How to Tangle with a Nested Radical

Susan Landau*
Computer Science Department
University of Massachusetts

July 13, 1993

1 Introduction

Like many an intriguing question in algebraic manipulation, the problem of denesting nested radicals had its origins with Ramanujan. That is not to say that no one had ever considered the problem of denesting radicals before he did. Certainly the fact that

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

is simple enough that it must have been known several centuries ago. Ramanujan [9] upped the ante. For each of the formulae below, he took the doubly nested radical on the left and simplified it to a combination of singly nested radicals on the right:

$$\begin{aligned}
\sqrt[3]{\sqrt[3]{2}-1} &= \sqrt[3]{1/9}-\sqrt[3]{2/9}+\sqrt[3]{4/9} \\
\sqrt{\sqrt[3]{5}-\sqrt[3]{4}} &= 1/3(\sqrt[3]{2}+\sqrt[3]{20}-\sqrt[3]{25}) \\
\sqrt[6]{7\sqrt[3]{20}-19} &= \sqrt[3]{5/3}-\sqrt[3]{2/3} \\
\sqrt[4]{\frac{3+2\sqrt[4]{5}}{3-2\sqrt[4]{5}}} &= \frac{\sqrt[4]{5}+1}{\sqrt[4]{5}-1}
\end{aligned}$$

*Supported by NSF grant DMS-8807202.

$$\sqrt{\sqrt[3]{28}-\sqrt[3]{27}} = 1/3(\sqrt[3]{98}-\sqrt[3]{28}-1)$$
$$\sqrt[3]{\sqrt[5]{32/5}-\sqrt[5]{27/5}} = \sqrt[5]{1/25}+\sqrt[5]{3/25}-\sqrt[5]{9/25}.$$

What Ramanujan neglected to do was provide a theory for simplifying nested radicals. When computers came along, symbolic computation became important. There was a practical reason to find an algorithm for denesting nested radicals.

A machine has no problem with:

$$1, \sqrt{5+2\sqrt{6}}, 5+2\sqrt{6}, (5+2\sqrt{6})^{3/2}$$

as a basis for $Q(\sqrt{5+2\sqrt{6}})$ over Q. Most human beings seem to prefer the basis:

$$1, \sqrt{2}, \sqrt{3}, \sqrt{6}.$$

The difficulty is that there was no general method to go from the complex form of a nested radical to a simplified version. If Ramanujan had one, he never wrote it down.

Necessity has often been the mother of invention, and so it proved to be in this case. Although the general problem remains open, there are now solutions to a number of subproblems: for denesting real nested square roots [3], for radicals of a special form [10], [7], for radicals using roots of unity [6], [5]. We are interested in three questions: When does a simplification exist? Is there a technique for finding it? How long does it take? In this article, we will briefly present some recent theorems for radical simplification, and the algorithms they lead to. For proofs, and complete presentations, the reader is urged to read the original papers.

2 What Does it Mean to Denest a Radical?

We begin by making precise what we mean by simplifying a nested radical[1] . Assume all fields are characteristic 0. Following [3], a *formula* over a field k and its *depth of nesting* are defined recursively:

[1]This brief overview is taken from [6]. A more detailed discussion of these issues mays be found in [5].

- An element in the field k is a formula of depth 0 over k.

  Thus 17 is of depth 0 over Q, while $1+\sqrt{2}$ is of depth 0 over $Q(\sqrt{2})$.

- An arithmetic combination (A$\pm$ B, A $\times$ B, A/B) of formulas A and B is a formula whose depth over k is max(depth(A), depth(B)).

  Since we view $\sqrt{5+2\sqrt{6}}-\sqrt{2}-\sqrt{3}$ as a sequence of formal symbols, it has nesting depth 2 over Q.

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

  Finally, $\sqrt{\sqrt{5+2\sqrt{6}}}$ has depth 3 over Q.

We will call such a formula a nested radical. A nesting of α means any formula A that can take α as a value. But there are difficulties involved since an n^{th} root is a multivalued function. When we write the equation

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

it is unclear which $\sqrt{2}$ 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.

