Torsion closures of ideals

Kiran S. Kedlaya

Department of Mathematics, University of California, San Diego
kedlaya@ucsd.edu
http://kskedlaya.org/slides/

Global Virtual Sage Days 109
May 28, 2020

Supported by NSF (grant DMS-1802161) and UC San Diego (Warschawski Professorship). Thanks also to the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation.

The torsion closure of an affine scheme

By a **torsion point** of the affine space $\operatorname{Spec} K[x_1, \dots, x_n]$ over a field K of characteristic 0, we will mean a closed point underlying a geometric point where $x_1, \dots, x_n$ are all roots of unity (i.e., a torsion point of the torus $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$).

Let X be a subscheme of $\operatorname{Spec} K[x_1, \dots, x_n]$. We define the **torsion closure** of X to be the reduced closed subscheme corresponding to the Zariski closure of the set of torsion points of X.

Let I be an ideal of $K[x_1, \dots, x_n]$. The **torsion closure**[*] of I is the (radical) ideal whose zero set is the torsion closure of the zero set of I.

We may also define torsion closures for ideals of $K[x_1^{\pm}, \dots, x_n^{\pm}]$. If I is an ideal of $K[x_1, \dots, x_n]$, its torsion closure is saturated with respect to $(x_1 \cdots x_n)$ and its extension to $K[x_1^{\pm}, \dots, x_n^{\pm}]$ is the torsion closure of $IK[x_1^{\pm}, \dots, x_n^{\pm}]$.

[*] This is clearly an idempotent operation, so the term "closure" is appropriate.

The torsion closure of an affine scheme

By a **torsion point** of the affine space $\operatorname{Spec} K[x_1, \dots, x_n]$ over a field K of characteristic 0, we will mean a closed point underlying a geometric point where $x_1, \dots, x_n$ are all roots of unity (i.e., a torsion point of the torus $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$).

Let X be a subscheme of $\operatorname{Spec} K[x_1, \dots, x_n]$. We define the **torsion closure** of X to be the reduced closed subscheme corresponding to the Zariski closure of the set of torsion points of X.

Let I be an ideal of $K[x_1, \dots, x_n]$. The **torsion closure**[*] of I is the (radical) ideal whose zero set is the torsion closure of the zero set of I.

We may also define torsion closures for ideals of $K[x_1^{\pm}, \dots, x_n^{\pm}]$. If I is an ideal of $K[x_1, \dots, x_n]$, its torsion closure is saturated with respect to $(x_1 \cdots x_n)$ and its extension to $K[x_1^{\pm}, \dots, x_n^{\pm}]$ is the torsion closure of $IK[x_1^{\pm}, \dots, x_n^{\pm}]$.

[*] This is clearly an idempotent operation, so the term “closure” is appropriate.

The torsion closure of an affine scheme

By a **torsion point** of the affine space $\operatorname{Spec} K[x_1, \dots, x_n]$ over a field K of characteristic 0, we will mean a closed point underlying a geometric point where $x_1, \dots, x_n$ are all roots of unity (i.e., a torsion point of the torus $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$).

Let X be a subscheme of $\operatorname{Spec} K[x_1, \dots, x_n]$. We define the **torsion closure** of X to be the reduced closed subscheme corresponding to the Zariski closure of the set of torsion points of X.

Let I be an ideal of $K[x_1, \dots, x_n]$. The **torsion closure**[*] of I is the (radical) ideal whose zero set is the torsion closure of the zero set of I.

We may also define torsion closures for ideals of $K[x_1^{\pm}, \dots, x_n^{\pm}]$. If I is an ideal of $K[x_1, \dots, x_n]$, its torsion closure is saturated with respect to $(x_1 \cdots x_n)$ and its extension to $K[x_1^{\pm}, \dots, x_n^{\pm}]$ is the torsion closure of $IK[x_1^{\pm}, \dots, x_n^{\pm}]$.

