Overview of implementing algorithms for integral closures

1 Introduction

The problems I have encountered in the area of computing integral closures are legion. When I started, with a new algorithm in hand, I expected to find that there was a canonical way of presenting an integral closure of a ring (or at least the affine domains I cared about), and that I would be comparing my algorithm to more standard algorithms for time and storage requirements only. (I had hoped that my approach was at least competitive with others, and was curious to see just how well or poorly it performed by comparison.) What I found was almost the exact opposite; there was no idea of how (or whether) a presentation should be given, with some implementations giving only the fractions defining $C(A, Q(A))$ as a ring extension of A, others defining it using a (strict) affine A-algebra presentation, and some using a more generic quotient ring, with or without the input ring A explicitly as a subring. That meant there was no concensus as to the number of variables needed, the types of relations used, the type of presentation, any induced structure on $Q(A)$ or $C(A, Q(A))$, and certainly no idea that it might be more natural to view it over a Noether normalization, P, instead of the input ring A, as I naturally did.

Early on, there were actual mistakes in various implementations; that is, stopping criteria that output a ring somewhere short of the correct answer. Most such mistakes have been dealt with by now, though there are cases when one should expect various implementations never to give any answer at all, given the hopeless computational implementation of certain simple theoretical ideas. And there was/is certainly no evidence that any of this has been sufficiently tested by examples having expected input/output pairs *with a discussion of the usefulness of the output*. Typically one finds a test input and probably a running time, maybe with comments about excessive time and/or storage useage, but certainly with no expected output (save what the implementation being tested produces) and certainly no suggestion of what information the presentation gives, or indeed if it is even readable in any sense. (This appalling situation suggested to me that there was no one actually doing any qualitative testing at all in this area. I could be wrong about this; but I can hardly believe so, given what is still accepted as output in various examples.)

The accompanying examples (in separate files on this website) are an attempt to remedy this situation. So they contain examples done by both of my implementations (in MAGMA and MACAULAY2) of my *qth-power algorithm* with explanation of what is important about the form of the output, together with my working of the same examples using the various implementations I have found (in MAGMA, MACAULAY2 and SINGULAR) for comparison of both form and running time. People may obviously draw their own conclusions as to the value of each approach. This set of examples is biased in the sense that I have chosen examples that can be transformed into integral extensions A of rings of free variables $P := \mathbf{F}[x_n, \ldots, x_1]$, probably with weights as well, so that my highly-structured view can be applied to give an induced highly-structured presentation.

I welcome comments/criticism of this project, especially as it regards the actual uses of integral closures by algebraic geometers/ commutative algebraists, and as it regards any misguided views I have about the subject. Further examples, especially with both input and expected output are invaluable as well. And if I have made any unjust, unsupported, and/or incorrect criticism, I need to know so that I can edit the files appropriately. (Any remarks I make usually come from frustration at a lack of communication, and are not meant to be taken as personal attacks on anyone in any way. The goal of all should be to produce better mathematics, no matter how it is done.)

2 Finiteness

Computationally, we can deal with polynomial rings in a finite number of variables, ideals that are finitely-generated, overrings that can be described as rings using a finite number of new variables, and/or modules that have finite generating sets.

So, for instance, we can consider a quotient ring such as $A := \mathbf{F}[y,x]/\langle x^3y + y^3 + x\rangle$, since it can be described using a polynomial ring in 2 variables modulo an ideal having 1 generator (which we shall more often call a relation). It is fortunate that all ideals of a multivariate polynomial ring in n free variables over a field $\mathbf{F}$ are finitely generated.

If we wish to view A as a P-module for $P := \mathbf{F}[x]$, then there is a natural finite P-module basis $(1, y, y^2)$; whereas if $P := \mathbf{F}[y]$, then there is no such since x^i, $0 \leq i$ are all standard monomials. [Note that they are not independent since there are necessarily P-syzygies among any 4 such, but there is no finite subset B of A for which every element of A is a (finite) P-linear combination of the elements of B.]

What makes these two choices so different? Any integral extension of degree d will have a finite basis of size d, whereas any non-integral extension will not.

So from a finite module perspective, it is natural to think in terms of integral extensions; but, extensions of what? It may be that a change of variables needs to be made in order to think of A as an integral extension of some subring. And then there may be many choices. The natural smallest type of subring over which A can be integral would be a Noether normalization, a ring $P := \mathbf{F}[x_n, \ldots, x_1]$ in n free variables over the field $\mathbf{F}$. And there may be many of these as well.

Having fixed such a P, it is nice to realize that not only is $C(A, Q(A))$ the maximal subring of $1/\Delta A$ for some conductor element $\Delta \in P$, but that there is a finite P-module generating set with each of its elements having leading monomial of the form a/Δ for some $a \in A$.

And it is further possible to find best choices for P once $C(A, Q(A))$ has been computed relative to one such P.