We start with the input as a sequence of expressions of the form:

$$\begin{aligned}
\alpha_1 &= \sqrt[n_1]{q},\ q\in k\\
\alpha_2 &= \sqrt[n_2]{\tilde{p}_2(\alpha_1)}, \tilde{p}_2\in k[x_1]\\
\alpha_3 &= \sqrt[n_3]{\tilde{p}_3(\alpha_1,\alpha_2)}, \tilde{p}_3\in k[x_1,x_2]\\
&\vdots\\
\alpha_m &= \tilde{q}(\alpha_1,\ldots,\alpha_{m-1})+\sqrt[n_m]{\tilde{p}_m(\alpha_1,\ldots,\alpha_{m-1})}, \tilde{q} \text{ and } \tilde{p}_m\in k[x_1,\ldots,x_{m-1}],
\end{aligned}$$

with $\alpha=\alpha_m$. It is not hard to go from this complicated sequence to the minimal polynomial for α over k. We can do it by first determining a minimal polynomial for α_1 over k, then using that to determine a minimal polynomial for α_2 over k, etc. (see [6] for details). One must take careful note of the choices of roots as they are made.

Once we choose a particular n^{th} root for $\sqrt[n]{\alpha}$, the same value must be assigned to it each time it appears. If the roots are specified at the time a nested radical is given, we choose those roots. Whenever roots appear which have not been previously specified, we are free to arbitrarily pick a value for them, so long as after that we consistently choose the same value to represent the root.

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 $\alpha=\sqrt[3]{\sqrt[3]{2}-1}$.

For the purposes of this paper, we have chosen that 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.

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 and B have the same (real or complex) value. We will say that *A can be denested in the field L* if there is a formula B of lower nesting depth than A with all of the terms (subexpressions) of B lying in L, again with A and B having the same value. For any α, we define the depth of α over k 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 say that α can be denested.

For the remainder of this paper, we will assume that α has been given by its minimal polynomial over k, and the choice of roots in any ambiguous situation have been spelled out.

3 Denesting Nested Square Roots

Nested square roots form the simplest example of nested radicals. Since nested real square roots describe the Euclidean distance from one vertex on a polyhedron to another, an algorithm for their simplification is potentially

of practical value. With such concerns in mind, Borodin, Fagin, Hopcroft and Tompa [3] studied simplifying nested square roots. Their first result demonstrates when square roots suffice for denesting:

Theorem 3.1 *(Borodin, Fagin, Hopcroft and Tompa [3]) Let $Q \subseteq k$, and let a, b, r be elements of k, with $\sqrt{r}$ not in k. Then the following are equivalent:*

1. $\sqrt{a+b\sqrt{r}}$ *is in* $k(\sqrt{r}, \sqrt{a_1}, \ldots, \sqrt{a_k})$ *for some* $a_1, \ldots, a_k$ *in* k.
2. $\sqrt{a+b\sqrt{r}}$ *is in* $k(\sqrt{s}, \sqrt{r})$ *for some* $s \neq 0$ *in* k.
3. $\sqrt{a^2-b^2r}$ *is in* k.

Next they gave the conditions under which fourth roots may help:

Theorem 3.2 *(Borodin, Fagin, Hopcroft and Tompa [3]) Let $Q \subseteq k$, and let a, b, r be in k, with $\sqrt{r}$ not in k. Then the following are equivalent:*

1. $\sqrt{a+b\sqrt{r}}$ *is in* $k(\sqrt[4]{r}, \sqrt{a_1}, \ldots, \sqrt{a_k})$ *for some* $a_1, \ldots, a_k$ *in* k.
2. *Either* $\sqrt[4]{r}\sqrt{s}\sqrt{a+b\sqrt{r}}$ *or* $\sqrt{s}\sqrt{a+b\sqrt{r}}$ *is in* $k(\sqrt{r})$, *for some* $s \neq 0$ *in* k.
3. *Either* $\sqrt{r(b^2r-a^2)}$ *or* $\sqrt{a^2-b^2r}$ *is in* k.