[*] This is clearly an idempotent operation, so the term "closure" is appropriate.

The torsion closure of an affine scheme

By a **torsion point** of the affine space $\operatorname{Spec} K[x_1, \dots, x_n]$ over a field K of characteristic 0, we will mean a closed point underlying a geometric point where $x_1, \dots, x_n$ are all roots of unity (i.e., a torsion point of the torus $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$).

Let X be a subscheme of $\operatorname{Spec} K[x_1, \dots, x_n]$. We define the **torsion closure** of X to be the reduced closed subscheme corresponding to the Zariski closure of the set of torsion points of X.

Let I be an ideal of $K[x_1, \dots, x_n]$. The **torsion closure**[*] of I is the (radical) ideal whose zero set is the torsion closure of the zero set of I.

We may also define torsion closures for ideals of $K[x_1^{\pm}, \dots, x_n^{\pm}]$. If I is an ideal of $K[x_1, \dots, x_n]$, its torsion closure is saturated with respect to $(x_1 \cdots x_n)$ and its extension to $K[x_1^{\pm}, \dots, x_n^{\pm}]$ is the torsion closure of $IK[x_1^{\pm}, \dots, x_n^{\pm}]$.

[*] This is clearly an idempotent operation, so the term "closure" is appropriate.

A structure theorem for torsion closures

Theorem (Laurent, 1984)

An integral subscheme of $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$ *is equal to its torsion closure if and only if it is a translate of a subtorus by a torsion point. Consequently, the torsion closure of any subscheme can be written as a finite union of such subschemes.*

Aside: the analogous statement for abelian varieties is Raynaud's theorem (formerly the **Manin-Mumford conjecture**).

Over an algebraically closed field, this says that every associated prime of the torsion closure of an ideal is generated by binomials of the form $x_1^{e_1} \cdots x_n^{e_n} - \zeta$ where ζ is a root of unity.

Concrete reinterpretation: given a system of polynomial equations to be solved in roots of unity, one can always exhibit a minimal finite set of parametric solutions. How to do this practice?

A structure theorem for torsion closures

Theorem (Laurent, 1984)

An integral subscheme of $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$ *is equal to its torsion closure if and only if it is a translate of a subtorus by a torsion point. Consequently, the torsion closure of any subscheme can be written as a finite union of such subschemes.*

Aside: the analogous statement for abelian varieties is Raynaud's theorem (formerly the **Manin-Mumford conjecture**).

Over an algebraically closed field, this says that every associated prime of the torsion closure of an ideal is generated by binomials of the form $x_1^{e_1} \cdots x_n^{e_n} - \zeta$ where ζ is a root of unity.

Concrete reinterpretation: given a system of polynomial equations to be solved in roots of unity, one can always exhibit a minimal finite set of parametric solutions. How to do this practice?

A structure theorem for torsion closures

Theorem (Laurent, 1984)

An integral subscheme of $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$ *is equal to its torsion closure if and only if it is a translate of a subtorus by a torsion point. Consequently, the torsion closure of any subscheme can be written as a finite union of such subschemes.*

Aside: the analogous statement for abelian varieties is Raynaud's theorem (formerly the **Manin-Mumford conjecture**).

Over an algebraically closed field, this says that every associated prime of the torsion closure of an ideal is generated by binomials of the form $x_1^{e_1} \cdots x_n^{e_n} - \zeta$ where ζ is a root of unity.

Concrete reinterpretation: given a system of polynomial equations to be solved in roots of unity, one can always exhibit a minimal finite set of parametric solutions. How to do this practice?

A structure theorem for torsion closures

Theorem (Laurent, 1984)

An integral subscheme of $\operatorname{Spec} K[x_1^{\pm}, \dots, x_n^{\pm}]$ *is equal to its torsion closure if and only if it is a translate of a subtorus by a torsion point. Consequently, the torsion closure of any subscheme can be written as a finite union of such subschemes.*

Aside: the analogous statement for abelian varieties is Raynaud's theorem (formerly the **Manin-Mumford conjecture**).

