Simplification of Nested Radicals

Susan Landau

Computer and Information Science Department
University of Massachusetts

COINS Technical Report 90-113

Simplification of Nested Radicals

Susan Landau*
Mathematics Department
Wesleyan University

COINS
University of Massachusetts

November 14, 1990

Abstract

Radical simplification is an important part of symbolic computation systems. Until now no algorithms were known for the general denesting problem. If the base field contains all roots of unity, then we give necessary and sufficient conditions for a denesting, and our algorithm computes a denesting of α when it exists. If the base field does not contain all roots of unity, then we show how to compute a denesting that is within one of optimal over the base field adjoin a single root of unity. Throughout our paper, we choose to represent a primitive l^{th} root of unity by its symbol ζ_l, rather than as a nested radical. The algorithms require computing the splitting field of the minimal polynomial of α over k, and have exponential running time.

In his magic way, Ramanujan observed a number of striking relationships between certain nested radicals:

$$\sqrt[3]{\sqrt[3]{2}-1} = \sqrt[3]{1/9}-\sqrt[3]{2/9}+\sqrt[3]{4/9}$$

$$\sqrt{\sqrt[3]{5}-\sqrt[3]{4}} = 1/3(\sqrt[3]{2}+\sqrt[3]{20}-\sqrt[3]{25})$$

$$\sqrt[6]{7\sqrt[3]{20}-19} = \sqrt[3]{5/3}-\sqrt[3]{2/3}$$

*Supported by NSF grants DMS-8807202, and CCR-8802835. Part of this work was done while the author was visiting the Yale University Math Department.

These remained innocent curiosities for decades. Then symbolic computation came of age, and such manipulations assumed greater importance. Consider the equation:

$$\sqrt{5+2\sqrt{6}} = \sqrt{2}+\sqrt{3}$$

The field $Q(\sqrt{5+2\sqrt{6}})$ has $\{1, \sqrt{5+2\sqrt{6}}, 5+2\sqrt{6}, (\sqrt{5+2\sqrt{6}})^3\}$ as a basis over Q. But $\{1, \sqrt{2}, \sqrt{3}, \sqrt{6}\}$ is also a basis for $Q(\sqrt{5+2\sqrt{6}})$ over Q. In many cases, the first basis is preferable – it is of the form $\{1, \alpha, \alpha^2, \alpha^3\}$ – but people often find the second basis easier to understand. The basis $\{1, \sqrt{2}, \sqrt{3}, \sqrt{6}\}$ is simple to manipulate. An important issue then, is the "denesting" of radicals – a term which will be precisely defined in the next section. It is also independently interesting : under what circumstances can a radical be expressed in terms of radicals with a lower depth of nesting?

In 1985, Borodin, Fagin, Hopcroft, and Tompa [3] gave an efficient algorithm for decreasing the nesting depth of a class of expressions involving square roots. Also in 1985, Zippel [18] gave some conditions under which a radical could denest. The general case remained open. It was unknown how to determine whether a radical could be denested.

We show that if the base field contains all roots of unity, then a radical can be denested iff there is a denesting which occurs within its splitting field. That is, a radical α can be denested over a field k containing all roots of unity iff there is a denesting which occurs with each term of the denesting lying within the splitting field of the minimal polynomial of α over k. If the base field does not contain all roots of unity, we show that a denesting within one of optimal can be achieved over $k(\zeta_l)$, ζ_l a primitive l^{th} root of unity, where l is the lcm of the exponents of the derived series of the Galois group of the splitting field of $k(\alpha)$ over k. We also show how to acheive an optimal denesting by adjoining a root of unity which is dependent on the presentation of α. We represent a primitive l^{th} root of unity by the symbol ζ_l rather than as a nested radical itself. This presents certain problems, which we discuss in §3.

Following the measure of the size of univariate polynomials in the factoring problem, we define the size of the denesting problem to be the minimal polynomial for α over k. (In the appendix we show how to go from the presentation as a nested radical to its minimal polynomial.) Next we show that given a radical α over a field k of characteristic 0, we can denest α in the time it takes to compute the splitting field of the minimal polynomial of α over k. In worst case this takes exponential time.

Our paper is organized as follows: §1: Background, §2 Ambiguity, §3 Algebraic Structure, §4 The Algorithm and Running Time Analysis, and §5 Conclusions, and the Appendix.

1 Background

We begin with a brief review of some algebraic concepts. The reader who is unfamiliar with this material is advised to consult [8] or [14] for algebraic number theory and Galois theory, and [13] for group theory.

Let k be a field. Throughout this paper we assume that k is of characteristic 0. An element α *is algebraic over* k iff α satisfies a polynomial with coefficients in k. The *degree* of α is the degree of the minimal irreducible polynomial of α over k. If L has finite dimension $[L : k]$ over a subfield k, then $L = k(\theta)$ for some θ in L, where the degree of θ over k equals $[L : k]$.

An extension field L is *algebraic over a field* k iff every element of L is algebraic over k. It is well known that every finite extension of a field is algebraic; the finite extensions of Q are the called the *algebraic number fields*. If k is an algebraic number field, it contains a set of elements which satisfy integer monic polynomials over Z; this is called the *ring of integers of* k, and is frequently denoted O_k.

In 1982, Lenstra, Lenstra and Lovász [11] showed that $f(x)$ in $Q[x]$ can be factored in polynomial time. If $f(x)$ is a polynomial in O_k, where $k = Q[t]/g(t)$ is a number field, then there are polynomial time algorithms for factoring $f(x)$ over O_k [6], [10].

Following [3], a *formula* over a field k and its *depth of nesting* are defined as follows:

(1) an element of k is a formula of depth 0 over k,

(2) an arithmetic combination (A+B, A × B, A/B) of formulas A and B is a formula whose depth over k is max(depth(A), depth(B)), and

(3) a root $\sqrt[n]{A}$ of a formula A is a formula whose depth over k is 1+ depth(A).

We will call such a formula a nested radical. A nesting of α means any formula A that can take α as a value[1]. We will say the formula A *can be denested over the field* k if there is a formula B of lower depth than A such that A=B. We will say that *A can be denested in the field* L if there is a formula B=A of lower nesting depth than A with all of the terms (subexpressions) of B lying in L. For any α, we define the depth of α over k

[1] Of course, an n^{th} root is a multivalued function. See section 2 for a discussion on ambiguity.

to be the depth of the minimum depth expression for α. When we are given a formula A for α such that A can be denested, we will sometimes instead say that α can be denested. We will write a primitive n^{th} root of unity as a special symbol ζ_n rather than as a nested radical, and we will define the depth of nesting for a primitive root of unity that is not already in the field to be 1.

We return to the examples mentioned earlier. At first they seem mysterious. In fact, each equation is a result of a complex algebraic structure. Consider:

$$\sqrt[3]{\sqrt[3]{2}-1} = \sqrt[3]{1/9} - \sqrt[3]{2/9} + \sqrt[3]{4/9}.$$

A natural question to ask is: What is the relationship between $\sqrt[3]{2}$ and $\sqrt[3]{9}$? There is no obvious one. But there is a relationship between $\sqrt[3]{\sqrt[3]{2}-1}$ and $\sqrt[3]{9}$; $\sqrt[3]{9}$ is an element of the field $Q(\sqrt[3]{\sqrt[3]{2}-1})$. This fact is unexpected.

Equations 2 and 3 give similar results. A natural questions arises: In which field do denestings occur?

Let k be an algebraic number field, and let $f(x)$ be an irreducible polynomial of degree n with coefficients in k, and roots $\alpha_1, \ldots, \alpha_n$. Then $k(\alpha_i) \simeq k[x]/f(x) \simeq k(\alpha_j)$, but in general $k(\alpha_i) \neq k(\alpha_j)$. The field $L = k(\alpha_1, \ldots, \alpha_n)$ is called the *splitting field of f(x) over k*. This is the smallest field containing $k(\alpha_1)$ that is Galois[2] over k. We consider the set of automorphisms of L which leave k fixed. These form a group, called the *Galois group of* L over k. As we can think of these automorphisms as permutations of the α_i, this group is sometimes called the *Galois group of* $f(x)$ *over* k. The Galois group is *transitive* on $\{\alpha_1, \ldots, \alpha_n\}$, that is, for each pair of elements α_i, α_j, there is an element σ in G, with $\sigma(\alpha_i) = \alpha_j$. Galois' great insight was the discovery of the relationship between the subgroups of G and the subfields of L containing k.

Let H be a subgroup of G. We denote by L^H the set of elements of L which are fixed pointwise by each element of H. This set forms a field. Furthermore, H fixes k, so that we have:

$$k \subseteq L^H \subseteq L$$

Conversely suppose that $J = k(\beta_1, \ldots, \beta_r)$ is a field such that $k \subset J \subset L$. Then the β_i can be written as polynomials in $\alpha_1, \ldots, \alpha_n$, and H, the

[2]We say a field $F \supseteq k$ is Galois over k if every irreducible polynomial $p(x)$ in $k[x]$ which has a root in F splits completely in F.

subgroup of G which fixes J, consists of those elements of G which fix the β_i pointwise. The relationship between the fields and the groups can be formally stated as:

Theorem 1.1 *Fundamental Theorem of Galois Theory: Let k be a field, and let $f(x)$ in $k[x]$ be a polynomial of degree n, with roots $\alpha_1, \ldots \alpha_n$. Then:*

1. *Every intermediate field J, with $k \subset J \subset L = k(\alpha_1, \ldots, \alpha_n)$ defines a subgroup H of the Galois group G, namely the set of automorphisms of L which leave J fixed. Furthermore, L is Galois over J.*
2. *The field J is uniquely determined by H, for J is the set of elements of L which are invariant under the action of H.*
3. *The subgroup H is normal iff J over k is a Galois extension. In that case the Galois group of J over k is G/H.*
4. *$|G| = [L : k]$, and $|H| = [L : J]$.*

The Galois groups we will be looking at are rather special. Our field extensions are extensions by radicals. Thus the groups we are looking at are solvable, that is, there is a sequence of subgroups $G = G_0 \supset G_1 \supset \ldots \supset G_t$, with G_{i+1} normal in G_i (written $G_{i+1} \lhd G_i$), and G_i/G_{i+1} cyclic of prime order. We introduce:

Definition: The commutator subgroup DG of G is the subgroup $\langle \sigma\tau\sigma^{-1}\tau^{-1} | \sigma, \tau \in G \rangle$. We will denote $D^2G = D(DG)$, and $D^iG = D(D^{i-1}G)$ for $i > 2$.

If G is a solvable group, there is an s such that $D^sG = \{e\}$. The sequence $D^iG/D^{i+1}G$ is called the derived series of G. The following are well-known:

Lemma 1.2 *Let $G = H_0 \supset H_1 \supset \ldots \supset H_t$ be a sequence of groups such that $H_i \lhd H_{i-1}$ and H_{i-1}/H_i is abelian. Then $H_i \supset D^iG$.*

Lemma 1.3 *The groups D^iG are normal in G for all i.*

Lemma 1.4 *If N is normal in G, then $D^i(G/N) \simeq D^i(G)N/N$.*

We say that an extension K over k is abelian if it is Galois and the resulting Galois group is abelian. Similarly, we say an extension is cyclic if it is Galois, and resulting Galois group is cyclic. If $K = k(\beta)$ for some β which satisfies an irreducible polynomial of the $x^n - b$ for some b in k, we say that K is a simple radical extension of k. A tower of such extensions will be called a radical extension. Let ζ_n denote a primitive n^{th} root of unity. The following classical theorems will prove useful:

Theorem 1.5 *Let k be a field. The following are equivalent:*

1. *α is a nested radical over k*
2. *There exists a solvable Galois extension L over k with α in L.*
3. *The splitting field of $k(\alpha)$ over k has solvable Galois group.*

This will be made more precise in Lemma 3.1, where the depth of nesting of a radical expression for α is related to the length of the derived series for the Galois group of L over k.

Theorem 1.6 *Let k be a field, with K a cyclic extension of k of degree n, and suppose ζ_n is in k. Then there is a β in K such that $K = k(\beta)$, and β satisfies $x^n - b$ for some b in k.*

Theorem 1.7 *(Hilbert's Theorem 90) Let K be a cyclic extension over k, and let σ be a generator of the Galois group G. For every element β in K with norm 1, there is an element $\gamma \neq 0$ in K such that $\beta = \gamma/\sigma(\gamma)$.*

Theorem 1.8 *Let k be a field with ζ_n in k, and suppose β is a root of $x^n - b$. Then $k(\beta)$ is cyclic over k of degree d, where d divides n, and β^d is an element of k.*

If ζ_n is not in k, the situation is not quite as simple:

Theorem 1.9 *Let k be a field, and n an integer ≥ 2. Let a be an element of k, $a \neq 0$. Assume that for all prime numbers p dividing n, that a is not a p^{th} power in k, and moreover, if $4 \mid n$, then a is not equal to $-4j^4$ for some j in k. Then $x^n - a$ is irreducible in $k[x]$.*

We are now ready to state our two main results:

Theorem 3.2 *Suppose α is a nested radical over k, where k is a field of characteristic 0 containing all roots of unity. Then there is a minimal depth nesting of α with each of its terms lying in the splitting field of the minimal polynomial of α over k.*

Theorem 3.7 *Suppose α is a nested radical over k, where k is a field of characteristic 0. Let L be the splitting field of $k(\alpha)$ over k, with Galois group G. Let l be the lcm of the exponents of the derived series of G. If there is a denesting of α such that each of the terms has depth no more than t, then*

there is a denesting of α over $k(\zeta_l)$ with each of the terms having depth no more than $t+1$ and lying in $L(\zeta_l)$.