Finally, they showed that for denesting expressions containing only real square roots only square roots or fourth roots play a role:

Theorem 3.3 *(Borodin, Fagin, Hopcroft and Tompa [3]) Let k be a real extension of Q, and let a, b, r be in k with $\sqrt{r}$ not in k. Let $n_1, \ldots, n_l \geq 1$, and let $a_1, \ldots, a_l$ in k be positive. If $\sqrt{a+b\sqrt{r}}$ is in $k(\sqrt[n_1]{a_1}, \ldots, \sqrt[n_l]{a_l})$, then $\sqrt{a+b\sqrt{r}}$ is in $k(\sqrt[4]{r}, \sqrt{p_1}, \ldots, \sqrt{p_l})$ for some $p_1, \ldots, p_l$ in k.*

In combination, these three theorems provide the backbone for an algorithm for denesting real nested square roots. Let r_0 be in Q, and inductively let $r_i = a_i + b_i\sqrt{r_{i-1}}$ with a_i and b_i in $Q(r_{i-1})$ for $i = 1, \ldots, n$. Consider the following tower of fields:

$$k_0 = Q \text{ and for all } i \geq 0, k_{i+1} = k_i(\sqrt{r_i}), \text{ where } r_i \in k_i.$$

There are two ways $\sqrt{r_n}$ can denest using only square roots. One is that $\sqrt{r_{n-1}}$ may denest. The other is that there is a field K satisfying (i) it contains a, b, r_{n-1}, (ii) it contains only elements of depth $n-1$, and (iii) $a^2 - b^2 r_{n-1}$ is a square.

To find the field K in which these conditions are satisfied, we attempt to denest $\sqrt{a^2 - b^2 r_{n-1}}$. If this is accomplished, we will have found a field $\hat{k}$ in which all elements have depth at most $n-2$. We will also have found an element s such that $\sqrt{a^2 - b^2 r_{n-1}}$ is in $\hat{k}(\sqrt{s}, \sqrt{r_{n-2}})$. Then $K = \hat{k}(\sqrt{s}, \sqrt{r_{n-2}})$. This idea leads to a recursive algorithm. If one is careful about how to handle the input and output, the running time of this algorithm is polynomial.

But this tells us only how to handle a nested square root which consists of roots recursively formed by $r_i = a_i + b_i\sqrt{r_{i-1}}$ with a_i and b_i in $Q(r_{i-1})$ for $i = 1, \ldots, n$, with r_0 in Q. If we want to denest all real nested square roots, we have to be able to handle linear combinations of nested square roots. We look first at the simpler case, in which none of the radicals are nested.

Theorem 3.4 *(Besicovitch [2]) Let $\{e_i\}$ denote the set of n^l radicals,*

$$\sqrt[n]{p_1^{m_1} p_2^{m_2} \ldots p_l^{m_l}}, 0 \leq m_i \leq n, 1 \leq i \leq l,$$

where $p_1, \ldots, p_l$ are the first l primes. Then the set $\{e_i\}$ is linearly independent over Q.

To check if a linear combination of square roots is equal to zero, one needs only to check if it is trivially equal to zero. There is the obvious solution of factoring all the integers under the square root sign in order to check if pieces cancel, but this is far from optimal. A much faster way to do the problem is to compute gcd's of the integers under the square root signs. Then where appropriate, pull out squares from under the radical signs. Combine like terms. Repeat until no further simplification is possible. If there is anything left, then the combination of square roots is different from 0.

One can implement a more complex version of this idea to simplify linear combinations of nested square roots. Borodin et al [3] consider the case where k is a real extension of Q, and where l_i, a_i, b_i and $\sqrt{r_i}$ are in k for $i = 1, \ldots, h$. Suppose $\Sigma_{i=1}^{h} l_i\sqrt{a_i + b_i\sqrt{r_i}}$ denests in k. They find the denesting as follows. First denest any single radical that can be denested using the criteria of Theorem 3.3. Next consider each pair of nested radicals,

$\sqrt{a_i + b_i\sqrt{r_i}}, \sqrt{a_j + b_j\sqrt{r_j}}$, and see if the product denests. Suppose it is equal to m in k. Replace $l_i\sqrt{a_i + b_i\sqrt{r_i}} + l_j\sqrt{a_j + b_j\sqrt{r_j}}$ by $(l_i + \frac{ml_j}{a_i + b_i\sqrt{r_i}})\sqrt{a_i + b_i\sqrt{r_i}}$ in $k(\sqrt{a_i + b_i\sqrt{r_i}})$. Iterate the process of looking for a pair of radicals that denests. If at any point the product of any pair of radicals cannot be further denested, then the combination of nested radicals cannot be further denested.

This settles the question of efficiently denesting real nested square roots, but it leads to:

Open Question 1: Is there a polynomial time algorithm for determining whether a linear combination of (not necessarily real) nested square roots is equal to zero?

