SIGACT News Complexity Theory Column 101

Lane A. Hemaspaandra
Dept. of Computer Science
University of Rochester
Rochester, NY 14627, USA

Introduction to Complexity Theory Column 101

This Issue

This issue's column brings to a close *the half-year of Bill* in the Complexity Theorem Column. In particular, the previous issue's Complexity Theory Column was Bill Gasarch's third P versus NP poll, and this issue's column is Bill, Erik, Jacob, and Scott's article on the muffin problem. Warmest thanks to Bill, Erik, Jacob, and Scott for their splendid baking! Quite naturally given that their article is rich in conjectures, this article is brought to you not just by the Complexity Theory Column, but in fact is a joint production with the Open Problems Column, the letter M, and the number $\frac{157}{286}$.

Have You Seen... the History of Computing?

Things I wish I had known of earlier:

1. The muffin problem.

2. Two lovely windows into the history of computing.

Bill, Erik, Jacob, and Scott's column made me aware of the first issue and their article provides a lovely invitation to the problem. But let me please also take a moment to point to the abovementioned lovely windows into the history of our field.

The first window into history is the 2018 book, "The Making of a New Science: A Personal Journey Through the Early Years of Theoretical Computer Science," by the terrific theoretical computer scientist Giorgio Ausiello, who is able to recount the history because he was there in the midst of it (and, as I personally know and deeply benefited from, he has had a tremendous influence on so many generations of the field, both through his own work and through his warm, kind, enthusiastic way of finding ways to bring researchers together).

The second window into history (which I did know of previously, though not how extensive it is) is the ACM Oral History Project, which is a series of interviews with many of the people whose work has shaped the field. At least some of these can be found via the ACM Digital Library, e.g., via the web page `https://dl.acm.org/citation.cfm?id=1141880&preflayout=flat#prox`, where you can for example find fascinating interviews with Juris Hartmanis, Michael Rabin, and many others across a broad range of the areas that computer science encompasses. Despite the fact that as of this writing there are no recent interviews listed at that URL and the page shows a 2006 proceedings-like date/entry (although some interviews are from as recently as 2012), I know of at least one interview—that of Dick Stearns—that has been conducted very recently and does not yet appear in that page's Table of Contents. And so with luck, maybe this valuable ACM effort is an ongoing project.

Coming Issues

Please stay tuned for future Complexity Theory Columns here from Emanuele Viola (topic: TBD); from Aviad Rubinstein (topic TBD); and from Sabine Broda, António Machiavelo, Nelma Moreira, and Rogério Reis (tentative title: Average Descriptional Complexity of Regular Languages).

And now please put on your oven mitts, because it is Muffin Time!

Guest Column: The Muffin Problem[1]

***William Gasarch*[2] *Scott Huddleston*[3] *Erik Metz*[4] *Jacob Prinz*[5]**



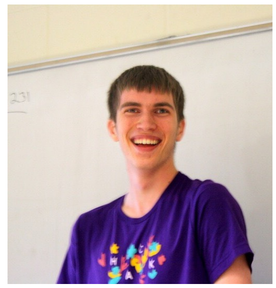


Abstract

You have m muffins and s students. You want to divide the muffins into pieces and give the shares to students such that every student has $\frac{m}{s}$ muffins and the minimum piece is maximized.

Let $f(m, s)$ be the minimum piece in the optimal protocol. We present, by examples, several methods to obtain upper and lower bounds on $f(m, s)$.

1 Introduction and Notation

At *The Twelfth Gathering for Gardner (2016)*, I (William Gasarch) saw a booklet:

Julia Robinson Mathematics Festival:
A Sample of Mathematical Puzzles
Compiled by Nancy Bachman.

On Page 2 was the following:

Begin Excerpt

The Muffin Puzzle
Invented by Recreational Mathematician Alan Frank
Described by Jeremy Copeland in
The New York Times Numberplay Online Blog
`wordplay.blogs.nytimes.com/2013/08/19/cake`

You have 5 muffins and 3 students. You want to divide the muffins evenly, but no student wants a tiny sliver. What division of muffins maximizes the smallest piece? What about 3 muffins and 5 students? 6 muffins and 10 students? 4 muffins and 7 students.

End Excerpt

Bill began working on the general problem of m muffins and s students and yada-yada-yada the following has transpired:

[2]Dept. of Computer Science, Univ. of MD at College Park, College Park, MD 20742, USA. `gasarch@cs.umd.edu`.

[3]13850 Harness Lane, Beaverton, OR 97008, USA. `c.scott.huddleston@gmail.com`.

[4]Dept. of Mathematics, Univ. of MD at College Park, College Park, MD 20742, USA. `emetz1618@gmail.com`.

[5]Dept. of Mathematics, Univ. of MD at College Park, College Park, MD 20742, USA. `jacobeliasprinz@gmail.com`.

1. Bill and a large set of co-authors have an article in the conference **Fun with Algorithms** titled **A Muffin Theorem Generator** [6]on **The Muffin Problem**.

2. Bill Gasarch, Erik Metz, Jacob Prinz, and Daniel Smolyak have a book coming out tentatively titled **The Mathematics of Muffins**.

3. Bill Gasarch, Scott Huddleston, Erik Metz, and Jacob Prinz have a survey in Lane Hemaspaandra's Complexity Column titled **The Muffin Problem**. You are reading it now.

4. Bill gave a talk at the MIT combinatorics seminar on **The Muffin Problem**.

5. At the seminar Bill met (this was pre-arranged) Alan Frank who brought muffins cut $\{\frac{5}{12}, \frac{7}{12}\}$ and was delighted that someone was working so hard on his problem.

Independent of the work of mentioned above, Scott Huddleston worked on the muffin problem and came up with an algorithm that seems to always work.

This article is a survey of some aspects of **The Muffin Problem** including Scott's work. Many details omitted here will be supplied in the book. So let's get started!

You have 5 muffins and 3 students. You want to divide the muffins evenly so each student gets $\frac{5}{3}$ muffins. You *could* do the following:

1. Leave 3 muffins uncut and cut the remaining 2 muffins $\{\frac{1}{3}, \frac{2}{3}\}$.

2. Give Alice and Bob $\{\frac{2}{3}, 1\}$.

3. Give Carol $\{\frac{1}{3}, \frac{1}{3}, 1\}$.

Note that the smallest piece is of size $\frac{1}{3}$. Can we do better? Yes:

Theorem 1.1 *There is a procedure that divides 5 muffins evenly among 3 students such that the smallest piece is $\frac{5}{12}$. There is no procedure that yields a larger smallest piece.*

Proof:
Part One: There is a procedure with smallest piece $\frac{5}{12}$.

The following procedure divides and distributes 5 muffins to 3 people with smallest piece $\frac{5}{12}$.

1. Divide 4 muffins $\{\frac{5}{12}, \frac{7}{12}\}$.

2. Divide 1 muffin $\{\frac{6}{12}, \frac{6}{12}\}$.

3. Give 2 students $\{\frac{6}{12}, \frac{7}{12}, \frac{7}{12}\}$.

4. Give 1 student $\{\frac{5}{12}, \frac{5}{12}, \frac{5}{12}, \frac{5}{12}\}$.

[6]A Muffin Theorem Generator, Ninth International Conference on Fun with Algorithms, 2018. Complete list of authors: Guangiqi Cui and John Dickerson and Naveen Durvasula and William Gasarch and Erik Metz and Jacob Prinz and Naveen Raman and Daniel Smolyak and Sung Hyun Yoo. This can be found at `http://drops.dagstuhl.de/opus/portals/lipics/index.php?semnr=16069`.

Part Two: Every procedure has a piece of size $\geq \frac{5}{12}$.