We also have an alternative version of this result in which we can achieve minimal depth at the expense of adjoining a primitive r^{th} root of unity where r is dependent upon the presentation of the input.

Corollary 3.8 *Let k, α, L, G, l, t be as in Theorem 3.7. Let m be the lcm of the (m_{ij}), where the m_{ij} runs over all the roots appearing in the given nested expression for α. Let r be the lcm of (m, l). Then there is a minimal depth nesting of α over $k(\zeta_r)$ with each of its terms lying in $L(\zeta_r)$.*

Roots of unity, and indeed the radicals themselves, are more complicated than they might first appear. Before we prove the theorems, we discuss ambiguity of radical expressions and primitive roots of unity more closely.

2 Ambiguity

When we write the equation

$$\sqrt{5+2\sqrt{6}} = \sqrt{2}+\sqrt{3},$$

it is ambiguous. Which $\sqrt{2}$ do we mean? Which $\sqrt{3}$? The usual interpretation is the positive real roots for all four choices in the equation above. Under those choices, the equation is correct; under others, it may not be. Similarly, when we talk about denesting a nested radical, we must be careful about what we mean.

For example, suppose we are interested in denesting the expression:

$$\sqrt[3]{\sqrt[3]{2}-1} - \sqrt[3]{1/9}.$$

The polynomial x^3-9 factors over the field $Q(\sqrt[3]{\sqrt[3]{2}-1})$. To denest $\sqrt[3]{\sqrt[3]{2}-1} - \sqrt[3]{1/9}$, we need to know to which root of x^3-9 we are referring in $Q(\sqrt[3]{\sqrt[3]{2}-1})$: the one which satisfies $x - \alpha^8 - 4\alpha^5 - 4\alpha^2$, or one of the two satisfying $x^2 + (\alpha^8 + 4\alpha^5 + 4\alpha^2)x + (3\alpha^4 + 6\alpha)$, where $\alpha = \sqrt[3]{\sqrt[3]{2}-1}$.

The problem, of course, is that $\sqrt[n]{\alpha}$ is a many-valued function. Once we choose which value it has, the same value must be assigned to it every time it appears. If the roots are specified at the time a nested radical is given, there is no difficulty, and we will choose those roots. If they are not, we have two choices: we could run the denesting algorithm for all possible choices of

the values of the roots and compute the denesting in each case, or we could arbitrarily pick values. For the sake of simplicity, in this paper we choose to do the latter. When we adjoin $\sqrt[n]{\alpha}$, we do so in a way that makes the smallest (in terms of degree) field extension possible. In the above example, we would choose the $\sqrt[3]{9}$ which is already in the field.

Sometimes radicals may be given in a reducible form, e.g. $\sqrt[4]{4}$. The obvious simplification to $\sqrt{2}$ (which omits $-\sqrt{2}$ as per above) is not right, since the minimal polynomial for $\sqrt[4]{4}$ over Q has two roots $\pm i\sqrt{2}$ which this "simplification" omits.

Let $\alpha = \sqrt[n]{a}$ be a reducible radical, that is suppose $x^n - a$ factors. By Theorem 1.9 this can happen in two ways. We might have a is a p^{th} power, where p divides n. In this case, $a = A^c$ where $gcd(c, n) = d$. Then the roots of $x^n - a$ are $\zeta_n^j \sqrt[n/d]{A^{c/d}}, j = 0, \ldots, n-1$. Note that $\sqrt[n/d]{A^{c/d}}$ is an irreducible radical. Thus what we can do is replace each instance of a reducible radical ($\sqrt[n]{A^c}$) by its irreducible cousin ($\sqrt[n/d]{A^{c/d}}$), and adjoin the appropriate ζ_n to the splitting field by letting the l of Theorem 3.7 be the lcm of the indices and the exponents of the derived series. The other situation, where $4 \mid n$ and $a = -4j^4$ for some j in the base field is handled similarly. Because of the way we handle roots of unity, this will not increase the depth of nesting.

Roots of unity present a more serious problem. We have chosen to write a primitive l^{th} root of unity as ζ_l, rather than as a nested radical. The reason is that we are trying to write expressions in as simple a form as possible. In many situations ζ_l is a more meaningful expression than the nested radical it represents.

There are difficulties with this approach. A primitive l^{th} root of unity may be of $\log l$ nesting depth. The symbol ζ_l hides that complexity. Adding roots of unity to k changes k in suprising ways. By the Kronecker-Weber Theorem, every abelian extension over Q can be embedded in an cyclotomic extension. Thus, we may have the ill fortune to be expressing $\sqrt[n]{\alpha}$ over $Q(\zeta_l)$ where $\sqrt[n]{\alpha}$ is an irrational number which happens to be in $Q(\zeta_l)$. Such is the case for $\sqrt{5}$ in the field $Q(\zeta_5)$. Thus, $\sqrt{5}$ will be represented as a polynomial in ζ_5, rather than the more usual expression $\sqrt{5}$. This type of simplification may drop us a single level of nesting.

A more serious problem is that in writing a root of unity as ζ_l we are some sense masking it. There are subtle ways in which we pay for that. For example, we can write $\sqrt{\sqrt{5} - 5/2}$ as $\zeta_5 - 1/\zeta_5$. Which symbol is easier to understand: $\sqrt{\sqrt{5} - 5/2}$ or $\zeta_5 - 1/\zeta_5$? That is certainly open to debate. In

choosing to express $\sqrt{\sqrt{5}-5/2}$ as $\zeta_5 - 1/\zeta_5$ it is not clear that we have really simplified things. This problem seems unavoidable in the current approach.

If $k = Q$, then the minimal polynomial of ζ_l is the unique irreducible factor of $x^l - 1$ of degree $\varphi(l)$. If $k \neq Q$, this minimal polynomial may factor. Since the primitive l^{th} roots of unity are all powers of one another, the fields $k(\zeta_l)$ are all equal. The denested expression we find for α is dependent on our choice of ζ_l however; a different ζ_l will give a different denesting expression. When we denest a nested radical α, we will always specify the minimal polynomial of the associated primitive l^{th} root of unity.

Now we are in a position to prove our main theorems, which we do in the next section.

3 Algebraic Structure

In this section we discuss the algebraic structure surrounding nested radicals. We begin with:

Lemma 3.1 *Let α be a nested radical, and suppose α can be denested in a field $L \supseteq k$. Suppose that the denested expression has nesting depth l. Let $\tilde{L}$ be the splitting field[3] of L over k, with Galois group G, and assume that all roots of unity of $\tilde{L}$ appear in k. Then there are subgroups $H_1, \ldots, H_l$ of G, with $G = H_0 \rhd H_1 \rhd \ldots \rhd H_l$ and H_i/H_{i+1} abelian for $i = 0, \ldots l-1$, with $H \supseteq H_l$, where H is the subgroup of G associated with $k(\alpha)$. Conversely, if there is such a sequence of subgroups, then α has nesting depth at most l.*

Proof If α can be denested, then $\alpha = q(\beta_1, \ldots, \beta_t)$, with each β_i having nesting depth $l_i \leq l$ over k. Each β_i has a nesting sequence of the form:

$$\begin{aligned}
\beta_{i1} &= \sqrt[m_{i1}]{p_{i1}},\ p_{i1} \in k, \\
\beta_{i2} &= \sqrt[m_{i2}]{p_{i2}(\beta_{11}, \ldots, \beta_{t1})}, \\
\beta_{i3} &= \sqrt[m_{i3}]{p_{i3}(\beta_{11}, \ldots, \beta_{t1}, \beta_{12}, \ldots, \beta_{t2})}, \\
&\vdots \\
\beta_i = \beta_{il_i} &= \sqrt[m_{il_i}]{p_{il_i}(\beta_{11}, \ldots, \beta_{tl_t-1})},
\end{aligned}$$

[3]This is usually called the normal closure of L over k, i.e. the minimal field $\tilde{L}$ containing L in which every polynomial in $k[x]$ which has a root in L splits completely in $\tilde{L}$.

where the p_{ij}'s and q are mulitvariate polynomials over k. (Since each β_i has a possibly different depth of nesting, without loss of generality, let $l_t = l$ be the maximal depth of nesting, and let $\beta_{ij} = \beta_i$ for $j > l_i$.) Without loss of generality assume that $\sqrt[m_{ij}]{p_{ij}(\beta_{11}, \ldots, \beta_{tj-1})}$ is of degree m_{ij} over $k(\beta_{11}, \ldots, \beta_{tj-1})$. Since $x^{m_{ij}} - p_{ij}(\beta_{11}, \ldots, \beta_{tj-1})$ has a root in $\tilde{L}$, it splits completely in $\tilde{L}$. Thus we know that $\zeta_{m_{ij}}$ is in k for every pair i, j. (Since all the necessary roots of unity lie in k, we do not run into problems with ambiguity.)

Number the extensions of k in the following way: $k_1 = k(\beta_{11}, \ldots, \beta_{t1}) \subseteq k_2 = k_1(\beta_{12}, \ldots, \beta_{t2}) \subseteq \ldots \subseteq k_l = k_{l-1}(\beta_{1l}, \ldots, \beta_{tl}) = k(\beta_{11}, \ldots, \beta_{tl})$. Let H_i be the subgroup of G corresponding to k_i. Observe that k_{i+1} is a Galois extension of k_i, since the appropriate roots of unity lie in k. Thus we have $G = H_0 \rhd H_1 \rhd \ldots \rhd H_l$. Furthermore, k_{i+1} is an abelian extension of k_i, since it is contained in the composite of cyclic extensions. Thus H_i/H_{i+1} is abelian. Finally, $k(\alpha) \subset k(\beta_{11}, \ldots, \beta_{tl_t})$, thus $H \supset H_l$.

The converse follows immediately from Theorems 1.5 and 1.6 and the fact that the needed roots of unity lie in k. ∎

We are now ready to proceed with:

Theorem 3.2 *Suppose α is a nested radical over k, a field of characteristic 0 which contains all roots of unity. Then there is a minimal depth nesting of α with each of its terms lying in the splitting field of the minimal polynomial of α over k.*

Proof Let L be the splitting field of $k(\alpha)$ over k, with Galois group G. Suppose α can be denested over k with all of its terms lying in K, a field. Let $\tilde{L}$ be the normal closure of K over k; then $\tilde{L} \supset L$. Let $\tilde{G}$ be the Galois group of $\tilde{L}$ over k; then the field $\tilde{L}$ is a Galois extension over L with group $\tilde{N}$. Furthermore, by the Fundamental Theorem of Galois Theory, $G \simeq \tilde{G}/\tilde{N}$. Let H be the subgroup of G associated with $k(\alpha)$, and let $\tilde{H}$ be the pullback

of H. Then $\tilde{H} \supset \tilde{N}$.

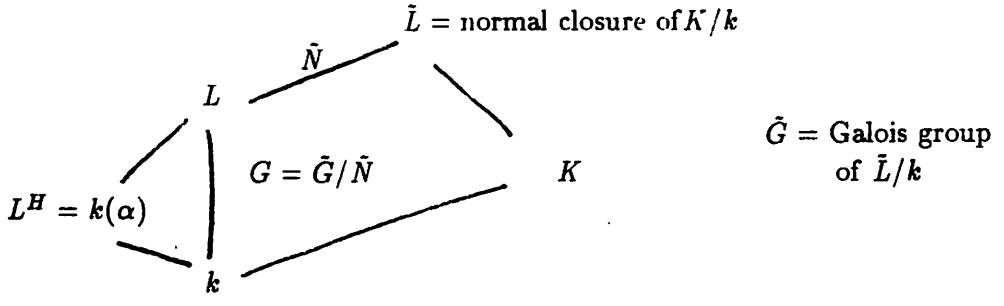

Now since α can be denested over $\tilde{L}$, by Lemma 3.1 there is a series of subgroups $\tilde{H}_1, \ldots, \tilde{H}_l$ of $\tilde{G}$ such that $\tilde{G} = \tilde{H}_0 \rhd \tilde{H}_1 \rhd \ldots \rhd \tilde{H}_l$, with $\tilde{H}_i/\tilde{H}_{i+1}$ abelian for $i = 0, \ldots, l-1$, and $\tilde{H} \supset \tilde{H}_l$. We will show how to pull this down to a sequence in G, thus showing that α can be denested in L.

Consider the sequence $\tilde{H}_1\tilde{N}, \tilde{H}_2\tilde{N}, \ldots, \tilde{H}_l\tilde{N}$. Clearly $\tilde{H}_i\tilde{N} \rhd \tilde{H}_{i+1}\tilde{N}$. That $\tilde{H}_i\tilde{N}/\tilde{H}_{i+1}\tilde{N}$ is abelian follows immediately from the facts that (i) $\tilde{H}_i/\tilde{H}_{i+1}$ is abelian, (ii) $\tilde{N}$ is normal in $\tilde{G}$. Let $a\tilde{H}_{i+1}\tilde{N}, b\tilde{H}_{i+1}\tilde{N}$ be elements of $\tilde{H}_i\tilde{N}/\tilde{H}_{i+1}\tilde{N}$. Then there are a_1, a_2, b_1, b_2 in $\tilde{H}_i$, n_1, n_2, n_3, n_4 in $\tilde{N}$ such that $a = n_1a_1 = a_2n_2$ and $b = n_3b_1 = b_2n_4$. Then $a\tilde{H}_{i+1}\tilde{N}b\tilde{H}_{i+1}\tilde{N} = b\tilde{H}_{i+1}\tilde{N}a\tilde{H}_{i+1}\tilde{N}$.