Open Question 2: Is there a polynomial time algorithm for determining whether a linear combination of nested radicals (not necessarily square roots) is equal to zero?

4 A Particularly Simple Form of Denesting

Zippel studied the equation

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

and found the beginnings of a pattern. While $\sqrt[3]{2} - 1$ is not a cube in $Q(\sqrt[3]{2})$, $9(\sqrt[3]{2} - 1)$ is. We can find a γ which satisfies $\gamma^3 = 9(\sqrt[3]{2} - 1)$ in $Q(\sqrt[3]{2})$ by factoring $x^3 - 9(\sqrt[3]{2} - 1)$ over $Q(\sqrt[3]{2})$. We find:

$$x^3 - 9(\sqrt[3]{2} - 1) = (x - (1 - \sqrt[3]{2} + \sqrt[3]{4}))(x^2 + (1 - \sqrt[3]{2} + \sqrt[3]{4})x + 3\sqrt[3]{4} - 3).$$

The first factor gives Ramanujan's denesting.

Zippel noticed that a similar situation arises with $5 + 2\sqrt{6}$. Again, this is not a square in $Q(\sqrt{6})$, but a multiple of it, $2(5 + 2\sqrt{6})$, is. Our task is to find γ, where $\gamma^2 = 2(5 + 2\sqrt{6})$ in $Q(\sqrt{6})$. We discover:

$$x^2 - 2(5 + \sqrt{6}) = (x - (2 + \sqrt{6}))(x - (-2 - \sqrt{6})).$$

In both cases, we found an element β in Q such that, although $\sqrt[d]{\alpha}$ is not

an element of $Q(\theta)$, $\sqrt[d]{\alpha\beta}$ is. We have the following picture:

$$L = K(\sqrt[d]{\alpha}) = KF$$

$$K \qquad\qquad F = k(\sqrt[d]{\alpha})$$

$$k = K \cap F,$$

In each case expressions of nesting depth n in the field L have been dropped to expressions of nesting depth $n-1$ in the subfield L. This idea generalizes to a theorem:

Theorem 4.1 *(Zippel [10][2]) Assume K is an extension of k, a field containing a primitive d^{th} root of unity. Let $L = K(\sqrt[d]{\alpha})$ be an extension of degree d, where α is in K. If there is a field F which is a Galois extension of $k = K \cap F$, and $L = KF$, then there is a β in k such that $\alpha\beta$ is a d^{th} power of an element of k. Furthermore, $F = k(\sqrt[d]{\beta})$.*

Zippel exploited some lucky guesses. If we want an algorithm, we will need something somewhat more deterministic than that. We will need an algorithm to determine whether such a β exists, and if so, how to find it. More precisely, given $\sqrt[d]{\alpha}$ in L, when is there a solution to $\alpha\beta = \gamma^d$, with β in k, a proper subfield of $k(\sqrt[d]{\alpha})$, and γ in K, a proper subfield of L. It is not hard to show that the following converse of Zippel's theorem holds:

Theorem 4.2 *(Landau [7]) Let α be an element of a field K. Suppose that $\sqrt[d]{\alpha}$ is of degree d over K, and that $\sqrt[d]{\alpha} = \lambda/\sqrt[d]{\beta}$ with λ in K, and β in $k \subset K$. Assume that the d^{th} roots of unity lie in k. Then the field $F = k(\sqrt[d]{\beta})$ satisfies: (i) F over k is Galois and the Galois group of F over $F \cap K$ is isomorphic to the group of FK over K, (ii) $FK = K(\sqrt[d]{\alpha})$, and (iii) $k = F \cap K$.*

Thus if we seek a "Zippel denesting", we are asking to find subfields of $K(\sqrt[d]{\alpha})$ satisfing Theorem 4.2. In [8], there is a polynomial time algorithm

[2]Zippel's original statement omitted, but implicitly assumed, the hypothesis that F is a Galois extension of k. A corrected version appears in [7].

to find maximal subfields of a field. We can use this for an algorithm for determining whether a Zippel denesting exists. The first step is to find all maximal subfields of L.