Assume there is a procedure for dividing up 5 muffins and distributing the shares to 3 students such that every student gets $\frac{5}{3}$ muffins. We show that some piece is $\leq \frac{5}{12}$. We can assume that every muffin is cut because if a muffin is uncut, we will cut it $\{\frac{1}{2}, \frac{1}{2}\}$ and give both halves to the intended recipient.
Case 1: Some muffin is split into ≥ 3 pieces. Then some piece is $\leq \frac{1}{3} < \frac{5}{12}$.
Case 2: All 5 muffins are cut into 2 pieces. Hence there are 10 pieces. Alice gets ≥ 4 shares (if everyone got ≤ 3 shares then there would be $\leq 9 < 10$ shares). One of those shares is $\leq \frac{5}{3} \times \frac{1}{4} = \frac{5}{12}$. ∎

Definition 1.2 Let $m, s \in \mathbb{N}$. An (m, s)-*procedure* is a procedure to cut m muffins into pieces and then distribute them to the s students so that each student gets $\frac{m}{s}$ muffins. An (m, s)-procedure is *optimal* if it maximizes the size of the smallest piece of any procedure. Let $f(m, s)$ be the size of the smallest piece in an optimal (m, s)-procedure. Theorem 1.1 can be restated as $f(5, 3) = \frac{5}{12}$.

It is not obvious that $f(m, s)$ exists or is rational. One can set up a Mixed Integer Program with rational coefficients whose answer is $f(m, s)$. Hence $f(m, s)$ exists, is rational, and is computable. This results in algorithm that runs in time 2^{ms}. We will later conjecture that $f(m, s)$ can be computed in time poly in m, s.

Notation 1.3 Let $m, s \in \mathbb{N}$ and $\alpha \in \mathbb{R}$.

1. $f(m, s) \geq \alpha$ means that there is an (m, s)-procedure with smallest piece $\geq \alpha$.
2. $f(m, s) \leq \alpha$ means that every (m, s)-procedure has a piece $\leq \alpha$.

We leave the following to the reader:

Theorem 1.4 *Let $m, s \in \mathbb{N}$.*

1. *s divides m if and only if $f(m, s) = 1$.*
2. *m divides s if and only if $f(m, s) = \frac{m}{s}$.*
3. *For all $k \in \mathbb{N}$, $f(km, ks) \geq f(m, s)$.*
4. *If s does not divide m then $f(m, s) \leq \frac{1}{2}$.*
5. *$f(m, s) \geq \frac{1}{s}$.*

Note that Theorem 1.4 *does not* imply

For all $k \in \mathbb{N}$, $f(km, ks) = f(m, s)$.

However, we conjecture that this is true:

Conjecture 1.5 *For all m, s, k, $f(km, ks) = f(m, s)$.*

Theorem 1.6 *For all $m \geq 1$ the following hold.*

1. $f(m,1)=1$.

2. *If m is even, then* $f(m,2)=1$.

3. *If m is odd, then* $f(m,2)=\frac{1}{2}$.

We leave the proof of the following theorem to the reader.

Theorem 1.7 *Let $m,s \in \mathbb{N}$. Then $f(s,m)=\frac{s}{m}f(m,s)$.*

Throughout the rest of the article we have the following conventions.

1. We assume $m \geq s$. This suffices to cover all cases by Theorem 1.7.

2. We will assume that m,s are relatively prime. By Theorem 1.4 we have $f(km,ks) \geq f(m,s)$. Every technique we have that finds a bound α such that $f(m,s) \leq \alpha$, also shows that, for all $k \in \mathbb{N}$, $f(km,ks) \leq \alpha$. Hence, our techniques provide evidence for Conjecture 1.5.

3. When we are trying to prove $f(m,s) \geq \alpha > \frac{1}{3}$ then we will assume every muffin is cut into exactly two pieces. We call this **by convention**.

The following definitions will be used throughout the paper

Definition 1.8 Let $m,s \in \mathbb{N}$ and assume there is an (m,s)-procedure.

1. A student who gets V shares is called a *V-student*.

2. If a share is given to a V-student, it is a *V-shares*.

Definition 1.9 Assume there is a procedure where every muffin is cut into exactly 2 pieces. Let x be a piece. *The buddy of x*, denoted $B(x)$, is $1-x$. This definition extends naturally to sets of pieces. We will be using it on intervals of pieces. We write $B(a,b)$ rather than the more proper $B((a,b))$. Similarly we use $B[a,b]$ rather than $B([a,b])$. Note that $B(x)=1-x$.

2 Generating Upper Bounds: The Floor-Ceiling Theorem

The proof of the upper bound in Theorem 1.1 generalizes to what we call *the Floor-Ceiling Theorem*:

Theorem 2.1 *Let $m,s \geq 1$ and s does not divide m.*

$$f(m,s) \leq \max\left\{\frac{1}{3}, \min\left\{\frac{m}{s} \times \frac{1}{\lceil \frac{2m}{s} \rceil}, 1-\frac{m}{s} \times \frac{1}{\lfloor \frac{2m}{s} \rfloor}\right\}\right\}.$$

Proof: Assume there is an (m,s)-procedure. Since s does not divide m every muffin is cut into ≥ 2 pieces.
Case 1: If some muffin is cut into ≥ 3 pieces then there is a piece $\leq \frac{1}{3}$.
Case 2: If all muffins are cut into ≤ 2 pieces, then by convention, all muffins are cut into 2 pieces; therefore, there are $2m$ pieces. Since there are s students both of the following hold:
Case 2a: Alice gets $\geq \lceil \frac{2m}{s} \rceil$ shares, therefore she gets a share $\leq \frac{m}{s} \times \frac{1}{\lceil \frac{2m}{s} \rceil}$.
Case 2b: Bob gets $\leq \lfloor \frac{2m}{s} \rfloor$ shares, therefore he gets a share $\geq \frac{m}{s} \times \frac{1}{\lfloor \frac{2m}{s} \rfloor}$. The buddy of that share is $\leq 1-\frac{m}{s} \times \frac{1}{\lfloor \frac{2m}{s} \rfloor}$. ∎

We define FC (Floor-Ceiling) to encompass both Theorem 2.1 and the trivial case where s divides m. This will make later exposition smoother without special cases.

Notation 2.2 If s divides m then $\mathrm{FC}(m,s) = 1$. Otherwise:

$$\mathrm{FC}(m,s) = \max\left\{\frac{1}{3}, \min\left\{\frac{m}{s} \times \frac{1}{\lceil \frac{2m}{s} \rceil}, 1 - \frac{m}{s} \times \frac{1}{\lfloor \frac{2m}{s} \rfloor}\right\}\right\}.$$

By the Floor-Ceiling Theorem, $f(m,s) \le \mathrm{FC}(m,s)$. We are interested in knowing if, for all m, s, $f(m,s) = \mathrm{FC}(m,s)$. The answer is NO. However (1) for $1 \le s \le 4$, $f(m,s) = \mathrm{FC}(m,s)$, and (2) for fixed s, for all but a finite number of m, $f(m,s) = \mathrm{FC}(m,s)$. We will not prove either of these.

3 Generating Upper Bounds: The Half Method

The smallest s for which $f(m,s) \ne \mathrm{FC}(m,s)$ is $s = 5$. The smallest m in the $s = 5$ case is $m = 11$. We show $f(11,5) = \frac{13}{30} < \mathrm{FC}(11,5) = \frac{11}{25}$.

Theorem 3.1 $f(11,5) = \frac{13}{30}$.

Proof: We show $f(11,5) \le \frac{13}{30}$. We leave the proof that $f(11,5) \ge \frac{13}{30}$ to the reader. In Section 6 we discuss how one might find the procedure systematically.

We derive the upper bound during the proof. We will denote it α.