Over an algebraically closed field, this says that every associated prime of the torsion closure of an ideal is generated by binomials of the form $x_1^{e_1} \cdots x_n^{e_n} - \zeta$ where ζ is a root of unity.

Concrete reinterpretation: given a system of polynomial equations to be solved in roots of unity, one can always exhibit a minimal finite set of parametric solutions. How to do this practice?

Example: a theorem of Conway–Jones

Theorem (Conway-Jones, 1979)

Let S be a set of at most 9 roots of unity with zero sum. Suppose that S does not contain $\{\alpha, -\alpha\}$ or $\{\alpha, \zeta_3\alpha, \zeta_3^2\alpha\}$ for any α. Then up to rotation, S is one of

$$
\begin{gathered}
\{1, \zeta_5, \zeta_5^2, \zeta_5^3, \zeta_5^4\} \qquad \{-\zeta_3, -\zeta_3^2, \zeta_5, \zeta_5^2, \zeta_5^3, \zeta_5^4\}, \\
\{1, \zeta_7, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, \zeta_7^6\}, \qquad \{1, \zeta_5, \zeta_5^4, -\zeta_3\zeta_5^2, -\zeta_3^2\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3\}, \\
\{-\zeta_3, -\zeta_3^2, \zeta_7, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, \zeta_7^6\}, \\
\{\zeta_5, \zeta_5^4, -\zeta_3, -\zeta_3^2, -\zeta_3\zeta_5^2, -\zeta_3^2\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3\}, \\
\{1, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, -\zeta_3\zeta_7, -\zeta_3^2\zeta_7, -\zeta_3\zeta_7^6, -\zeta_3^2\zeta_7^6\}, \\
\{1, -\zeta_3\zeta_5, -\zeta_3^2\zeta_5, -\zeta_3\zeta_5^2, -\zeta_3\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3, -\zeta_3\zeta_5^4, -\zeta_3^2\zeta_5^4\}.
\end{gathered}
$$

This improves a similar result of Włodarski, by providing a technique for resolving similar problems for **any** fixed number of roots of unity.

Example: a theorem of Conway–Jones

Theorem (Conway-Jones, 1979)

Let S be a set of at most 9 roots of unity with zero sum. Suppose that S does not contain $\{\alpha, -\alpha\}$ or $\{\alpha, \zeta_3\alpha, \zeta_3^2\alpha\}$ for any α. Then up to rotation, S is one of

$$\{1, \zeta_5, \zeta_5^2, \zeta_5^3, \zeta_5^4\} \qquad \{-\zeta_3, -\zeta_3^2, \zeta_5, \zeta_5^2, \zeta_5^3, \zeta_5^4\},$$
$$\{1, \zeta_7, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, \zeta_7^6\}, \qquad \{1, \zeta_5, \zeta_5^4, -\zeta_3\zeta_5^2, -\zeta_3^2\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3\},$$
$$\{-\zeta_3, -\zeta_3^2, \zeta_7, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, \zeta_7^6\},$$
$$\{\zeta_5, \zeta_5^4, -\zeta_3, -\zeta_3^2, -\zeta_3\zeta_5^2, -\zeta_3^2\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3\},$$
$$\{1, \zeta_7^2, \zeta_7^3, \zeta_7^4, \zeta_7^5, -\zeta_3\zeta_7, -\zeta_3^2\zeta_7, -\zeta_3\zeta_7^6, -\zeta_3^2\zeta_7^6\},$$
$$\{1, -\zeta_3\zeta_5, -\zeta_3^2\zeta_5, -\zeta_3\zeta_5^2, -\zeta_3\zeta_5^2, -\zeta_3\zeta_5^3, -\zeta_3^2\zeta_5^3, -\zeta_3\zeta_5^4, -\zeta_3^2\zeta_5^4\}.$$

This improves a similar result of Włodarski, by providing a technique for resolving similar problems for **any** fixed number of roots of unity.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Applications of Conway–Jones

A literature search turns up applications of Conway–Jones in areas such as:

- Euclidean and non-Euclidean geometry;
- operator algebras;
- representation theory of finite groups;
- Kähler geometry;
- knot theory;
- dynamical systems;
- graph theory;

and more.