Finally, since $\tilde{H} \supset \tilde{N}$ and $\tilde{H} \supset \tilde{H}_l$, we know that $\tilde{H} \supset \tilde{H}_l\tilde{N}$.

Now consider the sequence H_i in G defined by $\tilde{H}_i\tilde{N}/\tilde{N} = H_i$. That the sequence $G = H_0, H_1, \ldots, H_l$ satisfy $H_i \rhd H_{i+1}$ for $i = 0, \ldots, l-1$ follows from the third isomorphism theorem. That H_i/H_{i+1} is abelian follows from the fact that

$$H_i/H_{i+1} = \frac{\tilde{H}_i\tilde{N}/\tilde{N}}{\tilde{H}_{i+1}\tilde{N}/\tilde{N}} \simeq \tilde{H}_i\tilde{N}/\tilde{H}_{i+1}\tilde{N}.$$

The isomorphism is a consequence of the third isomorphism theorem. That $H \supset H_l$ is clear, since $H = \tilde{H}\tilde{N}/\tilde{N} \supset \tilde{H}_l\tilde{N}/\tilde{N} = H_l$.

Thus we have taken a chain of subgroups in $\tilde{G}$ and mapped it to one over G. The fact that all the needed roots of unity lie in k means that we can apply Theorem 1.8, and we have taken a denesting over $\tilde{L}$ and mapped it to one over L. The theorem is proved. ∎

This is a pleasing theorem. Although it is subsumed by Theorem 3.7, its proof is different, and it shows how any denesting expression for α can be mapped to one in which all the subexpressions are in L, the splitting field of α over k, assuming all roots of unity lie in k. By Lemma 1.2, a sequence

of minimal nesting depth for α is found by finding the derived series, and computing the associated fields. This will be explained in detail in section 4.

In general, it will not be the case that all roots of unity lie in k. Of course, one can adjoin them, but this may be an infinite extension. From a computational standpoint this is not a good approach. Instead we will find a single primitive root of unity we can adjoin to k which will give us a denesting. We begin with:

Definition The least positive integer n such that $G^n = \{e\}$ is the exponent of G.

Theorem 3.3 *Suppose $k \subset k_1 \subset \ldots \subset k_t$ is a series of abelian extensions (k_{i+1} an abelian extension of k_i), with $K = k(\alpha) \subset k_t$. Let L be the splitting field of K over k, and L_1 be the splitting field of k_t over k, with Galois group G. Then $L \subset L_1$, and let $N = Gal(L_1/L)$. If s is minimal such that $D^sG \subset N$, then $s \leq t$. Furthermore s is the length of the derived series for $G/N = Gal(L/k)$.*

Proof We begin with some notation. Let

$$Gal(L_1/K) = H_0$$

$$Gal(L_1/k) = G$$

$$Gal(L_1/k_i) = H_i, i = 1, \ldots, t$$

$$Gal(L_1/k_t) = H_t = H.$$

Observe that $G \supset H_1 \supset \ldots \supset H_t = H$, and that $H_i \supset D^iG$, since H_{i-1}/H_i is abelian. Furthermore we know that $\bigcap_{\sigma \in G} \sigma H \sigma^{-1} = \{e\}$, since L_1 is the smallest normal extension of k_t over k. Similarly, $\bigcap_{\sigma \in G} \sigma H_0 \sigma^{-1} = N$, since L is the smallest normal extension of K over k, and $N = Gal(L_1/L)$. Note that N is a normal subgroup of G.

Now let s be the least integer such that $D^sG \subset N$. We know that $\{e\} = \bigcap_{\sigma \in G} \sigma H_t \sigma^{-1} \supset \bigcap_{\sigma \in G} \sigma D^tG \sigma^{-1} = \bigcap D^tG = D^tG$. Thus $D^tG = \{e\}$. This implies that $D^tG \subset D^sG$, or that $s \leq t$. Furthermore note that $D^i(G/N) \simeq D^i(G)N/N$ which implies that $D^s(G/N) = \{e\}$. Since s is the least integer such that $D^sG \subset N$, we have s is the length of the derived sequence of $G/N = Gal(L/k)$. ∎

This theorem almost gives us a minimal denesting. Let l be the lcm of the exponents of the derived series of $G = Gal(L/K)$. Then if ζ_l, a primitive l^{th} root of unity is in k, by Theorems 1.5 and 1.8 all the extensions by the derived series can be achieved as radical extensions. If ζ_l is not in k, but is in L, we can still achieve a minimal denesting with all of the terms lying in L. This is because of:

Lemma 3.4 *Let L be the splitting field of the minimal polynomial of α over k, with Galois group G, and suppose ζ_l is in L. If $H = DG$ is the commutator subgroup of G, then $L^H \supseteq k(\zeta_l)$.*

Proof Let J be the subgroup of G associated with $k(\zeta_l)$. Now $k(\zeta_l)$ is an abelian extension of k, thus $J \lhd G$, and G/J is abelian . By Theorem 1.2, $J \supseteq DG$, and therefore $k(\zeta_l) = L^J \subseteq L^{DG}$. ∎

Of course, it will not always be the case that ζ_l lies in L. We can adjoin it to k. Thus the field extensions of k made by considering the fixed fields L^{D^iG} can all be made to be radical extensions. We will need to make use of:

Theorem 3.5 *(Lang [8], pp. 196-7.) Let K be a Galois extension of k, and F be an arbitrary extension of k. Then KF is Galois over F, and K is Galois over $K \cap F$. Let H be the Galois group of KF over F, and G the group of K over $K \cap F$. If σ is in H, then the restriction of σ to K is in G, and the mapping $\sigma \rightarrow \sigma\,|_K$ is an isomorphism.*

Lemma 3.6 *Suppose $k_1 = k(\alpha)$ is an abelian extension of k, with group G which has exponent l. Let ξ_l be a primitive l^{th} root of unity, and let $K = k(\zeta_l)$. Then $K(\alpha)$ is an abelian extension of K, and if $\tilde{G}$ is the associated Galois group, l is divisible by the exponent of $\tilde{G}$. If furthermore $k_1 = k(\alpha)$ is a cyclic extension of k, then $K(\alpha)$ is a radical extension of K, that is, $K(\alpha) = K(\beta)$, for some β which satisfies an irreducible polynomial of the form $x^n - b$ for some b in K.*

Proof Consider the following picture:

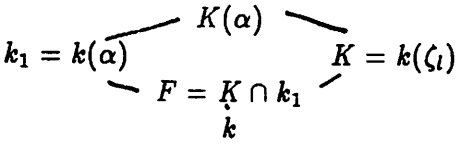

Now k_1 is an abelian extension over k; thus so is k_1 over F. Furthermore, by Lemma 3.5, $\tilde{G}$ is isomorphic to $Gal(k_1/F)$, which is just a subgroup of G. Thus the exponent of $\tilde{G}$ divides the exponent of G. If k_1 is a cyclic extension over k, so is k_1 over F. Thus so is $K(\alpha)$ over K. Again $\tilde{G}$ can be viewed as a subgroup of G. The exponent of $\tilde{G}$ will divide the exponent of G. But then, by Theorem 1.6, the extension $K(\alpha)$ over K can be realized as a radical extension. ∎

Theorem 3.7 *Suppose α is a nested radical over k, where k is a field of characteristic 0. Let L be the splitting field of $k(\alpha)$ over k, with Galois group G. Let l be the lcm of the exponents of the derived series of G. If there is a denesting of α such that each of the terms has depth no more than t, then there is a denesting of α over $k(\zeta_l)$ with each of the terms having depth no more than $t+1$, and lying in $L(\zeta_l)$.*

Proof Since α is a nested radical with nesting depth t, α can be expressed as a polynomial in $\beta_1, \ldots, \beta_r$, with each β_i having a nesting depth $t_i \leq t$ over k. Each β_i has a nesting sequence of the form:

$$\begin{aligned}
\beta_{i1} &= \sqrt[m_{i1}]{p_{i1}},\ p_{i1} \in k, \\
\beta_{i2} &= \sqrt[m_{i2}]{p_{i2}(\beta_{11}, \ldots, \beta_{r1})}, \\
\beta_{i3} &= \sqrt[m_{i3}]{p_{i3}(\beta_{11}, \ldots, \beta_{r1}, \beta_{12}, \ldots, \beta_{r2})}, \\
&\vdots \\
\beta_i = \beta_{it_i} &= \sqrt[m_{il_i}]{p_{il_i}(\beta_{11}, \ldots, \beta_{rt_{r-1}})}.
\end{aligned}$$

Of course, the symbol $\sqrt[m_{ij}]{p_{ij}(\beta_{11}, \ldots, \beta_{rj-1})}$ will mean the same thing each time it appears. Let $m = lcm(m_{11}, \ldots, m_{1l_1}, \ldots, m_{r1}, \ldots, m_{rl_r})$, and let $\hat{k} = k(\zeta_m)$.

Consider the sequence of fields $k_1 = \hat{k}(\beta_{11}, \ldots, \beta_{1r}) \subseteq k_2 = k_1(\beta_{21}, \ldots, \beta_{2r}) \subseteq \ldots \subseteq k_t = k_{t-1}(\beta_{1t}, \ldots, \beta_{rt})$. For each i, k_{i+1} is a composite of cyclic extensions of k_i. Thus for each i, k_{i+1} is an abelian extension of k_i. Let L be the splitting field of $k(\alpha)$ over k, with Galois group G. We know that L is contained in the normal closure of k_t over k. If the length of the derived series for G is s, then by Theorem 3.3 $s \leq t+1$.

Let $L_i = L^{D^iG}$. Then the sequence of fields $k = L_0 \subset L_1 \ldots \subset L_{s-1} \subset L_s = L$ is a tower of abelian extensions. The associated Galois group is $D^{i-1}G/D^iG$.

Each $D^{i-1}G/D^iG$ is an abelian group and can be written as a direct product of cyclic groups, say $D^{i-1}G/D^iG = J_{i1} \times \ldots \times J_{it_i}$. For each i and j, $i = 1, \ldots, s$, $j = 1, \ldots, t_i$, let $\tilde{J}_{ij} = \{e\} \times \ldots \{e\} \times J_{ij} \times \{e\} \ldots \{e\} \subseteq D^iG/D^{i-1}G$. Then $L_i = L^{\tilde{J}_{ij}}$ is a cyclic extension of L_{i-1}, and $L_i = L_i^{\tilde{J}_{i1}} L_i^{\tilde{J}_{i2}} \ldots L_i^{\tilde{J}_{it_i}}$ is a composite of cyclic extensions.

Now let $K_i = L_i(\zeta_l)$. By Lemma 3.6 the extension K_i over K_{i-1} is abelian, and the extension $K_{ij} = L_i^{\tilde{J}_{ij}}(\zeta_l)$ over K_i is a radical extension for every pair i, j. The maximal depth of nesting for any expression in K_s over k is s, the height of the derived series for G. The theorem is proved. ∎

Observe that the reason the depth of nesting achieved in Theorem 3.7 is one more than minimal is because the abelian tower created by the derived series begins with the field k rather than $k(\zeta_m)$. But ζ_m is, after all, just a root of unity. If we begin our tower of fields at $k(\zeta_m)$ instead, then we have:

Corollary 3.8 *Let k, α, L, G, l, t be as in Theorem 3.7. Let m be the lcm of (m_{ij}), where the m_{ij} runs over all the roots appearing in the given depth t nested expression for α. Let r be the lcm of (m, l). Then there is a minimal depth nesting of α over $k(\zeta_r)$ with each of its terms lying in $L(\zeta_r)$.*

Proof This proof is a variation on the one above. We want to show that there is a denesting of α over $k(\zeta_r)$ of minimal depth, with all of its terms lying in $L(\zeta_r)$. We do it by showing something slightly stronger, that there is a denesting achieving this with all of its terms lying in $\hat{L}$, the splitting field of the minimal polynomial of α over $k(\zeta_r)$.

We begin with some notation. Let $\hat{k} = k(\zeta_r)$, and let $\hat{G}$ be the Galois group of the splitting field of the minimal polynomial of α over $\hat{k}$. Note that by Theorem 3.5, $L(\zeta_r)$ over $\hat{k}$ is Galois. If $\tilde{G}$ is the group of that extension, then $\tilde{G}$ is isomorphic to a subgroup of G, the subgroup corresonding to L over $k(\alpha) \cap k(\zeta_r)$. (Note that $k(\alpha) \cap k(\zeta_r)$ is an abelian extension of k.) We have the following picture:

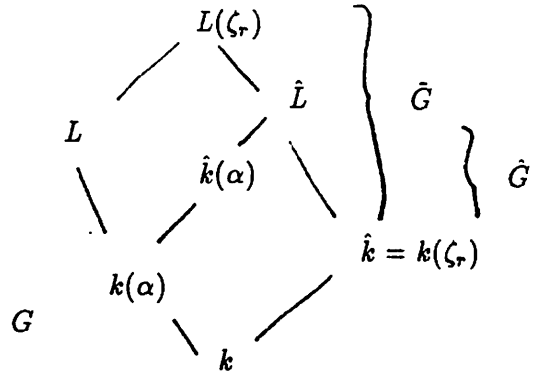

Consider the tower $\hat{k} \subset \ldots \hat{k}_t$, and let us compare it to a tower of abelian extensions in $\hat{L}$ over $\hat{k}$. Theorem 3.3 applies. If s be the length of the derived series for $\hat{G}$, we have $s \leq t$.

We let $\hat{L}_i = \hat{L}^{D^i\hat{G}}$. As in the previous theorem, the sequence of fields $\hat{k} \subset \hat{L}_1 \subset \ldots \subset \hat{L}_s$ is an abelian tower. Again, this can be transformed so that for each i, $\hat{L}_i$ is a composite of cyclic extensions of $\hat{L}_{i-1}$.