Assume there is an $(11,5)$-procedure with smallest piece α. By convention every muffin is cut into exactly 2 pieces. Hence there are 22 pieces. Note that there can be at most 11 pieces $> \frac{1}{2}$. This will be a key to getting a contradiction.

We make a leap here and assume there is some V such that every students is either a V-student or a $(V-1)$-student. (see Definition 1.8). With some trial and error we find that $V = 5$ works (we define *works* soon). We now proceed with the proof.

Case 1: If Alice gets ≥ 6 shares then some share is

$$\le \frac{11}{5 \times 6} = \frac{11}{30} \le \alpha.$$

We will need $\frac{11}{30} \le \alpha$.

Case 2: If Alice gets ≤ 3 shares then some share is

$$\ge \frac{11}{5 \times 3} = \frac{11}{15}$$

so its buddy is

$$\le 1 - \frac{11}{15} = \frac{4}{15} \le \alpha.$$

We will need $\frac{4}{15} \le \alpha$.

Case 3: Everyone is either a 4-student or a 5-student.

Let s_4 (s_5) be the number of 4-students (5-students). Since every muffin is cut into 2 pieces there are $11 \times 2 = 22$ pieces.

We have

$$\begin{aligned} 4s_4 + 5s_5 &= 22 \\ s_4 + s_5 &= 5. \end{aligned}$$

Hence $s_4 = 3$ and $s_5 = 2$. So there are twelve 4-shares and ten 5-shares. (This is why $V = 5$ *works*—it leads to the above equations having a natural number solution. No other choice of V does.) Since there are 11 muffins, each cut in half, there are at most 11 pieces $> \frac{1}{2}$. In particular not all 12 of the 4-shares can be $> \frac{1}{2}$. We will *derive* what α needs to be to ensure that all the 4-shares are $> \frac{1}{2}$. This will be our contradiction.

We want α such that there are no 4-shares $\le \frac{1}{2}$. Assume, by way of contradiction, that there is a 4-share $\le \frac{1}{2}$. The remaining 3 shares add up to $\ge \frac{11}{5} - \frac{1}{2} = \frac{17}{10}$; hence one of the other shares is $\ge \frac{17}{10} \times \frac{1}{3} = \frac{17}{30}$. The buddy of that share is $\le 1 - \frac{17}{30} = \frac{13}{30}$.

We will need $\frac{13}{30} \le \alpha$.

Putting it all together we need

$$\alpha = \max\left\{\frac{11}{30}, \frac{8}{30}, \frac{13}{30}\right\} = \frac{13}{30}$$

to get a contradiction. We have just proved $f(11,5) \le \frac{13}{30}$. ∎

In the proof that $f(11,5) \le \frac{13}{30}$ we guessed that everyone was either a 4-student or a 5-student and this turned out to be true. The following seems to be true empirically:

Conjecture 3.2 *(The V-Conjecture) Let $m \ge s$. Let $V = \lceil \frac{2m}{s} \rceil$. If $f(m,s) < \mathrm{FC}(m,s)$ then in every optimal (m,s)-procedure everyone is either a V-student or a $(V-1)$-student. (If we only insist that some optimal procedure has this property than we think we can drop the $f(m,s) < \mathrm{FC}(m,s)$ condition.)*

Note 3.3 We show that the $f(m,s) < \mathrm{FC}(m,s)$ condition is needed. $f(15,8) \le \frac{3}{8}$ by the Floor-Ceiling Theorem. Note that $V = \lceil \frac{30}{8} \rceil = 4$. $f(15,8) \ge \frac{3}{8}$ by the following two procedures:

Procedure One: This procedure has 3-students and 5-students. Hence it is a counterexample to the V-condition without the condition $f(m,s) < \mathrm{FC}(m,s)$.

1. Divide 15 muffins $\{\frac{3}{8}, \frac{5}{8}\}$.
2. Give 3 students $\{\frac{5}{8}, \frac{5}{8}, \frac{5}{8}\}$.
3. Give 5 students $\{\frac{3}{8}, \frac{3}{8}, \frac{3}{8}, \frac{3}{8}, \frac{3}{8}\}$.

Procedure Two: This procedure has 3-students and 4-students. Hence $f(15,8)$ is not a counterexample to the weaker V-conjecture where we only require that *some* procedure use only V-students and $(V-1)$-students.

1. Divide 6 muffins $\{\frac{3}{8}, \frac{5}{8}\}$.
2. Divide 9 muffins $\{\frac{4}{8}, \frac{4}{8}\}$.
3. Give 2 students $\{\frac{5}{8}, \frac{5}{8}, \frac{5}{8}\}$.

4. Give 6 students $\{\frac{3}{8}, \frac{4}{8}, \frac{4}{8}, \frac{4}{8}\}$.

One can (roughly) take the technique used to get $f(11,5) \le \frac{13}{30}$ and make it into a program that will, on input (m,s), output HALF(m,s) which is the bound you get from applying the method. Hence we have

$$f(m,s) \le \min\{\mathrm{FC}(m,s), \mathrm{HALF}(m,s)\}.$$

Is this an equality? No. But for $1 \le s \le 8$ and $m \ge s$ equality holds.

4 Generating Upper Bounds: Other Methods

There are four methods to *generate* an upper bound on $f(m,s)$ which we will not describe: Interval (INT), Midpoint(MID), Easy Buddy Match (EBM) and Hard Buddy Match (HBM). INT and MID are similar to, but harder than, HALF. There are two methods to *verify* an upper bound on $f(m,s)$: GAPS and TRAIN (we will describe GAP in the next section, though not TRAIN). Both are similar to though harder than HALT, INT, MID.

The Buddy-Match methods only work when there is a 2-student and uses the fact that if such a student has a share of size x, the other share is of size $\frac{m}{s} - x$.

When all is said and done we have:

$$f(m,s) \le \min\{\mathrm{FC}(m,s), \mathrm{HALF}(m,s), \mathrm{INT}(m,s), \mathrm{MID}(m,s), \mathrm{EBM}(m,s), \mathrm{HBM}(m,s)\}.$$

Is this an equality? No. But we have empirical evidence that, for $1 \le s \le 18$ and $m \ge s$, equality holds.

5 Verifying Upper Bounds: The Gap and Train Methods

For all of the prior methods discussed we end up with a function that, on input m, s, outputs an upper bound. The GAPS method has to be *given* α and then *verifies that* $f(m,s) \le \alpha$ *or fails to do so.* In Section 7 we will discuss how to use the method within a program to find $f(m,s)$.

We do one example of the GAPS method.

Theorem 5.1 $f(31,19) = \frac{54}{133}$.

Proof: We show $f(31,19) \le \frac{54}{133}$. We leave the proof that $f(31,19) \ge \frac{54}{133}$ to the reader. In Section 6 we discuss how one might find the procedure systematically.

We express everything with denominator 266. Note that $\frac{31}{19} = \frac{434}{266}$.

Assume there is a $(31,19)$-procedure with smallest piece $> \frac{108}{266}$. Then the biggest piece is of size $< 1 - \frac{108}{266} = \frac{158}{266}$. We leave it to the reader to show that there are fourteen 3-students, five 4-students, forty-two 3-shares, twenty 4-shares.

Assume Alice is a 3-student with shares $x \le y \le z$. If $x \le \frac{118}{266}$ then $y + z \ge \frac{434}{266} - \frac{118}{266} = \frac{316}{266}$. Hence $z \ge \frac{316}{266} \times \frac{1}{2} = \frac{158}{266}$. Hence all 3-shares are $> \frac{118}{266}$. By similar reasoning all 4-shares are $< \frac{110}{266}$. The following picture captures what we know:

$$\underset{\frac{108}{266}}{(} \quad \text{20 4-shs} \quad \underset{\frac{110}{266}}{)[} \quad 0 \quad \underset{\frac{118}{266}}{](} \quad \text{42 3-shs} \quad \underset{\frac{158}{266}}{)}$$

(here and below, "shs" is shorthand for "shares").