It is a simple matter to check if a field extension is Galois. One finds a primitive element for the larger field, possibly by resorting to the well-known construction that $k(\gamma, \rho) = k(\gamma + c\rho)$ for some $c \leq (\deg_k(\gamma)\deg_k(\rho))^2$. With a primitive element, say α, which has minimal polynomial $p(x)$, one can compute the action of the Galois group by observing that the factorization:

$$p(x) = (x - p_1(\alpha))(x - p_2(\alpha)) \ldots (x - p_m(\alpha))$$

gives the group table since $\sigma_i(\alpha_j) = (p_i(p_j(\alpha)) \pmod{p(x)}$.

Computing whether a candidate subfield satisfies parts (ii) and (iii) are even easier. This is just a matter of checking whether two fields are equal. This can be done by seeing whether the basis of one is contained in the other, and vice versa. Iterating the procedure in [8] will give an algorithm to find all subfields. Thus there is an algorithm to determine if a "Zippel denesting" exists.

It is not fast. There may be 2^d fields between k and L, and in worst case, we would have to check each one of them. But this exponential time algorithm can still be quite reasonable for small (< 10) values of d.

Zippel used his theorem to shed some light on the calculations and theorems of Borodin et al [3]. Let a, b, q be elements of a field k, and suppose we are hoping to denest $\sqrt{a + b\sqrt{q}}$. Assume there is a β in k such that

$$\beta(a + b\sqrt{q}) = (a_0 + \sqrt{q})^2.$$

Now $(a^2 - qb^2)$, the norm of $a + b\sqrt{q}$, is a square, d^2. This leads to:

$$\beta = \frac{2}{b^2}(a \pm d), a_0 = \frac{1}{b}(a \pm d).$$

Choosing the positive sign, we find:

$$\sqrt{a + b\sqrt{d}} = \sqrt{\frac{a+d}{2}} \pm \sqrt{\frac{a-d}{2}},$$

where the sign depends upon the sign of b.

If $a^2 - qb^2$ is not a square, then that means that that $\sqrt{a + b\sqrt{q}}$ does not denest in a quadratic extension $K = k(\sqrt{q})$, and we must look for a quartic extension in which it denests. We try looking for a β such that $\beta(a\sqrt{q} + bq)$ is a perfect square. That computation eventually leads to:

$$\sqrt{a + b\sqrt{q}} = \frac{1}{2(bq+d)}(\sqrt[4]{4q(bq+d)^2} + \sqrt[4]{(4q(bq+d)^2)^3}).$$

These two denesting formulae are the two shown in [3] to be the only ways in which expressions involving square roots can be denested.

5 What about the General Case?

Zippel's criteria are simple and elegant, but the conditions on Theorems 4.1 and 4.2 are sufficiently restrictive that they will not handle all cases. Trying to understand where the 9 came from in:

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

this author discovered the following subfields[3] of $Q(\sqrt[3]{\sqrt[3]{2} - 1})$:

$$Q(\sqrt[3]{\sqrt[3]{2} - 1}) = Q(\sqrt[3]{2}, \sqrt[3]{9})$$

$$Q(\sqrt[3]{2}) \qquad\qquad Q(\sqrt[3]{9})$$

$$Q$$

The denesting

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

[3]This diagram gives a partial explanation of where the 9 comes from. A more complete answer is that 9 divides the discriminant of $x^9 + x^6 + x^3 - 1$, the minimal polynomial of $\sqrt[3]{\sqrt[3]{2} - 1}$ over Q.

led to a similar tower of fields:

$$Q(\sqrt[6]{7\sqrt[3]{20}-19}) = Q(\sqrt[3]{20}, \sqrt[3]{18})$$

$$Q(\sqrt[3]{20}) \qquad Q(\sqrt[3]{18}) \qquad Q(\sqrt[3]{45})$$

$$Q$$

In fact, such pictures arose for *all* the simplifications this author had occasion to try. This was too beautiful to happen by accident. A natural place to search for a denesting is the smallest closed field in which the minimal polynomial of α factors completely – the splitting field. The answer almost turned out to be that surprisingly simple and elegant. A minimal depth expression for a nested radical can always be found in the splitting field, *provided all roots of unity lie in the base field.* More precisely:

Theorem 5.1 *(Landau [6]) 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.*

Proof (This proof is sufficiently simple to include a sketch.) Consider the following diagram:

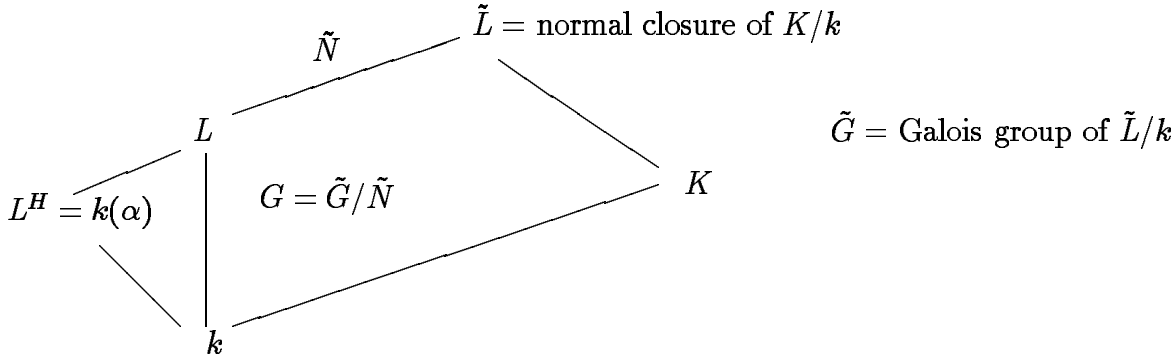

Let α be a nested radical over k, and suppose that L is the splitting field of $k(\alpha)$ over k. Let $\tilde{L}$ be the normal closure of K over k, where K is a field which contains a minimal depth nested expression for α. If we let $\tilde{G}$ be the Galois group of $\tilde{L}$ over k, and G be the group of L over k, then we note that $G = \tilde{G}/\tilde{N}$, where $\tilde{N}$ is the group of $\tilde{L}$ over L. (Note that $\tilde{N}$ is normal since $\tilde{L}$ is normal over L.) Since α can be denested over $\tilde{L}$, there is a sequence of subgroups $\tilde{H}_1, \ldots, \tilde{H}_l$ of $\tilde{G}$ with $\tilde{G} = \tilde{H}_0 \rhd \tilde{H}_1 \rhd \ldots \rhd \tilde{H}_l$, with $\tilde{H}_i/\tilde{H}_{i+1}$ abelian for $i = 0, \ldots, l-1$, and $\tilde{H} \supset \tilde{H}_l$. This sequence can be pulled down to a sequence in G. If all roots of unity are in k, the tower defined by the groups can be made into a tower of radicals extensions, thus showing there is a minimal depth denesting in L. ∎

Of course, this does not solve the original problem, which was to denest over an arbitrary field. Can one a priori add certain roots of unity to the base field so that a minimal depth nesting can be achieved? The answer is yes, so long as we are careful as to how we handle roots of unity.

All roots of unity can be expressed in terms of radicals. The problem is that the depth of nesting of a root of unity can be very deep indeed. In general, a p^{th} root of unity has nesting depth one more than the maximum of the nesting depths of the prime factors of $p-1$. Thus if there are arbitrarily long sequences of primes $p, 2p+1, 2(2p+1)+1, \ldots$ – a plausible, but unproved conjecture in number theory – then an n^{th} root of unity can have nesting depth $\log n$.

For this author the motivation for studying the denesting of radicals was to develop an algorithm for radical simplification. In many applications, writing a root of unity as ζ_n instead of the nested radical is a perfectly reasonable solution. This was the route taken here. But adding roots of unity to k does change the field in unexpected ways.

By the Kronecker-Weber Theorem, every abelian extension over Q can be embedded in an cyclotomic extension. When we attempt to write $\sqrt[n]{\alpha}$ in $Q(\zeta_l)$ we may find that $\sqrt[n]{\alpha}$ is an irrational number which is already 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 in some sense we are masking it. There are subtle ways in which we pay for that. For example, $\sqrt{\sqrt{5} - 5/2} = \zeta_5 - 1/\zeta_5$. Which symbol is easier to understand: $\sqrt{\sqrt{5} - 5/2}$ or $\zeta_5 - 1/\zeta_5$? That depends

on the application – or the mathematician.