Ingredients of the Conway–Jones theorem

The key lemma of Conway–Jones is the following.

Lemma (Conway–Jones)

Let S be a (multi)set of roots of unity with zero sum, which is minimal: no nonempty proper (multi)subset of S sums to zero.

(a) *There exists a rotation of S consisting of roots of unity of squarefree order.*

(b) *Let N be the minimal integer for which S admits a rotation consisting of N-th roots of unity. Then*

$$\#S \geq 2 + \sum_{p|N} (p-2).$$

Both statements follow from basic facts about cyclotomic fields.

Ingredients of the Conway–Jones theorem

The key lemma of Conway–Jones is the following.

Lemma (Conway–Jones)

Let S be a (multi)set of roots of unity with zero sum, which is minimal: no nonempty proper (multi)subset of S sums to zero.

(a) *There exists a rotation of S consisting of roots of unity of squarefree order.*

(b) *Let N be the minimal integer for which S admits a rotation consisting of N-th roots of unity. Then*

$$\#S \geq 2 + \sum_{p|N} (p-2).$$

Both statements follow from basic facts about cyclotomic fields.

Ingredients of the Conway–Jones theorem

The key lemma of Conway–Jones is the following.

Lemma (Conway–Jones)

Let S be a (multi)set of roots of unity with zero sum, which is minimal: no nonempty proper (multi)subset of S sums to zero.

(a) *There exists a rotation of S consisting of roots of unity of squarefree order.*

(b) *Let N be the minimal integer for which S admits a rotation consisting of N-th roots of unity. Then*

$$\#S \geq 2 + \sum_{p|N} (p-2).$$

Both statements follow from basic facts about cyclotomic fields.

Ingredients of the Conway–Jones theorem

The key lemma of Conway–Jones is the following.

Lemma (Conway–Jones)

Let S be a (multi)set of roots of unity with zero sum, which is minimal: no nonempty proper (multi)subset of S sums to zero.

(a) *There exists a rotation of S consisting of roots of unity of squarefree order.*

(b) *Let N be the minimal integer for which S admits a rotation consisting of N-th roots of unity. Then*

$$\#S \geq 2 + \sum_{p|N} (p-2).$$

Both statements follow from basic facts about cyclotomic fields.

Ingredients of the Conway–Jones theorem

The key lemma of Conway–Jones is the following.

Lemma (Conway–Jones)

Let S be a (multi)set of roots of unity with zero sum, which is minimal: no nonempty proper (multi)subset of S sums to zero.

(a) *There exists a rotation of S consisting of roots of unity of squarefree order.*

(b) *Let N be the minimal integer for which S admits a rotation consisting of N-th roots of unity. Then*

$$\#S \geq 2 + \sum_{p|N} (p-2).$$

Both statements follow from basic facts about cyclotomic fields.

An algorithm for computing torsion closures

The Conway-Jones lemma can **in principle** be used to classify minimal additive relations among any number of roots of unity, and hence to give an algorithm to compute torsion closures (Leroux).

However, this does not seem to be feasible in practice except in very simple cases. The record extension of the Conway-Jones classification is for 12 roots of unity (Poonen–Rubinstein, 1998).

An algorithm for computing torsion closures

The Conway-Jones lemma can **in principle** be used to classify minimal additive relations among any number of roots of unity, and hence to give an algorithm to compute torsion closures (Leroux).

However, this does not seem to be feasible in practice except in very simple cases. The record extension of the Conway-Jones classification is for 12 roots of unity (Poonen–Rubinstein, 1998).