There are no shares in $(\frac{110}{266}, \frac{118}{266})$. Hence, by buddying, there are no shares in $(\frac{148}{266}, \frac{156}{266})$. The following picture captures what we know:

$$\begin{array}{ccccccccccc} (& \text{20 4-shs} &)[& 0 &](& \text{22 3-shs} &)[& 0 &](& \text{20 3-shs} &) \\ \frac{108}{266} & & \frac{110}{266} & & \frac{118}{266} & & \frac{148}{266} & & \frac{156}{266} & & \frac{158}{266} \end{array}.$$

There are the same number of shares in $(\frac{118}{266}, \frac{133}{266})$ as in $(\frac{133}{266}, \frac{148}{266})$ by buddying. There are an even number of shares of size $\frac{133}{266}$. We will think of them as being divided between the two intervals. Hence there are 11 shares in each of the two intervals.

The following picture captures what we know about the 3-shares:

$$\begin{array}{ccccccccc} (& \text{11 3-shs} & | & \text{11 3-shs} &)[& 0 &](& \text{20 3-shs} &) \\ \frac{118}{266} & & \frac{133}{266} & & \frac{148}{266} & & \frac{156}{266} & & \frac{158}{266} \end{array}.$$

We define the following intervals.

1. $I_1 = (\frac{118}{266}, \frac{133}{266})$.
2. $I_2 = (\frac{133}{266}, \frac{148}{266})$ $(|I_1| = |I_2| = 11)$.
3. $I_3 = (\frac{156}{266}, \frac{158}{266})$ $(|I_3| = 20)$.

There may be shares of size $\frac{133}{266}$ but this will not affect our proof.

We need a finer classification of 3-students. We need to know how many shares from I_1, I_2, and I_3 they have.

Notation 5.2

1. If $1 \le i \le 3$ then *an I_i-share* is a share from I_i.
2. Let $1 \le j_1 \le j_2 \le j_3 \le 3$. A (j_1, j_2, j_3)-*student* is a student who has an I_{j_1}-share, an I_{j_2}-share, and an I_{j_3}-share. The j's could be equal.
3. y_{j_1,j_2,j_3} is the number of students of type (j_1, j_2, j_3).
4. If a proof that a student is impossible is exact then we put a * on it. This will be clearer when we do it.

Claim 1:

1. The following are the only students who are allowed.
   (a) $(1, 2, 3)$.
   (b) $(1, 3, 3)$.
   (c) $(2, 2, 2)$.
   (d) $(2, 2, 3)$.

2. There are no shares in $[\frac{122}{266}, \frac{128}{266}]$.

3. There are no shares in $[\frac{138}{266}, \frac{144}{266}]$ (this follows from the last item and buddying).

Proof of Claim 1:
1) We establish that some students are impossible.
A $(1,2,2)$-student has less than $\frac{133}{266} + 2 \times \frac{148}{266} = \frac{429}{266} < \frac{434}{266}$.
A $(1,1,3)$-student has less than $2 \times \frac{133}{266} + \frac{158}{266} = \frac{424}{266} < \frac{434}{266}$.
A $(2,3,3)$-student has more than $\frac{133}{266} + 2 \times \frac{156}{266} = \frac{445}{266} > \frac{434}{266}$.
The result follows.
2) How big does an I_1-share have to be?
A $(1,2,3)$-student has I_1-share $> \frac{434}{266} - \frac{306}{266} = \frac{128}{266}$.
A $(1,3,3)$-student has I_1-share $< \frac{434}{266} - 2 \times \frac{156}{266} = \frac{122}{266}$.
These are the only students who use I_1-shares. The result follows.
End of Proof of Claim 1

We define new intervals.

1. $I_1 = (\frac{118}{266}, \frac{122}{266})$.

2. $I_2 = (\frac{128}{266}, \frac{133}{266})$.

3. $I_3 = (\frac{133}{266}, \frac{138}{266})$ ($|I_2| = |I_3|$).

4. $I_4 = (\frac{144}{266}, \frac{148}{266})$ ($|I_1| = |I_4|$).

5. $I_5 = (\frac{156}{266}, \frac{158}{266})$ ($|I_5| = 20$).

(We also know that $|I_1| + |I_2| = |I_3| + |I_4| = 11$ but this is not needed.)
Claim 2: The following are the only students who are allowed.

1. $(1,5,5)$.

2. $(2,4,5)$.

3. $(3,4,5)$.

4. $(4,4,4)$.

Proof of Claim 2:

We establish that some students are impossible. (Regarding the two "*" symbols below, please recall the final part of Notation 5.2.)
A $(1,4,5)$-student has less than $\frac{122}{266} + \frac{148}{266} + \frac{158}{266} = \frac{428}{266}$.
A $(3,4,4)$-student has less than $\frac{138}{266} + 2 \times \frac{148}{266} = \frac{434}{266}$.*
A $(3,3,5)$-student has less than $2 \times \frac{138}{266} + \frac{158}{266} = \frac{434}{266}$.*
A $(2,5,5)$-student has more than $\frac{128}{266} + 2 \times \frac{156}{266} = \frac{440}{266} > \frac{434}{266}$.
A $(4,4,5)$-student has more than $2 \times \frac{144}{266} + \frac{156}{266} = \frac{444}{266} > \frac{434}{266}$.
The result follows.
End of Proof of Claim 2

Let x be the number of $(1,5,5)$-students. Note that since $|I_2| = |I_3|$, there are an equal number of $(2,4,5)$-students and $(3,4,5)$-students. Call that number y. Finally, let z be the number of $(4,4,4)$-students.

Since $|I_1| = |I_4|$,

$$x = 2y + 3z.$$

Since there are $s_3 = 14$ students,

$$x + 2y + z = 14.$$

If you substitute $x = 2y + 3z$ into the equation $x + 2y + z = 14$, you get

$$y + z = \frac{7}{2}$$

which is a contradiction since $y, z \in \mathbb{N}$. ▮

There is an extension of the GAPS method called the TRAIN method. We omit it.

Let GAPS(m, s, α) be computed as follows: try to apply the GAPS method to show $f(m, s) \leq \alpha$. If it succeeds output YES, otherwise NO. Similar for TRAIN(m, s, α).

6 Verifying Lower Bounds: Linear Algebra

In this section we derive procedures systematically. From what we present the reader can easily devise an algorithm that will *given* m, s, α either find a procedure for $f(m, s) \geq \alpha$ or fail to do so. Failure does not mean that there is no such procedure.

We will be assuming the following conjecture:

Conjecture 6.1 *Let $m \geq s$. Let $a, b \in \mathbb{N}$. Let $d = gcd(s, b)$. If $f(m, s) \geq \frac{a}{b}$ then there is an (m, s)-procedure where all of the shares are multiples of $\frac{1}{d}$.*

We will need the following notation.

Definition 6.2 A *multiset* is a set where elements may be repeated. A *multisubset* of a set B is a multiset where every elements is in B. Note that if an elements appears n times in B then it can appear $\leq n$ times in a multisubset of B.

Definition 6.3 Let B be a set of L numbers, and let S be a multisubset of B. The *vector representation* of S is a vector $(n_1, n_2, \ldots, n_L)$ where n_i is the number of times that the i^{th} element of B appears in S (when B is ordered least to greatest).

Definition 6.4 The *sum of a multiset* simply means the sum of the elements in the multiset.

We now derive the procedure for $f(5, 3) \geq \frac{5}{12}$. We do it in steps.