Since $\hat{G}$ is a quotient group of $\tilde{G}$, the lcm of the exponents of the derived series for $\hat{G}$ divides the lcm of the exponents for the derived series of $\tilde{G}$, which in turn divides the lcm of the exponents of the derived series for G. This divides r. Thus $\hat{k}$ contains the roots of unity needed to make the field extensions of $\hat{L}$ over $\hat{k}$ corresponding to the derived series into composites of radical extensions. The corollary follows.

∎

Although Corollary 3.8 appears to give a better bound than Theorem 3.7, it does so in an dissatisfying way, by introducing roots of unity which have to do with the input and are not necessarily a genuine part of the problem. For example, in denesting $\sqrt[3]{\sqrt{5}+2} - \sqrt[3]{\sqrt{5}-2}$ over Q, Corollary 3.8 indicates that one would want to add a primitive sixth root of unity to Q to denest. In fact $\sqrt[3]{\sqrt{5}+2} - \sqrt[3]{\sqrt{5}-2} = 1$, and thus no new root of unity is needed (a fact which would be discovered in computing the minimal polynomial for $\sqrt[3]{\sqrt{5}+2} - \sqrt[3]{\sqrt{5}-2}$ over Q). For this reason we have chosen to use the Theorem 3.7 variety of denesting as the basis for our algorithms. The algorithms can be easily modified to handle the Corollary 3.8 version if desired.

Now we are in a position to denest general radicals. Theorem 3.7 tells us which root of unity we must add. Before we procede with the algorithm, we briefly discuss computing the minimal polynomial of α; details may be found in the appendix.

From an algebraic point of view, a sensible measure of the size of the denesting problem is the size of the minimal polynomial of α. There is some disagreement about this issue. Borodin et al [3] define the size of α to be m, the depth of nesting of α. Under this measure, they gave an algorithm for denesting a nested class of square roots. This algorithm has polynomial arithmetical complexity, but exponential bit complexity. The latter comes from the doubling of bits which occurs at each extension. We believe that this problem is inherent in any denesting scheme, and that the measure defined by [3] is too conservative. Under our measure, their running time is polynomial.

We can view a nested radical α as a sequence of polynomials $\tilde{p}_i, \tilde{q}$ in $k[x_1, \ldots, x_{m-1}]$ such that

$$\begin{aligned}
\alpha_1 &= \sqrt[n_1]{\tilde{p}_1},\ \tilde{p}_1 \in k \\
\alpha_2 &= \sqrt[n_2]{\tilde{p}_2(\alpha_1)} \\
\alpha_3 &= \sqrt[n_3]{\tilde{p}_3(\alpha_1, \alpha_2)} \\
&\vdots \\
\alpha_m &= \tilde{q}(\alpha_1, \ldots, \alpha_{m-1}) + \sqrt[n_m]{\tilde{p}_m(\alpha_1, \ldots, \alpha_{m-1})},
\end{aligned}$$

with $\alpha = \alpha_m$. In the appendix we show how to go from this description to minimal polynomial of $\alpha = \alpha_m$ over k.

4 The Algorithm and Running Time Analysis

At this point we have presented all the new ideas which were needed in order to compute denestings. In this section we show that we can effectively compute the denesting, and we give a running time analysis. There are some computations we have not yet shown how to do, the most important of which is how to go from an arbitrary description of a radical extension (eg $K = k[x]/f(x)$, $f(x)$ irreducible) to a description as a radical extension ($K = k[x]/(x^n - b)$, with $x^n - b$ irreducible). The answer is Lagrange resolvents, and we use these and Artin's normal basis to achieve what we want.

We begin this section with a brief – but complete – description of the algorithm. The rest of the section is concerned with the details of running time analysis and coefficient blow-up. Because it is remarkably easy to miss the forest for the trees, we urge the first-time reader to read the brief – but complete – description of the algorithm, and then to read about normal bases and Lagrange resolvents. A second reading can cover the running time analyses, and the other algorithms presented. We end the section with a more precise version of the algorithm, and, of course, a running time analysis.

Suppose we wish to denest the nested radical α. We begin by computing the minimal polynomial of α over k. We construct the splitting field L of the minimal polynomial of α over k. We compute $G = Gal(L/k)$, and the series of commutator subgroups $D^iG, i = 1, \ldots, s$, where $D^sG = \{e\}$. We also compute l, the lcm of the exponents of the derived series of G.

Next, for each $i, i = 1, \ldots, s$, we compute $D^{i-1}G/D^iG = J_{i1} \times \ldots \times J_{it_i}$ as a direct product of cyclic groups. Let $\tilde{J}_{ij} = \{e\} \times \ldots \times \{e\} \times J_{ij} \times \{e\} \times \ldots \times \{e\}$, and let $L_i = L^{D^iG}$. Thus for each i, $L_i = L_i^{\tilde{J}_{i1}} \ldots L_i^{\tilde{J}_{it_i}}$ is a composite of cyclic extensions of L_{i-1}. For each i and j, we compute $\tilde{\beta}_{ij}$ such that $L_i^{\tilde{J}_{ij}} = L_{i-1}(\tilde{\beta}_{ij})$. Thus $L_i = L_{i-1}(\tilde{\beta}_{i1}, \ldots, \tilde{\beta}_{it_i})$.

We write $K_0 = k(\zeta_l)$, where ζ_l is a primitive l^{th} root of unity. By Lemma 3.6, $K_{ij} = K_{i-1}(\tilde{\beta}_{ij})$ can be written as a radical extension of K_{i-1}, and each $K_i = K_{i1} \ldots K_{it_i}$ is a composite of radical extensions of K_{i-1}. We achieve the radical extensions as follows.

Following Artin, we construct a polynomial $s_{ij}(x)$ whose roots $\theta_{ij1}, \ldots, \theta_{ijr_{ij}}$ form a "normal" basis for K_{ij} over K_{i-1}. The degree of $s_{ij}(x)$ is $r_{ij} = [K_{ij} : K_{i-1}]$, and its roots are linearly independent over K_{i-1}. Then we will use Lagrange resolvents to find a β_{ij} in K_{ij} such that $K_{ij} = K_{i-1}(\beta_{ij})$, where β_{ij} satisfies an irreducible polynomial of the form $x^{n_{ij}} - b_{ij}$ over K_{i-1}. In our construction we will actually have two descriptions of the fields K_i, namely $K_{i-1}(\beta_{i1}, \ldots, \beta_{it_i})$ and $K_{i-1}[x]/g_i(x)$, where $g_i(x)$ is irreducible. We will do computations using the latter representation.

In this section we show that all the computations described above can be carried out in time polynomial in the size of the splitting field of the minimal polynomial of α over k. Unfortunately, the splitting field can be large. In worst case, the splitting field is of exponential degree over the base field. Our bounds are chosen for the relative simplicity of presentation, and are not necessarily the tightest possible. We will restrict our running time analysis to $k = Q$ for the present.

In [6], it was observed that the splitting field and Galois group G can be computed in time polynomial in $r = |G| = [L : k]$ and $\log |f(x)|$. More precisely:

Theorem 4.1 *Let $f(x)$, an irreducible polynomial of degree m over k, be the minimal polynomial of α, a nested radical over k. The Galois group of $f(x)$ over k can be computed in $O(r^{12+\epsilon} \log^{4+\epsilon}(r|f(y)|))$ steps.*

Proof Suppose $f(x)$ has roots $\alpha_1, \ldots, \alpha_m$. The following algorithm computes L, the splitting field of $f(x)$ over k, and G, the associated Galois group.

Step 1: $g(x) \leftarrow f(x)$;
Step 2: WHILE $f(y)$ does not split completely in $k[x,y]/g(x)$ DO:
BEGIN
Factor $f(y) = \Pi f_i(y)$ in $k[x,y]/g(x)$;
Let $h(y)$ be some $f_i(y)$ of degree > 1;
Pick c such $g(x) \leftarrow Res_t(g(t), h(x - ct))$ is square free;
END;
Step 3: $\rho \leftarrow$ root of $g(x)$;
Factor $g(y) = \Pi(y - p_i(\rho))$ in $k[x,y]/g(x)$;
$\sigma_i(\rho) \leftarrow p_i(\rho)$;
FOR each σ_i in G, compute σ_i acting on ρ_j by
$\sigma_i(\rho_j) = \sigma_i(p_j(\rho)) = p_j(\sigma_i(\rho)) = p_j(p_i(\rho))$.

By Theorem A.4, it is not hard to see that this algorithm computes a polynomial $g(x)$ such that $L = k[x]/g(x)$. The roots of this polynomial are the conjugates of

$$\gamma_1 = (c_1 \cdots c_{m-1})^2 \alpha_1 + (c_1 \cdots c_{m-2})^2 \alpha_2 + \ldots + \alpha_m,$$

where

$$c_i < d_i = [k(\alpha_1, \ldots, \alpha_i) : k(\alpha_1, \ldots, \alpha_{i-1})].$$

How long the process takes depends on the structure of G. We will assume a worst case scenario. Note that not all worst case assumptions can occur simultaneously (and thus what we have is really an exaggerated worst case scenario).

We loop through Step 2 at most $m \leq \log r$ times. The maximal possible $g(x)$ occurs in the last iteration of the loop, and it is bounded by the size of

its roots. Now $g(x)$ is of degree r over k in the final iteration. Then

$$|g(x)| \leq r\binom{r}{r/2}||\gamma_1||^r < 2^r(r^2||\alpha_1||)^r < (r||\alpha_1||)^{2r} < (r|f(y)|)^{2r+\epsilon},$$

since $||\alpha_1|| < 1+|f(y)|$. Using Theorem A.4, and the fact that degree$(g(x))\times$ degree$(f(x)) = r$, we have factoring $f(y)$ over $k[x]/g(x)$ takes no more than

$$O((r/m)^{9+\epsilon}m^{7+\epsilon}\log^2(g(x))\log^{2+\epsilon}(||f(y)||(r|g(x)|)^r(rm)^m))$$

steps. We loop through at most m times, getting

$$\begin{aligned} &O((r/m)^{9+\epsilon}m^{7+\epsilon}\log^2(r|f(y)||)^{2r+\epsilon}\log^{2+\epsilon}(|f(y)|(r(r|f(y)||)^{2r+\epsilon})^r(rm)^r)m) \\ &\quad < O(r^{9+\epsilon}\log^2(r|f(y|)^{2r+\epsilon}\log^{2+\epsilon}(r|f(y)||)^{4r^2})) \\ &\quad < O(r^{9+\epsilon}\log^2(r|f(y|)^{2r+\epsilon}\log^{2+\epsilon}(r|f(y)||)^{4r^2})) \\ &\quad < O(r^{12+\epsilon}\log^{4+\epsilon}(r|f(y)||)) \end{aligned}$$

steps.

∎

Thus in time polynomial in the degree of the splitting field, and the size of the minimal polynomial of α over k, we can compute the Galois group. It has long been known how to compute commutator subgroups quickly from a group table; recently Babai, Luks and Seress [2] showed how to do so in $O(r^4\log^c r)$ steps, for a permutation group on r elements.

We say a polynomial $g(x)$ over a field k is normal iff it is irreducible and $k[x]/g(x)$ is a Galois extension of k (equivalently $g(y)$ splits completely in $k[x]/g(x)$). It is not hard to go from such a polynomial to one which splits completely and whose roots are linearly independent over k. Artin gave an effective method for calculating such a polynomial. We reproduce it here.

Theorem 4.2 *(Artin [1]) If L is a Galois extension of k, and $\sigma_1, \ldots, \sigma_r$ are the elements of the group G, then there is an element θ in L such that the r elements $\theta_1 = \sigma_1(\theta), \ldots, \theta_r = \sigma_r(\theta)$ are linearly independent with respect to k.*

Proof Let $L = k(\rho)$, and let $g(x)$ be the minimal polynomial for ρ over k. Let $\sigma_i(\rho) = \rho_i$. We let $h(x) = g(x)/(x-\rho)g'(\rho)$, and $h_i(x) = \sigma_i(h(x)) = g(x)/(x-\rho_i)g'(\rho_i)$. Then $h_j(x)$ is a polynomial over L having ρ_i as a root for $j \neq i$, and thus $h_i(x)h_j(x) = 0 \pmod{g(x)}$ for $i \neq j$. In the equation

$$h_1(x) + h_2(x) + \ldots + h_r(x) - 1 = 0, \tag{1}$$

the left side is of degree at most $r-1$. If the equation were true for r different values of x, the left side would be identically 0. But we have r such values $\rho_1, \ldots, \rho_r$, since $h_i(\rho_i) = 1$, and $h_j(\rho_i) = 0$ for $j \neq i$. Multiplying this equation by $h_i(x)$, and using the fact that $h_i(x)h_j(x) = 0 \pmod{g(x)}$ for $i \neq j$, we find that

$$(h_i(x))^2 = h_i(x) \pmod{g(x)}. \tag{2}$$

We next consider the determinant

$$D(x) = \mid \sigma_i\sigma_j(h(x)) \mid, \tag{3}$$

and prove $D(x)$ is not identically 0. If we square the corresponding matrix we get the identity matrix $\pmod{g(x)}$, because of equations 1-3.

Now $D(x)$ can have at most r^2 roots in k. If we avoid them, we can find a value a for x such that $D(a) \neq 0$. Let $\theta = h(a)$. Then

$$\mid \sigma_i\sigma_j(\theta) \mid \neq 0.$$

Consider any relation of the form $a_1\sigma_1(\theta) + \ldots + a_r\sigma_r(\theta) = 0$, where the a_i are in k. Applying the automorphisms σ_j to it would lead to r homogeneous equations for the r unknowns a_i. Then $a_i = 0$, and we are done. ∎

Algorithm: Compute Normal Basis

input: $g(x)$, a normal polynomial over k with root ρ
output: θ, an element in $k[x]/g(x)$ whose conjugates form a normal basis for $k[x]/g(x)$ over k.