Taking these concerns into consideration, we find:

Theorem 5.2 *(Landau [6]) 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, and let us write a primitive l^{th} root of unity as ζ_l, and not simply as a nested radical. 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 5.3 *(Landau [6]) Let k, α, L, G, l, t be as in Theorem 5.2. 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 $L(\zeta_r)$.*

These theorems tell us that the splitting field is the right place to look. They also lead naturally to an algorithm. If we wish to denest the nested radical α, we begin by computing the minimal polynomial of α over k. From that we construct the splitting field L of the minimal polynomial of α over k. Next we compute $G = Gal(L/k)$. We have already seen how to do these computations. What is the shortest sequence of nested radicals that will give α? It will come from the shortest sequence of groups in the Galois group, the series of commutator subgroups $D^iG, i = 1, \ldots, s$, where $D^sG = \{e\}$. Good algorithms for group computations have existed for over fifteen years. Having a group table, or an equivalant, for G, it is not hard to compute the commutator series of the group.

Next we also compute l, the lcm of the exponent s of the derived series of G. 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. Then $K_{ij} = K_{i-1}(\tilde{\beta}_{ij})$ 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}. The crux of the matter is how to write these extensions as radical ones, that is, $K_i = K_{i-1}(\sqrt[q_i]{\alpha_{i-1}})$. 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}. That each of the extensions can be achieved as radical extensions stems from the fact that the appropriate roots of unity lie in the base field.

This is just a brief sketch of the algorithm, details of which can be found in [6]. But the point should be clear: there is an algorithm for simplifying nested radicals, assuming one allows roots of unity to creep into the expression.

How long does this take? Too long! If α is of degree n over k, its Galois group may be of size $n!$ Even groups which are exponentially large (S_n, A_n, etc.) have a small set of generators, but from a computational standpoint that does not seem to help. No one knows how to determine the generators of a Galois group of a general polynomial without first determining its splitting field. In general, computing the splitting field (that is, finding a minimal polynomial for a generator over the base field) is an exponential time computation. Allowing roots of unity written in shorthand, Theorem 5.1 limits where we have to search if we are seeking a denesting. Nonetheless, except for radicals with small degrees, the computation is presently infeasible. Until there are improvements in algorithms for splitting field and Galois group computations, the algorithms based on Theorems 5.1, 5.2 and 5.3 are useful only for nested radicals of small degree.

Open Question 3: Find a polynomial time algorithm to compute the Galois group of an irreducible polynomial over Q.

There is an improvement one can make to Corollary 5.3. Horng and Huang [5] have shown that:

Theorem 5.4 *(Horng and Huang [5]) Let k, α, L, G, l, t be as in Theorem 5.2. Let n be a natural number which is divisible by $[L : k]$ and the discriminant of L over Q. Then there is a minimal depth nesting of α over $k(\zeta_n)$ with each of its terms lying in $L(\zeta_n)$.*

In finding a root of unity to achieve minimal depth nesting for α, they eliminate the need for including anything which relies on the presentation of α, as in Theorem 5.3. However, they do so at the expense of introducing a root of unity of the discriminant. The discriminant is of exponential size in α. Their algorithm for denesting follows the algorithm described earlier.

The problem of simplification of nested radicals is far from completely solved. Although Theorems 5.1, 5.2, 5.3 and 5.4 require roots of unity for denesting, none of Ramanujan's denestings actually use them. Thus the question that started this all remains:

Open Question 4: Without a special encoding for roots of unity, given a nested radical, determine whether there is another nested radical of the same value, with lower nesting depth.

Open Question 5: Find it.

Finally there is a variation a la Borodin, Fagin, Hopcroft and Tompa:

Open Question 6: Given a real nested radical, determine whether there is another nested radical of the same value, of lower nesting depth.

Acknowledgments: I would like to thank Neil Immerman, Hendrik Lenstra, Carl Pomerance, Martin Tompa and Ann Trenk for their help.

References

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

[2] A. Besicovitch, On the Linear Independence of Fractional Powers of Integers, *J. London Math. Soc.*, 15, (1940), pp. 3-6.

[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,* pp. 329-338.

[5] G. Horng and M. Huang, "On Simplifying Nested Radicals and Solving Polynomials by Pure Nested Radicals of Minimum Depth," manuscript.

[6] S. Landau, " Simplification of Nested Radicals," to appear, *SIAM J. of Comput.*

[7] S. Landau, "A Note on 'Zippel Denesting,' " to appear, *J. Symb. Comput.*

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

[9] S. Ramanujan, *Problems and Solutions, Collected Works of S. Ramanujan*, Cambridge University Press, 1927.

[10] R. Zippel, Simplification of Expressions Involving Radicals, *J. Symb. Comp.,* 1 (1985), pp. 189-210.