1) By Conjecture 6.1 all of the pieces are in the set A:

$$A = \left\{ \frac{5}{12}, \frac{6}{12}, \frac{7}{12} \right\}.$$

Each muffin will be split into two piece sizes from A. We can represent this with a multiset of piece sizes. For example, if a muffin is split into two halves, then we may represent it by the multiset $\{\frac{6}{12}, \frac{6}{12}\}$ or with the vector representation $(0, 2, 0)$.

In order to avoid cluttering the page with denominators, we will multiply all piece sizes by 12. So instead, let

$$B = \{5, 6, 7\}.$$

If we want to know all ways to cut a muffin, rather than say *We need all multiset subsets of A that sum to 1* we say *We need all multiset subsets of B that sum to 12.*

2) Find all vectors that correspond to how a muffin can be cut. We need all subsets of B that sum to 12 (since $1 = \frac{12}{12}$).

It is easy to see that the sets are

$\{6, 6\}$ which is $(0, 2, 0)$. Let m_1 be the number of muffins cut this way.

$\{5, 7\}$ which is $(1, 0, 1)$. Let m_2 be the number of muffins cut this way.

3) We can also use multisets to represent the pieces which a student receives. Each student receives $\frac{5}{3}$ muffin total, so each of these multisets will sum to 20 (since $\frac{5}{3} = \frac{20}{12}$).

It is easy to see that the sets are

- $\{6, 7, 7\}$ which is $(0, 1, 2)$. Let s_1 be the number of students who get these shares.
- $\{5, 5, 5, 5\}$ which is $(4, 0, 0)$. Let s_2 be the number of students who get these shares.

4) Set up equations to find m_1, m_2, s_1, s_2.

The number of each piece size that the muffins give is equal to the number of each piece size that the students receives. Therefore, we get the equation:

$$m_1(0, 2, 0) + m_2(1, 0, 1) = s_1(0, 1, 2) + s_2(4, 0, 0).$$

This equation implies:

$$\begin{aligned} m_2 &= 4s_2 \\ 2m_1 &= s_1 \\ m_2 &= 2s_1. \end{aligned}$$

Since there are 5 muffins and 3 students we have:

$$\begin{aligned} m_1 + m_2 &= 5 \\ s_1 + s_2 &= 3. \end{aligned}$$

The 5 equations have one solution in the naturals: $m_1 = 1$, $m_2 = 4$, $s_1 = 2$, $s_2 = 1$. It is easy to take this solutions and make a procedure out of it.

The above technique can be used on any (m, s, α). Let FINDPROC(m, s, α) be the algorithm that uses the technique above to try to prove $f(m, s) \geq \alpha$. If it succeeds then output YES, otherwise output NO.

7 An Algorithm that Tries to Find $\mathbf{f(m,s)}$ But Probably Fails Sometimes

We give an algorithm tries to find $f(m,s)$. Does it always succeed? This is an open problem. We comment on it at the end of this section.

1. Input(m,s). If $s=1$ then output 1.
2. If $m<s$ then run $f(s,m)$ and output $\frac{m}{s}f(s,m)$.
3. If s divides m output(1) and halt.
4. If m,s are not relatively prime then let $d=gcd(m,s)$ and run the program on $(m/d,s/d)$. Once you get the answer α see what techniques were used and see if they apply to (m,s). They will. Output α.
5. (We can assume $m>s$, m,s relatively prime, $s\ge 2$.) Compute the min of the following and call it α:

$$\{\mathrm{FC}(m,s),\mathrm{HALF}(m,s),\mathrm{INT}(m,s),\mathrm{EBM}(m,s),\mathrm{MID}(m,s),\mathrm{HBM}(m,s)\}.$$

6. Run FINDPROC(m,s,α).
7. If the answer is YES then output α.
8. If the answer is NO then let $\alpha=\frac{a}{b}$. Compute FINDPROC$(m,s,\frac{a-1}{b})$, FINDPROC$(m,s,\frac{a-2}{b})$, ... until you get a YES. (In practice this has always been $\frac{a-1}{b}$.) We know

$$\frac{a-i}{b}\le f(m,s)\le \frac{a}{b}.$$

9. Look at all the numbers between $\frac{a-i}{b}$ and $\frac{a}{b}$ that have denominators of the form bx where $1\le x\le C$ where C is a constant or function of b chosen by the programmer (we only used denominators ≤ 600, however one can replace this with some increasing function of b).

   Sort them to produce

$$\alpha_1<\alpha_2<\cdots<\alpha_n.$$

   By using FINDPROC(m,s,α_i) and GAPS(m,s,α_i) do a binary search to find an α_i that is both an upper and lower bound. (You might also use TRAIN(m,s,α_i).)

This method might not be that fast since FINDPROC may be slow. However, in practice it is usually fast. Does it always work?

Conjecture 7.1 *The above algorithm always produces $f(m,s)$.*

We don't really believe this conjecture. At various times we thought (1) FC would always give the right answer, (2) the min of FC and HALF would always give the right answer,..., (7) the min of FC, HALF, INT, MID, EBM, HBM, GAP (using GAP together with FP) always gave the right answer, and then it did not work for $f(67, 21)$. We then came up with TRAIN which handled that and some other cases. Will this process ever end? Will we find (perhaps with Scott's algorithm from Section 8) some other m, s such that none of FC, HALF, INT, MID, EBM, HBM, GAP, or TRAIN give the value of $f(m, s)$? We do not know. So this is less a conjecture and more of a question. Sadly, even if this algorithm does work, it is not very fast. Formally we do not think this algorithm works in time polynomial in m, s.

8 An Algorithm that Tries to Find $f(m, s)$ and Probably Never Fails

Scott Huddleston has devised a method that does the following: on input m, s, generates an α and an (m, s)-procedure that shows $f(m, s) \geq \alpha$. His algorithm has three amazing properties:

- Recall that the algorithm in Section 7 will *verify* that $f(m, s) \geq \alpha$. Scott's algorithm will, given m, s, *find* an α such that $f(m, s) \geq \alpha$. This is not impressive in and of itself but the next two points make it impressive.

- Scott's algorithm is extremely fast. The algorithm in Section 7 currently can verify $f(m, s) \geq \alpha$ (given an α we suspect is an answer) for most m, s with $m, s \leq 100$, in about an hour. With some fine tuning and a faster language we suspect we can get it up to $m, s \leq 200$ in about an hour. Scott's algorithm can compute a lower bound for $f(m, s)$ for all m, s with $m + s \leq 100,000$ in a few minutes.

- The α Scott gets seems to always be $f(m, s)$. We believe that Scott's algorithm *does* compute $f(m, s)$ (not just a lower bound for it) though we have not been able to prove this.

In this exposition we omit much including the case where $f(m, s) = \frac{1}{3}$.

We sketch two examples of his method. We assume the V-conjecture throughout. Hence we will always assume that the everyone is either a $(V - 1)$-student or a V-student where $V = \lceil \frac{2m}{s} \rceil$.

8.1 Five Muffins, Three Students

We assume $f(5, 3) > \frac{1}{3}$. By convention every muffin is cut into 2 pieces, so there are 10 shares. We leave it to the reader to show that there are two 3-students, one 4-student, six 3-shares, and four 4-shares. So far this is all standard.

We introduce several new ideas that we use throughout this section.

New Idea One: Generalize The Problem

We first restate the (5,3)-problem:

1. We have 5 muffins that are worth 1 each and cut into 2 pieces. We denote this as $(5, 1, 2)$. In the future we will have muffins that have values other than 1.

