varphi \cap \gamma) = \left\{ \left(\begin{bmatrix} a & c & d \\ 0 & b & 0 \\ 0 & 0 & b \end{bmatrix}, \begin{bmatrix} b & -c & -d \\ 0 & a & 0 \\ 0 & 0 & a \end{bmatrix} \right) : a, b, c, d \in k \right\}. \tag{7.1}$$

If the restriction A of $\mathrm{Adj}(\varphi \cap \gamma)$ to the first component is the centralizer of some $S \subseteq \mathbb{M}_3(k)$, then $A = C_{\mathbb{M}_3(k)}(C_{\mathbb{M}_3(k)}(A))$. However, $C_{\mathbb{M}_3(k)}(A) = \{aI_3 : a \in k\}$, so that $C_{\mathbb{M}_3(k)}(C_{\mathbb{M}_3(k)}(A)) = \mathbb{M}_3(k) \neq A$.

On the other hand there are flat pairs of forms whose adjoint algebra is a centralizer subring. For example,

$$\mathrm{Adj}\left(\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \cap \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right) = \left\{ \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix} : a, b \in k \right\} = C_{\mathbb{M}_2(k)} \left(\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \right). \tag{7.2}$$

In fact this example shows that it is possible to have a flat pair $[\varphi, \gamma]$ and a sloped pair $[\varphi', \gamma']$ such that $\varphi \cap \gamma = \varphi' \cap \gamma'$.

Classical intersections over different fields. In general, intersections of classical groups over different fields are impossible to manage by slope methods. For instance, let $p > 3$, $k = \mathbb{Z}/p$, and let $K = k[\omega]$ be a quadratic field extension of k. Let $V = K$ and define $\varphi : V \times V \to k$ and $\gamma : V \times V \to K$, where

$$\varphi(u_0 + u_1\omega, v_0 + v_1\omega) = u_0v_0 + u_1v_1,$$
$$\gamma(u_0 + u_1\omega, v_0 + v_1\omega) = \left(u_0v_0 + \omega^2 u_1v_1\right) + (u_0v_1 + u_1v_0)\omega.$$

Then φ is a symmetric k-form and γ is a Hermitian K-form and both are defined on V. Observe that $\mathrm{Adj}(\varphi \cap \gamma) = \{v \mapsto sv : s \in \mathbb{Z}/p\}$, so that $\mathrm{Adj}(\varphi \cap \gamma)$ is a centralizer. For $p > 3$, however, it is not the centralizer of a single linear transformation. It is therefore impossible to construct $\mathrm{GO}^+(2, p) \cap \mathrm{GU}(1, p)$ by centralizing one transformation. We note that the number of field operations needed to centralize a bounded set of transformations containing more than one element is $O(\mathrm{MM}(d^2))$, so there is no advantage to be gained over our generic algorithm in such situations.

We remark that the generic algorithm of [BW1] handles each of the examples described above, though rather less efficiently than the new algorithm handles suitable inputs of comparable size.

Acknowledgments

The first author thanks the Department of Mathematics at the University of Auckland for its hospitality during his recent stay in which some of this research was conducted. We also thank G.A. Young for his advice on how best to perform the trials in a statistically useful fashion.

References

[BF] E. Bayer-Fluckiger, Principe de Hasse faible pour les systèmes de formes quadratiques, J. Reine Angew. Math. 378 (1987) 53–59.
[BCP] W. Bosma, J. Cannon, C. Playoust, The Magma algebra system. I. The user language, in: Computational Algebra and Number Theory, London, 1993, J. Symbolic Comput. 24 (3–4) (1997) 235–265.
[BO] P.A. Brooksbank, E.A. O'Brien, On intersections of classical groups, J. Group Theory 11 (4) (2008) 465–478.
[BW1] P.A. Brooksbank, J.B. Wilson, Computing isometry groups of Hermitian maps, Trans. Amer. Math. Soc. 364 (2012) 1975–1996.
[BW2] P.A. Brooksbank, J.B. Wilson, The nilpotent groups of co-rank 2, in preparation.
[Ge] M. Giesbrecht, Factoring in skew-polynomial rings over finite fields, J. Symbolic Comput. 26 (4) (1998) 463–486.
[GG1] D. Goldstein, R.M. Guralnick, Alternating forms and self-adjoint operators, J. Algebra 308 (1) (2007) 330–349.
[GG2] D. Goldstein, R.M. Guralnick, Intersections of Symplectic groups, in preparation.
[Iv] G. Ivanyos, Fast randomized algorithms for the structure of matrix algebras over finite fields (extended abstract), in: Proceedings of the 2000 International Symposium on Symbolic and Algebraic Computation (St. Andrews), ACM, New York, 2000, pp. 175–183.
[Ja] N. Jacobson, The Theory of Rings, Amer. Math. Soc. Math. Surveys, vol. 1, Amer. Math. Soc., New York, 1953.
[Pa] D.S. Passman, A Course in Ring Theory, The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1991.
[Tay] D.E. Taylor, The Geometry of the Classical Groups, Sigma Series in Pure Mathematics, vol. 9, Heldermann Verlag, Berlin, 1992.
[vzGG] J. von zur Gathen, J. Gerhard Modern, Computer Algebra, 2nd ed., Cambridge University Press, Cambridge, 2003.