Another algorithm for computing torsion closures

An approach more in the spirit of computational commutative algebra is suggested by papers of Beukers–Smyth and Aliev–Smyth.

Lemma (Beukers–Smyth, 2000)

Let I be a maximal ideal of $\mathbb{Q}[x_1, \dots, x_n]$ whose zero set is a torsion point.

(a) *If the torsion order is divisible by 4, then there exists*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \quad f(x_i) = (-1)^{e_i} x_i \, (e_i \in \{0, 1\})$$

such that $f(I) \subseteq I$.

(b) *If the torsion order is odd*[a], *then the homomorphism*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \qquad f(x_i) = x_i^2$$

has the property that $f(I) \subseteq I$.

[a] If the order is 2 mod 4, apply some map as in (a) to get an odd-order point.

Another algorithm for computing torsion closures

An approach more in the spirit of computational commutative algebra is suggested by papers of Beukers–Smyth and Aliev–Smyth.

Lemma (Beukers–Smyth, 2000)

Let I be a maximal ideal of $\mathbb{Q}[x_1, \dots, x_n]$ whose zero set is a torsion point.

(a) *If the torsion order is divisible by 4, then there exists*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \quad f(x_i) = (-1)^{e_i} x_i \, (e_i \in \{0, 1\})$$

such that $f(I) \subseteq I$.

(b) *If the torsion order is odd[a], then the homomorphism*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \qquad f(x_i) = x_i^2$$

has the property that $f(I) \subseteq I$.

[a] If the order is 2 mod 4, apply some map as in (a) to get an odd-order point.

Another algorithm for computing torsion closures

An approach more in the spirit of computational commutative algebra is suggested by papers of Beukers–Smyth and Aliev–Smyth.

Lemma (Beukers–Smyth, 2000)

Let I be a maximal ideal of $\mathbb{Q}[x_1, \dots, x_n]$ whose zero set is a torsion point.

(a) *If the torsion order is divisible by 4, then there exists*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \quad f(x_i) = (-1)^{e_i} x_i \, (e_i \in \{0, 1\})$$

such that $f(I) \subseteq I$.

(b) *If the torsion order is odd[a], then the homomorphism*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \qquad f(x_i) = x_i^2$$

has the property that $f(I) \subseteq I$.

[a] If the order is 2 mod 4, apply some map as in (a) to get an odd-order point.

Another algorithm for computing torsion closures

An approach more in the spirit of computational commutative algebra is suggested by papers of Beukers–Smyth and Aliev–Smyth.

Lemma (Beukers–Smyth, 2000)

Let I be a maximal ideal of $\mathbb{Q}[x_1, \dots, x_n]$ whose zero set is a torsion point.

(a) *If the torsion order is divisible by 4, then there exists*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \quad f(x_i) = (-1)^{e_i} x_i \, (e_i \in \{0, 1\})$$

such that $f(I) \subseteq I$.

(b) *If the torsion order is odd*[a]*, then the homomorphism*

$$f : \mathbb{Q}[x_1, \dots, x_n] \mapsto \mathbb{Q}[x_1, \dots, x_n], \qquad f(x_i) = x_i^2$$

has the property that $f(I) \subseteq I$.

[a] If the order is 2 mod 4, apply some map as in (a) to get an odd-order point.

A hybrid algorithm

The algorithm derived from Aliev–Beukers–Smyth is quite practical for $n = 2$, but seems infeasible even for simple examples with $n = 3$. We gain some extra ground by combining with Conway–Jones, or more precisely a mod-2 variant suggested by Poonen; this seems to work better for $n = 3$.

Motivating problem (joint work with Kolpakov–Poonen–Rubinstein): which tetrahedra in $\mathbb{R}^3$ have all six dihedral angles being rational multiples of π? This is a six-variable problem, but mod-2 considerations reduce it to a large collection of three-variable problems.