2. We have one 4-student who needs $\frac{5}{3}$ via 4 shares. We denote this as $(1, \frac{5}{3}, 4)$.

3. We have two 3-students who need $\frac{5}{3}$ via 3 shares. We denote this as $(2, \frac{5}{3}, 3)$.

We denote this problem as

$$scott\left[(5,1,2), \left(1, \frac{5}{3}, 4\right), \left(2, \frac{5}{3}, 3\right)\right].$$

We call it SC(5,3)-0. We will soon recast it as a problem about finding weights on edges in a graph. We will still call this recast problem SC(5,3)-0.

This is an example of a *Scott-Muffin Problem*. We now give the formal definition and conventions.

Definition 8.1 A *Scott-Muffin Problem* is a 3-tuple of 3-tuples:

$$\begin{aligned}&(n_m, v_m, p_m)\\&(n_{s_1}, v_{s_1}, p_{s_1})\\&(n_{s_2}, v_{s_2}, p_{s_2}),\end{aligned} \tag{1}$$

with the following meaning.

1. All three tuples are in $\mathbb{N} \times \mathbb{Q} \times \mathbb{N}$. All 9 numbers are ≥ 0.

2. The first tuple (n_m, v_m, p_m) means that there are n_m muffins, each with value v_m, and each cut into p_m pieces. Later p_m will be the degree of a muffin-vertex in a graph. All three of these numbers are > 0.

3. Both the second and third tuples represent students.

   (a) The second tuple $(n_{s_1}, v_{s_1}, p_{s_1})$ means that there are n_{s_1} students (these are not all of the students) who want muffins of value v_{s_1}, in p_{s_1} shares. These are called *the major students* (we'll see why in point c). Later p_{s_1} will be the degree of a student-vertex in a graph. All three of these numbers are > 0.

   (b) The third tuple $(n_{s_2}, v_{s_2}, p_{s_2})$ means there are n_{s_2} students who want muffins of value v_{s_2}, in p_{s_2} shares. These are called *the minor students* (we'll see why in point c). Later p_{s_2} will be the degree of a student-vertex in a graph. If all three numbers are 0 then we leave it off and in this case the Scott-Muffin problem only has two tuples—a muffin tuple and the major students.

   (c) Which student-tuple is major and which is minor is determined as follows: the tuple with the larger ratio $\frac{\texttt{degree}}{\texttt{value}}$ is the major muffins. In other words, $\frac{p_{s_1}}{v_{s_1}} > \frac{p_{s_2}}{v_{s_2}}$.

4. Be forewarned: you are used to thinking of pieces of muffins being given to students. We will often invert that and think of students giving pieces to muffins. The graphs we use will be undirected so either mentality is fine.

The Scott-Muffin problem is important since we will be taking a standard muffin problem and transforming it into smaller Scott-Muffin problem, and then (possibly) again into an even smaller Scott-Muffin Problem, until we get to a certain type of Scott-Muffin problem that is easy to solve optimally. We will then use that to solve all the problems in the sequence (conjecturally) optimally.

So solving Scott-Muffin problems is an example of that trope: *it's sometimes easier to solve a harder problem.*

New Idea Two: Represent the Problem as a Graph

Since the 4-student only uses 4 shares and there are 5 muffins, there must be a muffin that is shared *among only the 3-students*. Since each muffin is cut in two pieces, there will be two 3-students who share a muffin. We represent this in Graph 1 where the massive magenta[7] vertices are muffins and the small cyan vertices are students.

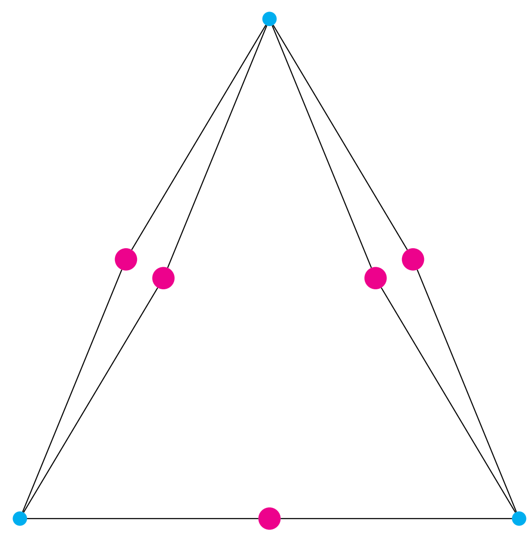

Graph 1: Five Muffins, Three Students, SC(5,3)-0

We will present many more graphs (actually multigraphs) where (1) vertices are either students or muffins, and (2) a muffin vertex is connected to a student vertex if that student gets a piece of that muffin. We state the conventions for such graphs.

Convention 8.2 In all of our graphs, the following hold.

1. Muffins are Massive Magenta (reddish) colored dots (M for Muffin, Massive, and Magenta).
2. Students are Small Cyan (blueish) colored dots (S for Student, Small, and (sort of) Cyan).
3. A muffin and a student are connected if a student has a piece of that muffin. Since muffins can only be connected to students and vice versa, students and muffins are the two parts of a bipartite graph. We do not draw the graphs as bipartite since that would be a mess.
4. Given a Scott-Muffin problem

$$\begin{array}{l}(n_m, v_m, p_m)\\(n_{s_1}, v_{s_1}, p_{s_1})\\(n_{s_2}, v_{s_2}, p_{s_2}),\end{array} \tag{2}$$

[7]Depending on the medium you are reading this in you may or may not see the colors.

we will associate a graph. This graph is *not unique*. That is, there may be more than one graph that represents the problem. This will end up not mattering since the graphs are visual aids and not used in the actual algorithm. The graph will have n_m muffin-vertices of degree p_m, n_{s_1} student-vertices of degree p_{s_1}, and n_{s_2} student-vertices of degree p_{s_2} (we leave out for now how to determine the edges). The problem is to assign nonnegative weights to the edges such that every muffin-vertex has weighted degree v_m, that every major-student-vertex has weighted degree v_{s_1}, and every minor-student-vertex has weighted degree v_{s_2}. It is easy to see how these weights can be used to obtain a solution to the Scott-Muffin problem.

5. Note that (1) all of the muffin-nodes are of degree p_m, (2) all of the major-student-nodes are of degree p_{s_1}, and (3) all of the minor-student-nodes are of degree p_{s_2}.

6. Note that the graph itself does not specify the entire Scott-Muffin problem. We often say things like *this graph captures some of the Scott-Muffin Problem.*

7. The muffin-vertices for a standard muffin problem will have degree 2 since each muffin is cut into exactly 2 pieces. For a Scott-Muffin problem where the muffins may have values other than 1 and may be cut into more than 2 pieces, the muffin-vertices may have higher degree.

8. The weights on the edges represent the size of the piece that the muffin gave to the student. To re-iterate: we will often invert that and think of a student giving pieces to a muffin.

Since there are two 3-students who share a muffin, and one 4-student, Graph 1 captures some of what we know.

Note that a (5,3)-procedure is a way to assign nonnegative weights to the edges of Graph 1 such that

- The weighted degree of each muffin vertex is 1.
- The weighted degree of each student vertex is $\frac{5}{3}$.

We call the problem of finding such weights SC(5,3)-0.

New Idea Three: Transform the Problem into a Smaller One—Clusters are Students

We need a notation for a certain part of the graph.

Definition 8.3 Let $L \geq 0$. An *L-cluster* is a sequence of length $2L+1$ of the form student-muffin-$\cdots$-student that has L minor students, together with all the other muffins attached to the students. The muffins in the student-muffin-$\cdots$-student sequence are called *internal muffins* whereas the muffins that are not in that sequence but are attached to the students are called *external muffins*. The muffins in the sequence might have other students attached to them but those students are not part of the cluster.

Graph 1 has a 1-cluster consisting of the 2 students and 1 internal muffin (at the bottom) *together with* the four external muffins that are adjacent to the students. Here is the big new idea:

We will transform the problem by regarding this 1-cluster as being a student.