Line 1: BEGIN
Line 2: $h(x) \leftarrow g(x)/(x-\rho)g'(\rho)$;
Line 3: FOR $i = 1, \ldots, r$ DO:
FOR $j = 1, \ldots, r$ $a_{ij} \leftarrow \sigma_i\sigma_j(h(x))$;
Line 4: $D(x) \leftarrow \mid a_{ij} \mid$;
Line 5: Pick a in k such that $D(a) \neq 0$;
Line 6: $\theta \leftarrow h(a)$;
Line 7: END.

Theorem 4.3 Algorithm [Compute Normal Basis] *computes a normal polynomial whose roots form a basis for $k[x]/g(x)$ over k. It does so in $O(r^{7+\epsilon}\log^{2+\epsilon}|f(y)|)$ steps.*

Proof That it does so correctly is clear from the proof of Artin's Theorem. We already know by Theorem 4.1 that

$$|g(x)| < (r|f(y)|)^{2r+\epsilon}$$

Thus

$$|g'(x)| < (r|f(y)|)^{2r+\epsilon'}$$

and

$$|h(x)| < (r|f(y)|)^{4r+\epsilon}$$

We know the entries in the matrix are bounded by $|h(x)|$, and that $D(x)$ is a polynomial of degree at most r^2. By Hadamard's inequality [5], the coefficients of $D(x)$ are bounded by $(r|f(y)|)^{9r^2}$. There is an a less than r^2 for which $D(a)$ is non-zero. Now

$$\theta = h(a) < r \cdot (r^2)^r (r|f(y)|)^{4r+\epsilon} < (r|f(y)|)^{8r}$$

Lines 3 and 5 dominate the computations. Line 3 can be computed in $O(r^2 \cdot r^3 \log^{2+\epsilon} |h(x)|) = O(r^{7+\epsilon} \log^{2+\epsilon}(r|f(y)|))$ steps, as can Line 5.

Observe that if $s(x)$ is the minimal polynomial for θ over k, then

$$|s(x)| < r\binom{r}{r/2} ||\alpha||^r < r2^r (r|f(y)|)^{8r} < (r|f(y)|)^{8r+\epsilon}.$$

∎

We are now ready to compute fixed fields. Without loss of generality, suppose $\{\sigma_1, \ldots, \sigma_l\}$ are the elements of H. Let $\Sigma a_j \theta_j$, with a_j in k, be an arbitrary element of L^H. For each σ_i in H, we have

$$\Sigma a_j \theta_j = \sigma_i(\Sigma a_j \theta_j) = \Sigma a_j \sigma_i(\theta_j).$$

Since the Galois group sends roots of $s(x)$ to roots of $s(x)$, we have $\sigma_i(\theta_j) = \theta_l$ for some l. Because the θ_j are linearly independent over k, the only way the above equation can be satisfied is if $a_i = a_l$. Each element σ_i in H gives rise to equalities amongst the a_i, which leads to a series of relationships amongst the θ_i. That is, if $a_i = a_l$, then in L^H any expression containing θ_i must have θ_l appearing with the same coefficient. Computing all such relationships gives us exactly the fixed field associated with H. We make this precise in the following algorithm.

Algorithm: Compute Fixed Field

input: $\theta = h(a)$, an element in $k[x]/g(x)$ whose conjugates form a normal basis for $k[x]/g(x)$ over k, and $\{\sigma_1, \ldots, \sigma_l\} = H$, a subgroup of Galois group $G = \{\sigma_1, \ldots, \sigma_s\}$ of $k[x]/g(x)$ over k.
output: $\{\gamma_1, \ldots, \gamma_r\}$, where $(k[x]/s(x))^H = k(\gamma_1, \ldots, \gamma_r)$.

Line 1: BEGIN
Line 2:

$$a_{ij} = \begin{cases} 1 & \text{if } \sigma_i(h(a)) = \sigma_j(h(a)) \\ 0 & \text{otherwise} \end{cases}$$

Line 3: $T \leftarrow$ transitive closure of (a_{ij});
Line 4: Let $C_1, \ldots, C_n$ be the connected components of the graph given by A

$$\gamma_i = \Sigma_{j \in C_i} \sigma_j(h(a))$$

Line 5: END.

Theorem 4.4 Algorithm [Compute Fixed Field] *computes the fixed field of a given subgroup H. It does so in $O(r^5 \log^2(r|f(y)|))$ steps.*

Proof If $K = L^H$, then $K = k(\gamma_1, \ldots, \gamma_r)$, where the γ_j are sums of the θ_i. (Note that the γ_j partition the θ_i.) In Algorithm Fixed Fields, we set up an $n \times n$ matrix (a_{ij}), with 1's on the diagonal, 0's elsewhere. If σ_i in H takes θ_j to θ_k, then a_{jk} and a_{kj} are both assigned a 1. This is shorthand for saying θ_j and θ_k appear with the same coefficient in K.

In Line 2 we set up the matrix. We compute the transitive closure in Line 3. Then in Line 3 we set up the γ_i, $i = 1, \ldots, n$, to be the sum of the θ_i which all have the same coefficient. We do this by computing the transitive closure of A. Clearly $K = k(\gamma_1, \ldots, \gamma_r)$.

The computation is dominated by the calculations of Line 2 which takes $O(r^2 \cdot r \log^2 |s(x)|) = O(r^5 \log^2(r|f(y)|))$ steps since $|s(x)| < (r|f(y)|)^{8r+\epsilon}$. ∎

At this point we observe that it is easy to calculate the minimal polynomial for a primitive l^{th} root of unity. If $k = Q$, it will be the irreducible factor of $x^l - 1$ of degree $\varphi(l)$. If $l(t)$ is the minimal polynomial for ζ_l over k, then $|l(t)| < 2^t$. If $k \neq Q$, the minimal polynomial is a little more complicated to find, and there may be more than one choice for it (that is, there are non-conjugate primitive l^{th} roots of unity in that case), but the bound on degree and coefficient size remains the same.

All that remains is to show how to transform a cyclic extension of degree l over k, where ζ_l, a primitive l^{th} root of unity, is in k, into a radical extension. The technique was developed by Lagrange, two centuries ago.

Definition: Let $L = k(\alpha)$ be a cyclic extension of degree l, and suppose ζ_l is in k. Let σ be a generator of $G = Gal(L/k)$. For any element τ in L, define

$$(\zeta_l, \tau) = \tau + \zeta_l \sigma\tau + \ldots + \zeta_l^{l-1}\sigma^{l-1}\tau.$$

The element (ζ_l, τ) is called a *Lagrange resolvent.*

It is not hard to show that the σ's are linearly independent over k, thus there is a τ which makes the Lagrange resolvent non-zero. Such an element will generate L over k, and in particular will satisfy an irreducible polynomial over k of the form $x^l - b$. Observe that if $\theta_1, \ldots, \theta_l$ is a normal basis for L, then (ζ_l, θ) is non-zero, and this will suffice.

Algorithm: Compute Denesting

input: $f(x)$, an irreducible polynomial of degree n over k, and n, the lcm of the indices of the reducible radical expressions for α.
output: A sequence of fields K_i, $i = 1, \ldots, s$, and an expression for α, where

1. the expression for α, a root of $f(x)$, in K_s that is within one of minimal nesting depth over $k(\zeta_l)$,

2. and l =lcm (n, exponent G/DG, exponent $DG/D^2G, \ldots$, exponent $D^{s-1}G/D^sG$), and ζ_l is a primitive l^{th} root of unity.

Line 1: BEGIN
Line 2: Compute L, the splitting field of $f(x)$ over k;
Line 3: Compute G = Galois group of L over k;
Line 4: $D^0 \leftarrow G$;
Line 5: $i \leftarrow 0$
Line 6: WHILE $D^iG \neq \{e\}$ DO:
Line 7: BEGIN
Line 8: $i \leftarrow i + 1$;
Line 9: $D^i \leftarrow \langle \sigma\tau\sigma^{-1}\tau^{-1} \mid \sigma, \tau \in D^iG \rangle$
Line 10: END;
Line 11: $l \leftarrow$ lcm (n, exponent G/DG,exponent $DG/D^2G, \ldots,$ exponent $D^{s-1}G/D^sG$);
Line 12: Find $l(x)$, minimal polynomial for ζ_l over k;
Line 13: $K_0 \leftarrow k(\zeta_l)$;
Line 14: FOR $i = 1$ TO s DO:
Line 15: BEGIN
Line 16: $L_i \leftarrow L^{D^iG}$;
Line 17: $K_i \leftarrow L_i(\zeta_l)$;
Line 18: Write $D^iG/D^{i-1}G$ as a product of cyclic groups, $J_{i1} \times \ldots \times J_{it_i}$;
Line 19: FOR $j = 1$ TO t_i DO:
Line 20: BEGIN
Line 21: $\tilde{J}_{ij} \leftarrow \{e\} \times \ldots \{e\} \times J_{ij} \times \{e\} \times \ldots \{e\} \subset D^iG/D^{i-1}G$;
Line 22: $L_{ij} \leftarrow L_i^{\tilde{J}_{ij}}$;
Line 23: $K_{ij} \leftarrow L_{ij}(\zeta_l)$;
Line 24: $l_{ij} \leftarrow [K_{ij} : K_{i-1}]$;

Line 25: Compute a normal basis for K_{ij} over K_{i-1}, $\theta_{ij1}, \ldots, \theta_{ijl_{ij}}$;
Line 26: Pick $\sigma_{(ij)}$, a generator of $\tilde{J}_{ij}$;
Line 27: $\beta_{ij} \leftarrow (\zeta_{l_{ij}}, \theta_{ij1})$;
Line 28: $K_{ij} \leftarrow K_{i-1}(\beta_{ij})$;
Line 29: END;
Line 30: $K_i \leftarrow K_{i-1}(\beta_{i1}, \ldots, \beta_{it_i})$;
Line 31: END;
Line 32: Express α in K_s;
Line 33: END.

Theorem 4.5 *Let $f(x)$ of degree m be the minimal polynomial for α, a nested radical over k, and let L be the splitting field $f(x)$ over k, with Galois group G. Let D^iG represent the i^{th} commutator subgroup of G, with $D^sG = \{e\}$, and let $l = lcm(n, exponent\ G/DG, \ldots, exponent\ D^{s-1}G/D^sG)$. On input $f(x)$ the algorithm* **Compute Denesting** *computes a denesting of α that is within depth one of optimal over $k(\zeta_l)$, where ζ_l is a primitive l^{th} root of unity. It does so in $O(r^{12+\epsilon}\log^{4+\epsilon}(r|f(y)|))$ steps, where r is the order of the Galois group G.*

Proof As always, we prove correctness, then analyze running time. We want a minimal depth denesting for α. By Theorem 3.7 there is a denesting of α over $k(\zeta_l)$ that is within depth one of minimal and which has each of its terms lying in $L(\zeta_l)$, where L is the splitting field of the minimal polynomial of α over k.

Thus what we must do is calculate L, the splitting field of $f(x)$ over k, G, its Galois group, the derived series, and l, the lcm of the indices of the reducible radicals in the original exression for α, and the exponents of the derived series. Then if $L_1 \subset L_2 \subset \ldots \subset L_s \supset L(\alpha)$ is a series of abelian extensions, so is $K_1 = L_1(\zeta_l) \subset L_2 = k_2(\zeta_l) \subset \ldots \subset K_s = L_s(\zeta_l)$ by Lemma 3.6. More to the point, by Lemma 3.6 K_i is a composite of radical extensions of K_{i-1} for each i. In particular, α will have an expression that is within one of minimal nesting depth in K_s.

In Line 2 we calculate the Galois group. It is straightforward to compute G, its derived series, the exponent l. This takes us to Line 12. Line 13 simply begins the process that will create the fields $K_i = L_i(\zeta_l)$, by setting $K_0 = k(\zeta_l)$.

In Line 16 we compute the fields $L_i = L^{D^iG}$ which correspond to the subgroups of the derived series. In Line 16 we create the corresponding tower of fields $K_i = L_i(\zeta_l)$. What we next need to do is compute the tower of fields K_i, and express them as radical extensions.

We begin by computing the groups $D^iG/D^{i-1}G$ as a product of cyclic groups $J_{i1} \times \ldots \times J_{it_i}$. The we cycle through and compute first the groups $\tilde{J}_{ij} = \{e\} \times \{e\} \times J_{ij} \times \{e\} \times \ldots \{e\}$, and then the fixed field associated with each such $\tilde{J}_{ij}, L_{ij}$. By Lemma 3.6, the field calculated in Line 23, $K_{ij} = L_{ij}(\zeta_l)$ is a radical extension of L_{i-1}. In Line 25 we compute the normal basis $\theta_{i1}, \ldots, \theta_{il_j}$ for K_{ij} over K_{j-1}, and in Line 26 we pick an element from the associated Galois groups $\tilde{J}_{ij}$. The element $\beta_{ij} = \theta_{ij_1} + \zeta_{l_{ij}}\sigma_{ij}^{l_{ij}}\theta_{ij_1} + \ldots + \zeta_{l_{ij}}^{l_{ij}-1}\sigma_{ij}\theta_{i1}$ for some ordering of the θ's. Since the θ's are linearly independent, the element is perforce non-zero, hence by the theory of Lagrange resolvents, generates K_{ij} over K_{j-1}. In particular, β_{ij} satisfies an irreducible cyclic polynomial over K_{j-1}.