A hybrid algorithm

The algorithm derived from Aliev–Beukers–Smyth is quite practical for $n = 2$, but seems infeasible even for simple examples with $n = 3$. We gain some extra ground by combining with Conway–Jones, or more precisely a mod-2 variant suggested by Poonen; this seems to work better for $n = 3$.

Motivating problem (joint work with Kolpakov–Poonen–Rubinstein): which tetrahedra in $\mathbb{R}^3$ have all six dihedral angles being rational multiples of π? This is a six-variable problem, but mod-2 considerations reduce it to a large collection of three-variable problems.

An implementation in Sage

I have a satisfactory (for this problem) implementation of the hybrid approach in Sage. It uses:

- Singular for the underlying commutative algebra (Gröbner bases, ideal membership testing, saturation, associated primes);
- Pari for arithmetic on univariate polynomials over cyclotomic fields;
- Cython for a few low-level steps (e.g., making toric changes of coordinates on a Laurent polynomial ring).

An implementation in Sage

I have a satisfactory (for this problem) implementation of the hybrid approach in Sage. It uses:

- Singular for the underlying commutative algebra (Gröbner bases, ideal membership testing, saturation, associated primes);
- Pari for arithmetic on univariate polynomials over cyclotomic fields;
- Cython for a few low-level steps (e.g., making toric changes of coordinates on a Laurent polynomial ring).

An implementation in Sage

I have a satisfactory (for this problem) implementation of the hybrid approach in Sage. It uses:

- Singular for the underlying commutative algebra (Gröbner bases, ideal membership testing, saturation, associated primes);
- Pari for arithmetic on univariate polynomials over cyclotomic fields;
- Cython for a few low-level steps (e.g., making toric changes of coordinates on a Laurent polynomial ring).

An implementation in Sage

I have a satisfactory (for this problem) implementation of the hybrid approach in Sage. It uses:

- Singular for the underlying commutative algebra (Gröbner bases, ideal membership testing, saturation, associated primes);
- Pari for arithmetic on univariate polynomials over cyclotomic fields;
- Cython for a few low-level steps (e.g., making toric changes of coordinates on a Laurent polynomial ring).

Some issues that arose: Laurent polynomials

- Ideals in Laurent polynomial rings are not implemented! I developed an *ad hoc* approach, but some further thought is required (see trac #29512). In particular, ideals in Laurent polynomial rings correspond to saturated ideals in ordinary polynomial rings, but saturation can be quite expensive; it is useful to postpone it when possible.
- Profiling suggests that element creation for Laurent polynomials is rather inefficient compared to ordinary polynomials.
- Just as for ordinary polynomials, there is a distinction between a univariate Laurent polynomial ring and a multivariate Laurent polynomial ring which happens to have only one generator. Fair enough, but the latter sometimes spontaneously changes into the former (e.g., when calling `change_ring`).
- There is plenty more missing parallelism between ordinary and Laurent polynomials (see trac #29474, #29688, and more to follow).

Some issues that arose: Laurent polynomials

- Ideals in Laurent polynomial rings are not implemented! I developed an *ad hoc* approach, but some further thought is required (see trac #29512). In particular, ideals in Laurent polynomial rings correspond to saturated ideals in ordinary polynomial rings, but saturation can be quite expensive; it is useful to postpone it when possible.
- Profiling suggests that element creation for Laurent polynomials is rather inefficient compared to ordinary polynomials.
- Just as for ordinary polynomials, there is a distinction between a univariate Laurent polynomial ring and a multivariate Laurent polynomial ring which happens to have only one generator. Fair enough, but the latter sometimes spontaneously changes into the former (e.g., when calling `change_ring`).
- There is plenty more missing parallelism between ordinary and Laurent polynomials (see trac #29474, #29688, and more to follow).