The internal and external muffins are part of the cluster. The 2 students in the cluster need $2 \times \frac{5}{3} = \frac{10}{3}$. There is 1 internal muffin and there are 4 external muffins so the cluster has 5 muffins. We are now going to view the students as having muffin pieces to give to the muffins. Hence the cluster can be viewed as *a student who has an excess of* $5 - \frac{10}{3} = \frac{5}{3}$ (the fact that this is $\frac{5}{3}$ is an accident, do not let that confuse you). The degree of the cluster is 4. Hence we can view the 1-cluster as a student of value $\frac{5}{3}$ and degree 4.

What to make of the remaining student? We have already used up all of the muffins, so that student can be viewed as needing $\frac{5}{3}$ but not having any muffins. So its value is $-\frac{5}{3}$. Rather than think of negative numbers we instead think of this student as being *a muffin who needs* $\frac{5}{3}$. Note also that this vertex (which now represents a muffin) has degree 4. Hence we can view this as the following Scott-Muffin Problem:

$$scott\left[\left(1, \frac{5}{3}, 4\right), \left(1, \frac{5}{3}, 4\right)\right].$$

We call it SC(5,3)-1 and it is represented by Graph 2. The problem is now to put nonnegative weights on the edges such that

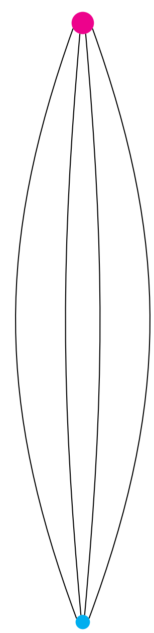

Graph 2: Five Muffins, Three Students, SC(5,3)-1

- The weighted degree of the muffin vertex is $\frac{5}{3}$.
- The weighted degree of each student vertex is $\frac{5}{3}$.

We now state a conjecture with two parts. One part we use now, one we use later.

Conjecture 8.4 *If either (1) there are no minor students, or (2) no set of clusters contains all of the minor students, then all the pieces being given to the major muffins will be the same size.*

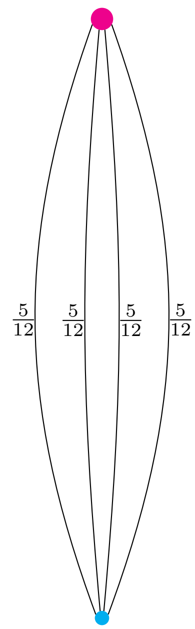

Graph 3: Five Muffins, Three Students, Solution to SC(5,3)-1

Using this conjecture, and the fact that there are no minor students, the problem is now easy: Assign each edge $\frac{5}{12}$ as in Graph 3.

How do we go from the solution of SC(5,3)-1 to a solution of SC(5,3)-0? The bottom node is really a cluster of two student-vertices and an internal muffin-vertex. Recall that the weights were how much these students were going to give away. Consider one of these students, Alice. Alice is connected to 2 external muffin-nodes. Using the solution to SC(5,3)-1 (Graph 3) we see that, for each of these muffins, she gives away $\frac{5}{12}$ and hence keeps $\frac{7}{12}$. Hence each student keeps $2 \times \frac{7}{12} = \frac{7}{6}$. They now need to split the internal muffin so that each one gets $\frac{5}{3}$. Hence they each need $\frac{5}{3} - \frac{7}{6} = \frac{1}{2}$. Wow! We give the solution to SC(5,3)-0 in Graph 4.

In this case obtaining the solution to the SC(5,3)-0 from the SC(5,3)-1 was easy. It is not always so easy and we do not always split internal muffins $(\frac{1}{2}, \frac{1}{2})$.

In summary we transformed SC(5,3)-0

$$\mathit{scott}\left[(5,1,2)\,,\left(1,\frac{5}{3},4\right),\left(2,\frac{5}{3},3\right)\right]$$

into the easier problem SC(5,3)-1

$$\mathit{scott}\left[\left(1,\frac{5}{3},4\right),\left(1,\frac{5}{3},4\right)\right].$$

We then solved SC(5,3)-1 and used its solution to solve SC(5,3)-0.

8.2 Thirty-Five Muffins, Thirteen Students

We now do the problem of $f(35,13)$. We will use the ideas from Section 8.1; hence we will assume familiarity with the definitions and ideas presented there. We will need a few new ideas as well.

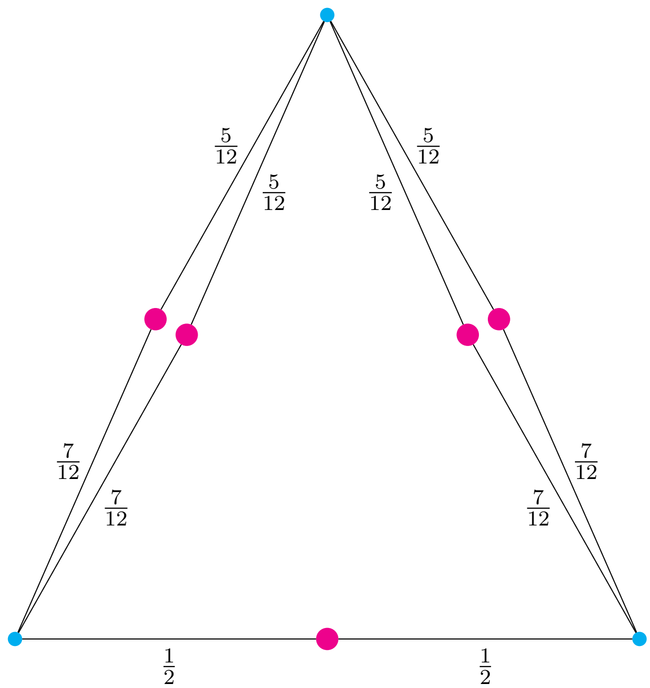

Graph 4: Five Muffins, Three Students, Solution to SC(5,3)-0

We assume $f(35,13) > \frac{1}{3}$. By convention every muffin is cut into 2 pieces, so there are 70 pieces. We leave it to the reader to show that there are eight 5-students, five 6-students, forty 5-shares, and thirty 6-shares. We express this as the following Scott-Muffin problem:

$$scott\left[\,(35,1,2)\,,\left(5,\frac{35}{13},6\right),\left(8,\frac{35}{13},5\right)\right].$$

We call this problem SC(35,13)-0. It is partially represented by Graph 6, though we need to explain why that graph represents the problem, so we will look at Graph 5 first.

Since there are thirty 6-shares, the 6-students can only use pieces of 30 muffins. Hence there are 5 muffins that are used entirely by the 5-students.

We will assume that the 5 muffins that are shared by the 5-students form one 1-cluster and two 2-clusters. Graph 5 shows those clusters (without the external muffins—that would make a mess) along with the 6-students. It turns out that essentially every muffin problem begins this way: (1) find V so that everyone is either a V-student or a $(V-1)$-student ($V = \lceil \frac{2m}{s} \rceil$), (2) find that the $(V-1)$-students must share m' muffins between them, (3) find an L such that the muffin-sharing can be represented by clusters of length L and $L-1$.

Given the Scott-Muffin problem and the clusters, Graph 6 represents it (other graphs might also).

As in Section 8.1 we will transform SC(35,13)-0 into a smaller problem. Look at Graph 6.

1. The five 6-students will be viewed as not having any muffins adjacent to them (these are the external muffins of the clusters in the next two items) hence these five 6-students need $\frac{35}{13}$ each and have nothing to begin with. These are now viewed as muffins and denoted $(5, \frac{35}{13}, 6)$.