How long does the computation take? We claim no more than $O(r^{12+\epsilon}\log^{4+\epsilon}(r|f(y)|))$ steps, where $r = [L : k]$. Certainly the computations in Lines 1-11 are dominated by the time it takes to compute the Galois group; by Theorem 4.1 that can done in $O(r^{12+\epsilon}\log^{2+\epsilon}(r^2|f(y)|))$ steps. Lines 12 and 13 are a simple computation (even when $k \neq Q$). We run through the loop in loop in Lines 14-29 at most $O(\log r)$ times. Its running time is dominated by Line 25, which takes $O(r^{8+\epsilon}\log^{2+\epsilon}|f(y)|)$ steps. The computation in Line 32 takes no time at all.

To simplify the computations, we view the fields K_i as simple extensions of K, and do the computation in terms of a primitive element, with minimal polynomial $g_{ij}(t)$ over k. However, since what we want is a basis for L with a minimal depth of nesting, we will store two different bases for K_{ij}; one with the β_{ij} which come from the computations, and then $K_{ij} = k[t]/g_{ij}(t)$ written as an extension by a primitive element. We do not run into coefficient blowup in computing $g_{ij}(t)$ because each of the fields K_{ij} can be viewed as a subfield of $L(\zeta_l)$. (In the interests of making the algorithm as clear as possible, we have not included those fine points in the algorithm.) ∎

We have presented the running time analysis for the algorithm over Q. The only reason for that is simplicity. The only requirement for the base field is that one needs to be able to factor polynomials over k. Theorem A.4 shows how a factorization over k can be raised to a factorization over $k(\alpha)$. We can generalize **Algorithm Compute Polynomial** to compute a minimal polynomial for α over k. Of course, we can generalize Theorem 4.1,

and **Algorithms Compute Fixed Field** and **Compute Denesting**.

Note also that the running time for the algorithm **Compute Denesting** can be simply stated as the time required to compute the splitting field of $f(x)$ over Q, which is in turn the time needed to factor $f(x)$ over its splitting field.

5 Conclusions

There are a number of open questions which remain.

- Can the bound in Theorem 3.7 be made optimal, that is, can we improve it to $s \leq t$ without involving the form of the input (as per Corollary 3.8)? What is the tradeoff between this and the roots of unity?

- Suppose α is a nested radical over a field k, and suppose there is a denesting expression for α involving roots of unity. When is there a way to transform that to an expression of the same depth which avoids using the roots of unity? (The expression $\sqrt{\sqrt{5} - 5/2} = \zeta_5 - 1/\zeta_5$ makes it clear it that will not always be possible to do so.)

Our algorithm is not fast. Thus the other important issue in this problem is speed. To this we have several comments:

- Is there a way to perform these computations via a straightline program? The encoding in the straightline version would avoid the problem of coefficient blow-up. The main difficulty seems to be in computing the Galois group. We do not see how to determine the group without determining both $f(x)$, the minimal polynomial of α over k, and the splitting field of $f(x)$ over k. If a way could be found around these problems, then the entire algorithm could be sped up. We do not think this will be easy. In particular, it is likely that any technique which works here will also work to determine the Galois group in the solvable case. No polynomial time algorithms are presently known, and this would be a surprising and exciting breakthrough.

- It is important to remember that the bounds given here are really *upper* bounds. The running time is exponential *iff* the splitting field is of exponential degree over k. If the degree of the splitting field is polynomially bounded, then so is the running time of the algorithm.

Algorithm: Compute Denesting

input: $f(x)$, an irreducible polynomial of degree n over k, and n, the lcm of the indices of the reducible radical expressions for α.
output: A sequence of fields K_i, $i = 1, \ldots, s$, and an expression for α, where

1. the expression for α, a root of $f(x)$, in K_s that is within one of minimal nesting depth over $k(\zeta_l)$,

2. and l =lcm (n, exponent G/DG, exponent $DG/D^2G, \ldots$, exponent $D^{s-1}G/D^sG$), and ζ_l is a primitive l^{th} root of unity.

Line 1: BEGIN
Line 2: Compute L, the splitting field of $f(x)$ over k;
Line 3: Compute G = Galois group of L over k;
Line 4: $D^0 \leftarrow G$;
Line 5: $i \leftarrow 0$
Line 6: WHILE $D^iG \neq \{e\}$ DO:
Line 7: BEGIN
Line 8: $i \leftarrow i + 1$;
Line 9: $D^i \leftarrow \langle \sigma\tau\sigma^{-1}\tau^{-1} \mid \sigma, \tau \in D^iG \rangle$
Line 10: END;
Line 11: $l \leftarrow$ lcm (n, exponent G/DG,exponent $DG/D^2G, \ldots,$ exponent $D^{s-1}G/D^sG$);
Line 12: Find $l(x)$, minimal polynomial for ζ_l over k;
Line 13: $K_0 \leftarrow k(\zeta_l)$;
Line 14: FOR $i = 1$ TO s DO:
Line 15: BEGIN
Line 16: $L_i \leftarrow L^{D^iG}$;
Line 17: $K_i \leftarrow L_i(\zeta_l)$;
Line 18: Write $D^iG/D^{i-1}G$ as a product of cyclic groups, $J_{i1} \times \ldots \times J_{it_i}$;
Line 19: FOR $j = 1$ TO t_i DO:
Line 20: BEGIN
Line 21: $\tilde{J}_{ij} \leftarrow \{e\} \times \ldots \{e\} \times J_{ij} \times \{e\} \times \ldots \{e\} \subset D^iG/D^{i-1}G$;
Line 22: $L_{ij} \leftarrow L_i^{\tilde{J}_{ij}}$;
Line 23: $K_{ij} \leftarrow L_{ij}(\zeta_l)$;
Line 24: $l_{ij} \leftarrow [K_{ij} : K_{i-1}]$;

Line 25: Compute a normal basis for K_{ij} over K_{i-1}, $\theta_{ij1}, \ldots, \theta_{ijl_{ij}}$;
Line 26: Pick $\sigma_{(ij)}$, a generator of $\tilde{J}_{ij}$;
Line 27: $\beta_{ij} \leftarrow (\zeta_{l_{ij}}, \theta_{ij1})$;
Line 28: $K_{ij} \leftarrow K_{i-1}(\beta_{ij})$;
Line 29: END;
Line 30: $K_i \leftarrow K_{i-1}(\beta_{i1}, \ldots, \beta_{it_i})$;
Line 31: END;
Line 32: Express α in K_s;
Line 33: END.

Theorem 4.5 *Let $f(x)$ of degree m be the minimal polynomial for α, a nested radical over k, and let L be the splitting field $f(x)$ over k, with Galois group G. Let D^iG represent the i^{th} commutator subgroup of G, with $D^sG = \{e\}$, and let $l = lcm(n, \text{exponent } G/DG, \ldots, \text{exponent } D^{s-1}G/D^sG)$. On input $f(x)$ the algorithm* **Compute Denesting** *computes a denesting of α that is within depth one of optimal over $k(\zeta_l)$, where ζ_l is a primitive l^{th} root of unity. It does so in $O(r^{12+\epsilon}\log^{4+\epsilon}(r|f(y)|))$ steps, where r is the order of the Galois group G.*

Proof As always, we prove correctness, then analyze running time. We want a minimal depth denesting for α. By Theorem 3.7 there is a denesting of α over $k(\zeta_l)$ that is within depth one of minimal and which has each of its terms lying in $L(\zeta_l)$, where L is the splitting field of the minimal polynomial of α over k.

Thus what we must do is calculate L, the splitting field of $f(x)$ over k, G, its Galois group, the derived series, and l, the lcm of the indices of the reducible radicals in the original exression for α, and the exponents of the derived series. Then if $L_1 \subset L_2 \subset \ldots \subset L_s \supset L(\alpha)$ is a series of abelian extensions, so is $K_1 = L_1(\zeta_l) \subset L_2 = k_2(\zeta_l) \subset \ldots \subset K_s = L_s(\zeta_l)$ by Lemma 3.6. More to the point, by Lemma 3.6 K_i is a composite of radical extensions of K_{i-1} for each i. In particular, α will have an expression that is within one of minimal nesting depth in K_s.

In Line 2 we calculate the Galois group. It is straightforward to compute G, its derived series, the exponent l. This takes us to Line 12. Line 13 simply begins the process that will create the fields $K_i = L_i(\zeta_l)$, by setting $K_0 = k(\zeta_l)$.

In Line 16 we compute the fields $L_i = L^{D^iG}$ which correspond to the subgroups of the derived series. In Line 16 we create the corresponding tower of fields $K_i = L_i(\zeta_l)$. What we next need to do is compute the tower of fields K_i, and express them as radical extensions.

We begin by computing the groups $D^iG/D^{i-1}G$ as a product of cyclic groups $J_{i1} \times \ldots \times J_{it_i}$. The we cycle through and compute first the groups $\tilde{J}_{ij} = \{e\} \times \{e\} \times J_{ij} \times \{e\} \times \ldots \{e\}$, and then the fixed field associated with each such $\tilde{J}_{ij}, L_{ij}$. By Lemma 3.6, the field calculated in Line 23, $K_{ij} = L_{ij}(\zeta_l)$ is a radical extension of L_{i-1}. In Line 25 we compute the normal basis $\theta_{i1}, \ldots, \theta_{il_j}$ for K_{ij} over K_{j-1}, and in Line 26 we pick an element from the associated Galois groups $\tilde{J}_{ij}$. The element $\beta_{ij} = \theta_{ij_1} + \zeta_{l_{ij}}\sigma_{ij}^{l_{ij}}\theta_{ij_1} + \ldots + \zeta_{l_{ij}}^{l_{ij}-1}\sigma_{ij}\theta_{i1}$ for some ordering of the θ's. Since the θ's are linearly independent, the element is perforce non-zero, hence by the theory of Lagrange resolvents, generates K_{ij} over K_{j-1}. In particular, β_{ij} satisfies an irreducible cyclic polynomial over K_{j-1}.

How long does the computation take? We claim no more than $O(r^{12+\epsilon}\log^{4+\epsilon}(r|f(y)|))$ steps, where $r = [L : k]$. Certainly the computations in Lines 1-11 are dominated by the time it takes to compute the Galois group; by Theorem 4.1 that can done in $O(r^{12+\epsilon}\log^{2+\epsilon}(r^2|f(y)|))$ steps. Lines 12 and 13 are a simple computation (even when $k \neq Q$). We run through the loop in loop in Lines 14-29 at most $O(\log r)$ times. Its running time is dominated by Line 25, which takes $O(r^{8+\epsilon}\log^{2+\epsilon}|f(y)|)$ steps. The computation in Line 32 takes no time at all.

To simplify the computations, we view the fields K_i as simple extensions of K, and do the computation in terms of a primitive element, with minimal polynomial $g_{ij}(t)$ over k. However, since what we want is a basis for L with a minimal depth of nesting, we will store two different bases for K_{ij}; one with the β_{ij} which come from the computations, and then $K_{ij} = k[t]/g_{ij}(t)$ written as an extension by a primitive element. We do not run into coefficient blowup in computing $g_{ij}(t)$ because each of the fields K_{ij} can be viewed as a subfield of $L(\zeta_l)$. (In the interests of making the algorithm as clear as possible, we have not included those fine points in the algorithm.) ∎

We have presented the running time analysis for the algorithm over Q. The only reason for that is simplicity. The only requirement for the base field is that one needs to be able to factor polynomials over k. Theorem A.4 shows how a factorization over k can be raised to a factorization over $k(\alpha)$. We can generalize **Algorithm Compute Polynomial** to compute a minimal polynomial for α over k. Of course, we can generalize Theorem 4.1,

and **Algorithms Compute Fixed Field** and **Compute Denesting.**

Note also that the running time for the algorithm **Compute Denesting** can be simply stated as the time required to compute the splitting field of $f(x)$ over Q, which is in turn the time needed to factor $f(x)$ over its splitting field.

5 Conclusions

There are a number of open questions which remain.

- Can the bound in Theorem 3.7 be made optimal, that is, can we improve it to $s \leq t$ without involving the form of the input (as per Corollary 3.8)? What is the tradeoff between this and the roots of unity?

- Suppose α is a nested radical over a field k, and suppose there is a denesting expression for α involving roots of unity. When is there a way to transform that to an expression of the same depth which avoids using the roots of unity? (The expression $\sqrt{\sqrt{5} - 5/2} = \zeta_5 - 1/\zeta_5$ makes it clear it that will not always be possible to do so.)

Our algorithm is not fast. Thus the other important issue in this problem is speed. To this we have several comments:

- Is there a way to perform these computations via a straightline program? The encoding in the straightline version would avoid the problem of coefficient blow-up. The main difficulty seems to be in computing the Galois group. We do not see how to determine the group without determining both $f(x)$, the minimal polynomial of α over k, and the splitting field of $f(x)$ over k. If a way could be found around these problems, then the entire algorithm could be sped up. We do not think this will be easy. In particular, it is likely that any technique which works here will also work to determine the Galois group in the solvable case. No polynomial time algorithms are presently known, and this would be a surprising and exciting breakthrough.

- It is important to remember that the bounds given here are really *upper* bounds. The running time is exponential *iff* the splitting field is of exponential degree over k. If the degree of the splitting field is polynomially bounded, then so is the running time of the algorithm.

Recall also that nearly all of the bounds follow from the factoring bound given in [11], and that bound is almost certainly not optimal. What is important to remember is that what is theoretically slow may be practically fast, and vice versa. In particular, the exponential time polynomial factorization algorithm of [9] is more practical than any of the current polynomial time algorithms for polynomial factorization. Any improvements in the running time for factoring polynomials will lead to great improvements in our bounds.