Some issues that arose: Laurent polynomials

- Ideals in Laurent polynomial rings are not implemented! I developed an *ad hoc* approach, but some further thought is required (see trac #29512). In particular, ideals in Laurent polynomial rings correspond to saturated ideals in ordinary polynomial rings, but saturation can be quite expensive; it is useful to postpone it when possible.
- Profiling suggests that element creation for Laurent polynomials is rather inefficient compared to ordinary polynomials.
- Just as for ordinary polynomials, there is a distinction between a univariate Laurent polynomial ring and a multivariate Laurent polynomial ring which happens to have only one generator. Fair enough, but the latter sometimes spontaneously changes into the former (e.g., when calling `change_ring`).
- There is plenty more missing parallelism between ordinary and Laurent polynomials (see trac #29474, #29688, and more to follow).

Some issues that arose: Laurent polynomials

- Ideals in Laurent polynomial rings are not implemented! I developed an *ad hoc* approach, but some further thought is required (see trac #29512). In particular, ideals in Laurent polynomial rings correspond to saturated ideals in ordinary polynomial rings, but saturation can be quite expensive; it is useful to postpone it when possible.
- Profiling suggests that element creation for Laurent polynomials is rather inefficient compared to ordinary polynomials.
- Just as for ordinary polynomials, there is a distinction between a univariate Laurent polynomial ring and a multivariate Laurent polynomial ring which happens to have only one generator. Fair enough, but the latter sometimes spontaneously changes into the former (e.g., when calling `change_ring`).
- There is plenty more missing parallelism between ordinary and Laurent polynomials (see trac #29474, #29688, and more to follow).

Some issues that arose: Laurent polynomials

- Ideals in Laurent polynomial rings are not implemented! I developed an *ad hoc* approach, but some further thought is required (see trac #29512). In particular, ideals in Laurent polynomial rings correspond to saturated ideals in ordinary polynomial rings, but saturation can be quite expensive; it is useful to postpone it when possible.
- Profiling suggests that element creation for Laurent polynomials is rather inefficient compared to ordinary polynomials.
- Just as for ordinary polynomials, there is a distinction between a univariate Laurent polynomial ring and a multivariate Laurent polynomial ring which happens to have only one generator. Fair enough, but the latter sometimes spontaneously changes into the former (e.g., when calling `change_ring`).
- There is plenty more missing parallelism between ordinary and Laurent polynomials (see trac #29474, #29688, and more to follow).

Some issues that arose: cyclotomic fields

- This algorithm ends up doing computations over cyclotomic fields. Coercions between cyclotomic fields are broken: if $K_1 \subset K_2$ is a proper inclusion (and $K_1 \neq \mathbb{Q}$), then elements of K_1 coerce into K_2, but elements in the image do not coerce back into K_1 (see trac #29511).
- Univariate polynomial arithmetic over a large cyclotomic field is a severe bottleneck. Can Nemo help? This might mean porting the whole algorithm over to Julia using Singular.jl, but this should be doable.

Some issues that arose: cyclotomic fields

- This algorithm ends up doing computations over cyclotomic fields. Coercions between cyclotomic fields are broken: if $K_1 \subset K_2$ is a proper inclusion (and $K_1 \neq \mathbb{Q}$), then elements of K_1 coerce into K_2, but elements in the image do not coerce back into K_1 (see trac #29511).
- Univariate polynomial arithmetic over a large cyclotomic field is a severe bottleneck. Can Nemo help? This might mean porting the whole algorithm over to Julia using Singular.jl, but this should be doable.

Some issues that arose: cyclotomic fields

- This algorithm ends up doing computations over cyclotomic fields. Coercions between cyclotomic fields are broken: if $K_1 \subset K_2$ is a proper inclusion (and $K_1 \neq \mathbb{Q}$), then elements of K_1 coerce into K_2, but elements in the image do not coerce back into K_1 (see trac #29511).
- Univariate polynomial arithmetic over a large cyclotomic field is a severe bottleneck. Can Nemo help? This might mean porting the whole algorithm over to Julia using Singular.jl, but this should be doable.