3 Monomial orderings and Normal Forms

Picking an appropriate monomial ordering is crucial in dealing with ideals in any polynomial ring, in that different monomial orderings may produce different leading monomials and different Gröbner bases, highlighting different aspects of the ring or ideal in question. What seems to be overlooked in this area is that there may be natural meaningful ways to extend a monomial ordering on P to one on an integral extension A, that in turn can be extended to elements of $Q(A)$, hence to $C(A, Q(A))$. Of course, if one is starting from A with no clue that there is a subring P to consider...

Curves in special position (or one-point form) have a natural monomial ordering extending the natural weight, the pole-order at the special point P_∞ where all the variables (representing rational homogeneous functions with no poles except possibly at P_∞) have their poles. The latter example above can be put in special position by using the variable $z := xy$ in place of x to get the integral extension $\mathbf{F}[z; y]/\langle z^3 + zy + y^5\rangle$ with $M := \begin{pmatrix} 5 & 3 \\ 1 & 0 \end{pmatrix}$ defining a weight-over-grevlex monomial ordering with $wt(z) = 5$ and $wt(y) = 3$ corresponding to the respective pole orders. The important property here is that since $wt(z^3) = 15 = wt(y^5)$, there must be an $\mathbf{F}$-linear combination of them with smaller (non-negative) weight. Clearly $wt(z^3 + y^5) = wt(-zy) = 8 < 15$. This in turn means there is at most one canonical monomial for any pole order. Here that means that z^3 is not the canonical monomial of order 15, since $NF(z^3) = -y^5 - zy \neq z^3$; whereas y^5 must be the canonical monomial of order 15 because $NF(y^5) = y^5$.

Note that in a default grevlex monomial ordering the reverse would be true, as what is being highlighted there is total degree rather than pole order. And lex orderings would be highlighting either y or z ignoring pole orders as well.

The point of this discussion is that it used to be that implementations completely ignored the monomial ordering on A, while now at least those that wish to see A as an explicit subring of $C(A, Q(A))$ have to use the input monomial ordering, but still choose a default (grevlex or lex) monomial ordering on the new variables produced, blissfully ignorant of any natural induced monomial ordering imposed on the output from the monomial ordering on the input.

From my perspective, this induced monomial ordering is more important than the input ring itself, which need not be so explicitly a part of the presentation of the output as it currently is in many implementations.

Elements b in a Gröbner basis for an ideal I should have meaning; that is, they necessarily describe how to reduce $LT(b)$ to something smaller, namely $LT(b) - b$, the first step in computing $NF(LT(b))$. In particular, if one wishes to view $C(A, Q(A))$ (which we shall henceforth call $C(P, Q(A))$ when there is a specific P and affine P-algebra presentation in mind) as an algebra over P, there should be quadratic relations over P that tell one how to multiply, that is how to reduce $y_i y_j$ to a P-lniear combination $\sum_k c_k y_k$. There may also be some P-syzygies necessary when the P-module generators are not free over P.

4 Canonical representatives for elements of $Q(A)$

Membership problems for ideals I of a multivariate polynomial ring R is traditionally done by computing a (minimal, reduced) Gröbner basis B for I, then computing a canonical remainder $NormalForm(f, I)$ after division by said basis, with $NF(f, I) = 0$ iff $f \in I$.

To see whether an element $a/b \in Q(A)$, it would be nice to compute $NF(a, \overline{I})$ and see if it is divisible by b. This is feasible if $b \in P$ and we are given a P-module generating set for $\overline{A}$ with all of its elements of the form g_i/b for the same $b \in P$ and $g_i \in A$; otherwise probably not without computing something similar.

Consider the example in the `Singular 3.7.3` file below. There

$$A := \mathbf{F}[x, y, z]/\langle x^6 + x^3z - y^3z^2 \rangle.$$

SINGULAR's `normal` function produces fractions

$$x^3/x^3,\ x^2yz/x^3,\ xy^2z/x^3,\ y^3z/x^3.$$

Since there is a common denominator $x^3 \notin P := \mathbf{F}[y, z]$, it is relatively easy to determine whether $a/x^3 \in \overline{A}$ for $a \in A$. But to do this for $(x^5 + x^2z)/(y^2z)$ or even $(y^3z)/(x^3 + z)$ is more challenging.

MACAULAY2's `icFractions` function gives

$$x,\ y,\ z,\ (x^3 + z)/z,\ y^2z/x^2.$$

Here it is even harder to show that $(x^5 + x^2z)/(y^2z) \in \overline{A}$, since it is $(y^2z/x^2)^2 - x^2y$ a non-trivial combination of the ring generators found.

With MAGMA's `Normalisation` function, it takes some work to remove excess elements to even see that $(x^4 + xz)/(yz) \in \overline{A}$, let alone anything harder.