- Our algorithm does not specifically examine the issue of determining if the expression equals zero. (We answer it, of course, in computing the minimal polynomial for α over k, but not efficiently.) Examining it leads to some interesting open questions on linear combinations of nested radicals. The case of nested square roots was treated by Borodin et al [3].

Acknowledgements: Warm thanks to Tsuneo Tamagawa, for help with Theorem 3.7, and for many insights into Galois theory. Many thanks also to John Cremona whose thorough reading and many observations greatly improved this paper, and to an anonymous referee whose careful reading of an earlier version caught an error.

References

[1] E. Artin, *Galois Theory*, University of Notre Dame Press, 1942.

[2] L. Babai, E. Luks and Á. Seress, *Fast Management of Permutation Groups*, Proceedings of the 29th Annual Symposium on Foundations of Computer Science (1988), pp. 272-282.

[3] A. Borodin, R. Fagin, J. Hopcroft and M. Tompa, *Decreasing the Nesting Depth of Expressions Involving Square Roots*, J. Symb. Comput., 1 (1985), pp. 169-188.

[4] B. Caviness and R. Fateman, *Simplification of Radical Expressions*, Proc. SYMSAC 77.

[5] D. Knuth, The Art of Computer Programming, Volume 2: Seminumerical Algorithms, Addison-Wesley, Reading, Mass. 1969.

[6] S. Landau, *Factoring Polynomials over Algebraic Number Fields*, SIAM J. of Comput., 14, No. 1 (1985), pp. 184-195.

[7] S. Landau and G. Miller, *Solvability by Radicals is in Polynomial Time*, J. Comput. and Sys. Sci., 30 (1985), pp. 179- 208.

[8] S. Lang, *Algebra*, Addison-Wesley, Reading, Mass., 1971.

[9] A. K. Lenstra, *Lattices and Factorization of Polynomials*, Proceedings Eurocam '82, LNCS 144, pp. 32-39

[10] A.K. Lenstra, *Factoring Polynomials over Algebraic Number Fields*, Proceedings Eurocal (1983), Springer Verlag Lecture Notes in Computer Science, pp. 458-465.

[11] A.K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, *Factoring Polynomials with Rational Coefficients*, Mathematische Annelan, 261 (1982), pp. 513-534.

[12] M. Mignotte, *Some Inequalities about Univariate Polynomials*, Proceedings of the 1981 ACM Symposium on Algebraic and Symbolic Computation, pp. 195-199.

[13] J. Rotman, *The Theory of Groups*, Allyn and Bacon, Boston, Mass. (1973).

[14] P. Samuel, *Algebraic Theory of Numbers*, Hermann, Paris (1972).

[15] B. Trager, *Algebraic Factoring and Rational Function Integration*, Proceedings of the 1976 ACM Symposium on Symbolic and Algebraic Computations, pp. 219-226.

[16] P. Weinberger, and L. Rothschild, *Factoring Polynomials over Algebraic Number Fields*, ACM Transactions on Mathematical Software (1976), pp. 335-350.

[17] H. Weyl, *Algebraic Theory of Numbers*, Princeton University Press (1940).

[18] R. Zippel, *Simplification of Expressions Involving Radicals*, J. Symbolic Computation, 1 (1985), pp. 189-210.

A Appendix

As we stated earlier, we can view a nested radical α as a sequence of polynomials $\tilde{p}_i, \tilde{q}$ in $k[x_1, \ldots, x_{m-1}]$ such that

$$\alpha_1 = \sqrt[n_1]{\tilde{p}_1},\ \tilde{p}_1 \in k \tag{4}$$

$$\alpha_2 = \sqrt[n_2]{\tilde{p}_2(\alpha_1)} \tag{5}$$

$$\alpha_3 = \sqrt[n_3]{\tilde{p}_3(\alpha_1, \alpha_2)} \tag{6}$$

$$\vdots \tag{7}$$

$$\alpha_m = \tilde{q}(\alpha_1, \ldots, \alpha_{m-1}) + \sqrt[n_m]{\tilde{p}_m(\alpha_1, \ldots, \alpha_{m-1})}, \tag{8}$$

with $\alpha = \alpha_m$. We show how to compute the minimal polynomial of $\alpha = \alpha_m$ over k. This polynomial is polynomial size in $n = \Pi n_i$, $||\alpha_i||$, $||\tilde{q}(x_1, \ldots, x_{m-1})||$ and $||\tilde{p}_i(x_1, \ldots, x_{i-1})||$. Whether or not the minimal polynomial is polynomial in the length of the expression for α is an open question.

We begin with some background. In particular, we define the size of a polynomial over Z, and over an algebraic number field, and quote some important results on polynomial factorization in these domains.

Definition: Let $f(x) = a_n x^n + a_{n-1} x^{n-1} + \ldots + a_0$ be a polynomial in $Z[x]$. The size of $f(x)$, denoted $|f(x)|$, is defined to be $(\Sigma a_i^2)^{1/2}$.

Theorem A.1 *(Mignotte [12]) Let $f(t)$ be a primitive polynomial in $Z[t]$, and let $g(t)$, also primitive, be a factor of $f(t)$. Then $|g(t)| \leq 2^t |f(t)|$.*

Definition: Let $h(x_1, \ldots, x_n) = h_{\alpha_1, \ldots, \nu_1} x_1^{\alpha_1} \ldots x_n^{\nu_1} + \ldots + h_{\alpha_t, \ldots, \nu_t} x_1^{\alpha_t} \ldots x_n^{\nu_t}$ be a polynomial over $O_k[x_1, \ldots, x_n]$. Then the size of $h(x_1, \ldots, x_n)$, denoted $||h(x_1, \ldots, x_n)||$, is $(\Sigma h^2_{\alpha_j, \ldots, \nu_j})^{1/2}$.

Theorem A.2 *(Lenstra, Lenstra and Lovász [11]) Let $g(t) = a_n t^n + \ldots + a_0$ be a primitive polynomial of degree n over $Z[t]$. There is an algorithm to factor $g(t)$ into irreducible factors over $Z[t]$ which requires $O(n^{9+\epsilon} + n^{7+\epsilon} \log^{2+\epsilon} |g(t)|)$ steps.*

Definition: Let β be an algebraic number. We define the *size* of β, written $||\beta||$, to be the maximum of the absolute values of the conjugates of β.

Definition: Let $g(t)$ be a monic irreducible polynomial over Z, with roots $\alpha_1, \ldots, \alpha_m$. If $f(x) = \beta_n x^n + \ldots + \beta_0$, with $\beta_i = \Sigma_{j=0}^{m-1} b_{ij}\alpha^j$, then the *size* of $f(x)$, written $||f(x)||$, is defined to be the $max_i(\Sigma_{j=0}^{m-1} b_{ij}^2)^{1/2}$.

Theorem A.3 *(Weinberger & Rothschild [16]) Let $g(t)$ be a monic irreducible polynomial of degree m over Z, and let $K = Q[t]/g(t)$. Let β be a root of $f(x)$, a polynomial in $O_K[x]$. Then $||\beta|| \leq 1 + ||f(x)||$. Assume that $f(x)$ is monic, and that $h(x) = h_r x^r + \ldots h_0$ is a factor of $f(x)$ in $O_K[x]$. Then $||h(x)|| \leq m||f(x)||(m|g(t)|)^m$.*

Theorem A.4 *(Landau [6]) Let $g(t)$ be a monic irreducible polynomial of degree m over Z, and let $K = Q[t]/g(t)$. Let $f(x)$ be a polynomial of degree n over O_K. Then there is an algorithm to factor $f(x)$ into irreducible factors over O_K. It runs in $O(m^{9+\epsilon}n^{7+\epsilon}\log^2|g(t)|\log^{2+\epsilon}(||f(x)||(m|g(t)|)^n(mn)^n))$ steps.*

In order to calculate the minimal polynomial for α over k, we introduce two important concepts from number theory: the norm and the resultant. The norm relates elements in extension fields to elements in the base field. Let $\beta = a_0 + a_1\gamma + \ldots + a_{m-1}\gamma^{m-1}$ be an element of the field $k(\gamma)$. Then

$$Norm_{k(\gamma)/k}(\beta) = N_{k(\gamma)/k}(\beta) = \Pi_i(a_0 + a_1\gamma_i + \ldots + a_{m-1}\gamma_i^{m-1}),$$

where the product is over the conjugates of γ. If σ is an element of the Galois group of the minimal polynomial of γ over k, then $\sigma(\gamma) = \gamma_j$ for some conjugate γ_j of γ. Then

$$\begin{aligned}
\sigma_j(N_{k(\gamma)/k}(\beta)) &= \sigma_j(\Pi_i(a_0 + a_1\gamma_i + \ldots a_{m-1}\gamma_i^{m-1})) \\
&= \Pi_i\sigma_j(a_0 + a_1\gamma_i + \ldots + a_{m-1}\gamma_i^{m-1}) \\
&= \Pi_i(a_0 + a_1\gamma_i + \ldots + a_{m-1}\gamma_i^{m-1}) \\
&= N_{k(\gamma)/k}(\beta).
\end{aligned}$$

Since σ_j just permutes the $\gamma_i's$, we conclude that $N(\beta)$ is in k. Further, the norm is multiplicative. We can extend the definition of norm to include polynomials, by thinking of $f(x)$ in $k(\gamma)[x]$ as a polynomial in two variables: x and γ, and we denote it by $f_\gamma(x)$. Then

$$N_{k(\gamma)/k}(f(x)) = \Pi_i f_{\gamma_i}(x).$$

If $f(x)$ is in $k(\gamma)[x]$, then the norm of $f(x)$, $N_{k(\gamma)/k}(f(x))$, is in $k[x]$. We will need the following important property of norms of polynomials:

Lemma A.5 *(Weyl [17], Trager [15]) If $f(x)$ is an irreducible polynomial over $k(\gamma)$, then $Norm_{k(\alpha)/k}(f(x))$ is the power of an irreducible polynomial over k.*

How do we calculate the norm of a polynomial? The coefficients of the the norm are all symmetric functions in γ_i. A simple algorithm for computing norms using determinants has been known since the nineteenth century; this is the *resultant*. Let $g(t) = g_m t^m + g_{m-1}t^{m-1} + \ldots + g_0$, and $f(t) = f_n t^n + \ldots + t_0$. We define $Res_t(g(t), f(t)) =$

$$\begin{vmatrix}
f_n & 0 & 0 & \ldots & 0 & g_m & 0 & 0 & \ldots & 0 \\
f_{n-1} & f_n & 0 & \ldots & 0 & g_{m-1} & g_m & 0 & \ldots & 0 \\
f_{n-2} & f_{n-1} & f_n & \ldots & 0 & g_{m-2} & g_{m-1} & g_m & \ldots & 0 \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
f_0 & f_1 & f_2 & \cdots & f_{m-1} & g_{m-n} & g_{m-n+1} & g_{m-n+2} & \cdots & g_m \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdot & \cdot & \cdots & \cdot & \cdot & \cdot & \cdot & \cdots & \cdot \\
0 & 0 & \ldots & 0 & f_0 & 0 & 0 & \ldots & 0 & g_0
\end{vmatrix}$$

$$\underbrace{\qquad\qquad\qquad\qquad}_{m} \qquad \underbrace{\qquad\qquad\qquad\qquad}_{n}$$

The coefficients $g_i, f_j = 0$ whenever $i, j < 0$ respectively. It can be shown that $Res_x(g(x), f(x)) = g_m^n \Pi f(\gamma_i)$. Thus $N(f(x)) = \Pi f_{\gamma_i}(x) = (Res_t(g(t), f(x,t)))/g_m^n$, where $f(x,t)$ is $f(x)$ with $t's$ substituted in for $\gamma's$, and $g(t)$ is the minimal polynomial for γ over k.

The nesting expressions for α in equations 4-8 is general, but cumbersome. Instead of multivariate polynomials $\tilde{q}, \tilde{p}_i$, we would prefer to have univariate polynomials q, p_i. We can do this by using:

Lemma A.6 *(Trager [15]) Suppose $[k(\alpha,\beta) : k(\alpha)] = n$, and $[k(\alpha) : k] = m$. Then there is a $c < (mn)^2$ such that $k(\alpha,\beta) = k(\alpha + c\beta)$. If $h(x)$ is the minimal polynomial for β over $k(\alpha)$, then whenever $H(x) = Norm_{k(\alpha)/k}(h(x - c\alpha))$ is squarefree, $k(\alpha,\beta) \simeq k[x]/H(x)$.*

Lemma A.7 *(Landau [6]) Let k, α, β, c, m, and n be as above. If $g(x)$ is the minimal polynomial for α over Q, then $H(x)$, the minimal polynomial for $\alpha + c\beta$ over Q, has coefficients no larger than $(c|g(x)|||h(x)||)^{m+n}$.*

Thus for each α_i there is an associated γ_i, with $\gamma_1 = \alpha_1$, and $\gamma_i = \alpha_i + c_i\gamma_{i-1}$. Then $k(\alpha_1, \ldots, \alpha_i) = k(\gamma_i)$. The multivariate polynomials in equations 4-8 $\tilde{q}, \tilde{p}_i$ can be replaced by univariate polynomials q, p_i, and a nesting expression can be more simply written as:

$$\begin{aligned}
\alpha_1 &= \sqrt[n_1]{p_1},\ p_1 \in k, \\
\alpha_2 &= \sqrt[n_2]{p_2(\gamma_1)}, \gamma_1 = \alpha_1 \\
\alpha_3 &= \sqrt[n_3]{p_3(\gamma_2)}, \gamma_2 = \alpha_2 + c_2\gamma_1 \\
&\vdots \\
\alpha_m &= q(\gamma_{m-1}) + \sqrt[n_m]{p_m(\gamma_{m-1})}, \gamma_{m-1} = \alpha_{m-1} + c_{m-1}\gamma_{m-2},
\end{aligned}$$

with the $c_i < (\Pi_{j<i} n_j)^2$. Lemma A.6 gives a method for computing the minimal polynomial of γ_i over k. We will show how to compute these polynomials in **Algorithm [Compute Polynomial]**, where we also compute $f(x)$, the minimal polynomial for α over k. We will first sketch the idea behind the algorithm.

It is easy to see what the minimal polynomial for α_1 over k is. Suppose $\tilde{p}_i(x)$ is a d^{th} power of an element in k, say p_1, and let d be the largest integer dividing n_1 for which this statement is true. If we let $a_1 = n_1/d$, then $\sqrt[n_1]{\tilde{p}_1} = \sqrt[a_1]{p_1}$, and α_1 satisfies $f_1(x) = x^{a_1} - p_1$. If this polynomial is irreducible, we are done. We could factor to check irreduciblity, but a more efficient test is to use the criterion of Theorem 1.9. If the conditions of Theorem 1.9 are not satisfied (i.e. $4 \mid a_1$, ζ_{a_1} is not in k, $p_1 = -4j^4$ for some j in k), then $x^{a_1} - p_1$ is reducible. In this situation, we do factor $x^{a_1} - p_1 = \Pi j_i(x)$, and set $f_1(x)$ to be the $j_i(x)$ for which $j_i(p_1) = 0$. Of course, $g_1(t)$, the minimal polynomial for $\sqrt[a_1]{p_1(q)}$ over k, is equal to $f_1(t)$.

To compute the minimal polynomial for α_2 over k, we do nearly the same computation. We find d, the largest integer dividing n_2 such that $\tilde{p}_2(\alpha_1)$ is a d^{th} power of an element in $k(\gamma_1)$, say $p_2(q)$. Again, $a_2 = n_2/d$, $\sqrt[n_2]{\tilde{p}_2(\alpha_1)} = \sqrt[a_2]{p_2(q)}$, and α_2 satisfies $x^{a_2} - p_2(\gamma_1)$. We check irreducibility as before, and set $f_2(x)$ equal to the minimal polynomial of α_2 over $k(\gamma_1)$. By Theorem A.5, the polynomial

$$s_2(x) = N_{k(\gamma_1)/k}(f_2(x)) = Res_t(t^{a_1} - p_1(q), f_2(x,t))$$

is a power of an irreducible polynomial over k which α_2 satisfies. If $h_2(x)$ is that polynomial, then $h_2(x) = s_2(x)/\gcd(s_2(x), s_2'(x))$.

Next we compute the minimal polynomial of γ_2 over k. (Note that that is really equivalent to computing a primitive element of $k(\gamma_1, \alpha_2)$ over k.) By Lemma A.6, it suffices to find a c_2 such that $g_2(t) = N_{k(\gamma_1)/k}(f_2(t - c_2\gamma_1))$ is squarefree, and there is a $c < (a_1 a_2)^2$ which will work. We calculate the norm by

$$N_{k(\gamma_1)/k}(h_2(t - c_2\gamma_1)) = Res_y(g_1(y), h_2(t - c_2 y, y),$$

thinking of $h_2(t)$ in $k(\gamma_1)$ as a polynomial in t and y in $k[t, y]/g_1(y)$.

At this point, the pattern begins to emerge. What we are doing is first calculating the minimal polynomial of α_i over $k(\gamma_{i-1})$, either $x^{a_i} - p_i(\gamma_{i-1})$, or a divisor of it. Call it $f_i(x)$. Then $s_i(x) = Res_t(g_{i-1}(t), f_i(x, t))$ is the power of an irreducible polynomial over k, $h_i(x) = s_i(x)/\gcd(s_i(x), s_i'(x))$. We compute a primitive element of $k(\gamma_i) = k(\gamma_{i-1}, \alpha_i)$ over k. We do that by finding a c_i such that $g_i(t) = N_{k(\gamma_{i-1})/k}(f_i(t - c_i\gamma_{i-1}))$ is square free. We are then ready to repeat the process.

Algorithm: Compute Polynomial

input: $(\tilde{q}_i(x_1,\ldots,x_{i-1}), n_i, \tilde{p}_i(x_1,\ldots,x_{i-1})), i = 1,\ldots m$, where $\alpha_1 = \sqrt[n_1]{\tilde{p}_1}, \alpha_2 = \sqrt[n_2]{\tilde{p}_2(\alpha_1,\alpha_2)},\ldots,\alpha = \alpha_m = \tilde{q}(x_1,\ldots x_{m-1}) + \sqrt[n_m]{\tilde{p}_m(\alpha_1,\ldots,\alpha_{m-1})}$.
output: $f(x)$, the minimal polynomial of α over k, and $g_i(t)$, where $k(\alpha_1,\ldots,\alpha_i) \simeq k[t]/g_i(t)$ for $i = 1,\ldots,m$, and n =lcm of the indices of the reducible radicals $\hat{p}_i$.

Line 1: BEGIN
Line 2: Compute minimal polynomial for $\tilde{p}_1$, and call it $h_1(x), f_1(x)$;
Line 3: $g_1(t) \leftarrow f_1(t)$;
Line 4: $n \leftarrow 1$;
Line 5: FOR $i = 2$ TO $m-1$ DO:
Line 6: BEGIN
Line 7: Express $\alpha_1,\ldots,\alpha_{i-1}$ as elements in $k[t]/g_{i-1}(t)$
and Rewrite $\tilde{p}_i(\alpha_1,\ldots,\alpha_{i-1})$ as $p_i(t)$;
Line 8: Find $f_i(x)$, minimal polynomial for $p_i(t)$ over $k[t]/g_{i-1}(t)$;
Line 9: Write "α_i satisfies $f_i(t)$ over $k[t]/g_{i-1}(t)$";
(This is either $x^{n_i} - p_i(t)$, or a divisor of it;
we will pick a divisor of minimal degree.)
Line 10: If $\hat{p}_i$ is a reducible radical, $n \leftarrow \mathrm{lcm}(n_i, n)$;
Line 11: $s_i(x) \leftarrow Res_t(g_{i-1}(t), f_i(x))$;
Line 12: $h_i(x) \leftarrow s_i(x)/\gcd(s_i(x), s_i'(x))$;
Line 13: Find c_i such that $g_i(t) \leftarrow Res_y(g_{i-1}(y), f_i(t - c_i y, y))$
is square free;
Line 14: END;
Line 15: Rewrite $\tilde{p}_m(\alpha_1,\ldots,\alpha_{m-1})$ as $p_m(t)$,
and rewrite $\tilde{q}(\alpha_1,\ldots,\alpha_{m-1})$ as $q(t)$;
Line 16: Find $f_m(x)$, minimal polynomial for $q(t) + \sqrt[n_m]{p_m(t)}$
over $k[t]/g_{m-1}(t)$;
(This is either $(x - \tilde{q}(t))^{n_m} - \tilde{p}_m(t)$, or a divisor of it.)
Line 17: $s_m(x) \leftarrow Res_t(g_{m-1}(t), f_m(x))$;
Line 18: $f(x) \leftarrow s_m(x)/\gcd(s_m(x), s_m'(x))$;
Line 19: Find c_m such that $g_m \leftarrow Res_t(g_{m-1}(y), f_m(t - c_m y, y))$
is square free;
Line 20: END.

Theorem A.8 *On input the sequence* $\alpha_1 = \sqrt[n_1]{\tilde{p}_1},\ldots,\alpha_m = \tilde{q}(\alpha_1,\ldots,\alpha_{m-1}) + \sqrt[n_m]{\tilde{p}_m(\alpha_1,\ldots,\alpha_{m-1})}$, *where the* $\tilde{p}_i, \tilde{q}$ *are polynomials in* $O_k[x_1,\ldots,x_m]$ *for*

$i = 2, \ldots, m$, *and* $\tilde{p}_1$ *is in* O_k, *the* **Algorithm [Compute Polynomial]** *computes* $f(x)$, *the minimal polynomial for* $\alpha = \alpha_m$ *over* k, *and the polynomials* $g_i(t)$, $i = 1, \ldots, m$, *where* $k(\alpha_1, \ldots, \alpha_i) \simeq k[t]/g_i(t)$. *It does so in* $O(n^{21} \log^{4+\epsilon}(\mathcal{AP}))$ *steps, where* $\mathcal{A} = max_i||\alpha_i||$, *and* $\mathcal{P} = max(||\tilde{p}_i(x_1, \ldots, x_{i-1})||, ||\tilde{q}(x_1, \ldots, x_{m-1})||)$.

Proof As usual, we begin with a proof of correctness. We will prove that $g_i(t)$ in $O_k[t]$ is an irreducible polynomial generating $k(\alpha_1, \ldots, \alpha_i)$, i.e. that $k(\alpha_1, \ldots, \alpha_i) \simeq k[t]/g_i(t)$, that $f_i(x)$ is the minimal polynomial of α_i over $k[t]/g_{i-1}(t)$, and that $h_i(t)$ is the minimal polynomial for α_i over k.

That $f_1(x), h_1(x)$ and $g_1(t)$ calculated in Lines 2 and 3 satisfy these conditions is clear. We prove by induction that the $f_i(x), h_i(x)$ and $g_i(t)$ calculated in Lines 4-14 do so also. It is clear that α_i satisfies $x^{a_i} - p_i(t)$. By using the criterion of Theorem 1.9, and factoring if necessary, we find $f_i(x)$, the minimal polynomial of α_i over $k[t]/g_{i-1}(t)$.

Since $g_{i-1}(t)$ generates $k(\alpha_1, \ldots, \alpha_{i-1})$ over k, we have $Res_t(g_{i-1}(t), f_i(x)) = N_{k(\alpha_1,\ldots,\alpha_{i-1})/k}(f_i(x))$ is a polynomial over k which α_i satisfies. Furthermore, by Lemma A.5, $s_i(x) = Res_t(g_{i-1}(t), f_i(x))$ is the power of an irreducible polynomial over k. Thus $h_i(x)$ is the minimal polynomial of α_i over k. Finally, by Lemma A.6, the $g_i(t)$ calculated in Line 13 is the minimal polynomial of $\alpha_i + c_i\gamma_{i-1}$, where γ_{i-1} is a root of $g_{i-1}(t)$, and thus generates $k(\alpha_1, \ldots, \alpha_i)$, that is $k(\alpha_1, \ldots, \alpha_i) \simeq k[t]/g_i(t)$.

How long does the computation take? In order to discuss the time issue, we need to know how big the coefficients of $f_i(x), s_i(x), h_i(x)$ and $g_i(t)$ can be. We analyze that first. Our bounds are not tight, but are chosen for the relative simplicity of the analysis.

Let $n =$ degree $(f(x))$,and let $\mathcal{A} = max_i||\alpha_i||$. Observe that $h_i(x)$ has root α_i over k, and that thus each coefficient of $h_i(x)$ is less than $2^n||\alpha_i||^n$ by the binomial theorem. Therefore $|h_i(x)| < (n(2^n||\alpha_i||^n)^2)^{1/2} < (2\mathcal{A})^{n+\epsilon}$. Then $\gamma_i = c_1 \ldots c_i\alpha_1 + c_1 \ldots c_{i-1}\alpha_2 + \ldots + \alpha_{i-1}$, a root of $g_i(t)$, is less than $n^{2n}\mathcal{A}$, and the coefficients of $g_i(t)$ are less than $2^n(n^{2n}\mathcal{A})^n$. Thus $|g_i(t)| < (n(2^n(n^{2n}\mathcal{A})^n)^2)^{1/2} < (n\mathcal{A})^{4n^2+\epsilon}$.

It is easy to analyze the size of $s_i(x)$. Observe that all the irreducible factors of $s_i(x)$ appear in $h_i(x)$. Thus $|s_i(x)| < 2^n|h_i(x)|^n < (4\mathcal{A})^{n^2+\epsilon}$.

Finally we are ready to tackle $f_i(x)$. Rewriting $\tilde{p}_i(\alpha_1, \ldots, \alpha_{i-1})$ as $p_i(t)$ will increase the coefficient size only slightly. This is because all we are doing is a matrix computation. Let $N_i = \Pi_{j \le i} n_j$. The coefficients of $f_i(x)$ are bounded by $(|g_i(t)|||\tilde{p}_i(x_1, \ldots, x_{i-1})||)^{N_j} N_j! < n!\mathcal{A}^{4n^2+\epsilon}||\tilde{p}_i||^n$.

The running time of the algorithm is dominated by the factorizations of Lines 2, 8, 11, 13 and 16. The factorizations of Lines 8 and 16 are over algebraic extensions of k, and are thus most expensive. They take at most

$$O(n^{9+\epsilon}n^{7+\epsilon}\log^2(n\mathcal{A}^{4n+\epsilon})\log^{2+\epsilon}(n!\mathcal{A}^{4n^2+\epsilon}||\tilde{\mathcal{P}}||^n(n(n\mathcal{A})^{4n+\epsilon})^n(n^2)^n))) < O(n^{20+\epsilon}\log^{4+\epsilon}(\mathcal{AP}))$$

steps, where $\mathcal{P}=max(||\tilde{p}_i(x_1,\ldots,x_{i-1})||,||\tilde{q}(x_1,\ldots,x_{m-1})||)$. The loop is repeated at most m times, hence the bound in the theorem. ∎