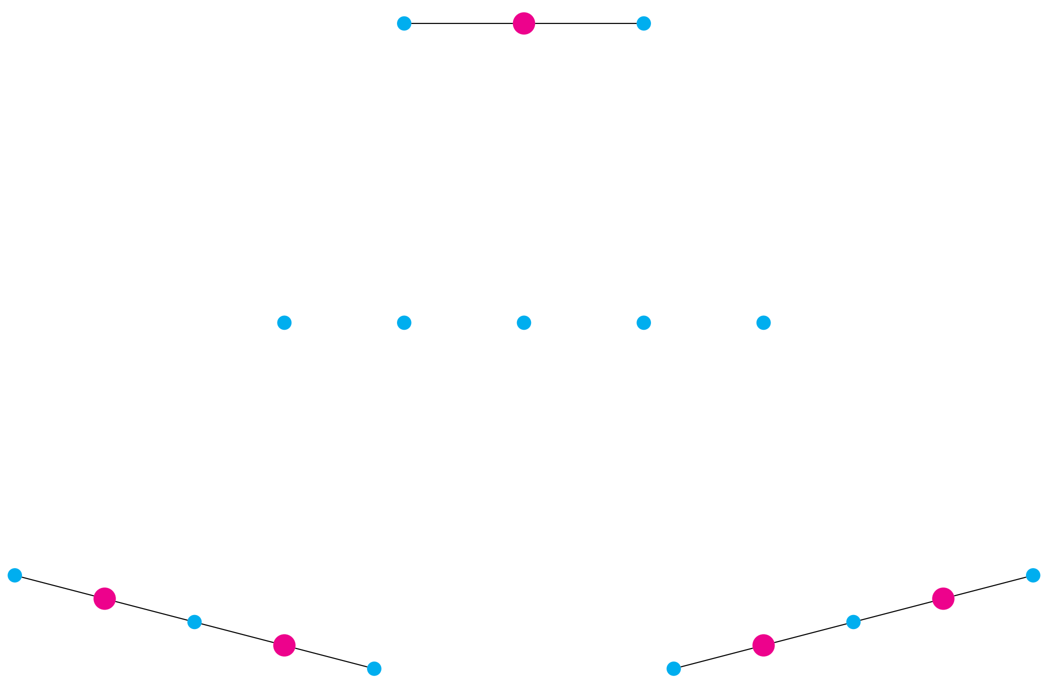

Graph 5: Thirty-Five Muffins, Thirteen Students: Clusters

2. There are two 2-clusters of 5-students (they are at the bottom of both Graph 5 and 6). We focus on one of them; however, the same goes for the other one. The three students need $3 \times \frac{35}{13} = \frac{105}{13}$ muffins. The cluster has 2 internal muffins and 11 external muffins for a total of 13 muffins. Hence the cluster becomes a student of value $13 - \frac{105}{13} = \frac{64}{13}$. Note that there are 11 edges coming out of the cluster. Since there are 2 of these clusters we denote this $(2, \frac{64}{13}, 11)$. Note that $\frac{\texttt{degree}}{\texttt{value}} = \frac{11}{64/13} \sim 2.23$. This ratio of degree to value is larger than the one in the next item, so these are the major students.

3. There is one 1-cluster of 5-students (it is at the top of both Graph 5 and 6). The two students need $2 \times \frac{35}{13} = \frac{70}{13}$ muffins. The cluster has 1 internal muffin and 8 external muffins for a total of 9 muffins. Hence the cluster becomes a student of value $9 - \frac{70}{13} = \frac{47}{13}$. Note that there are 8 edges coming out of the cluster. Since there is only 1 of these clusters we denote this $(1, \frac{47}{13}, 8)$. Note that $\frac{\texttt{degree}}{\texttt{value}} = \frac{8}{47/13} \sim 2.21$. This ratio of degree to value is smaller than the one in the prior item, so these are the minor students.

Hence we have the following Scott-Muffin problem:

$$scott\left[\left(5, \frac{35}{13}, 6\right), \left(2, \frac{64}{13}, 11\right), \left(1, \frac{47}{13}, 8\right)\right].$$

We call this problem SC(35,13)-1. It is partially captured by Graph 7.
While we have spilled a lot of ink, all we've done so far is transformed SC(35,13)-0:

$$scott\left[(35, 1, 2), \left(5, \frac{35}{13}, 6\right), \left(8, \frac{35}{13}, 5\right)\right]$$

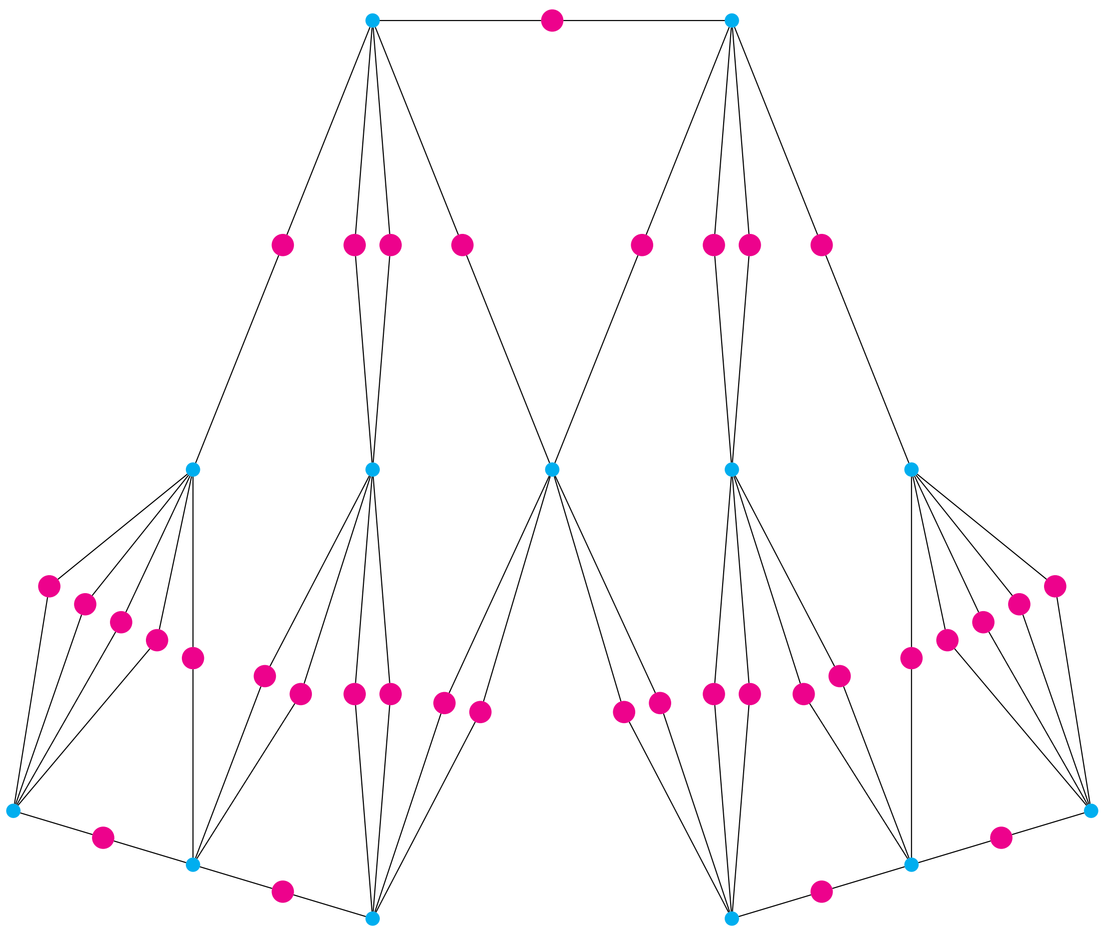

Graph 6: Thirty Five Muffins and Thirteen Students, SC(35,13)-0

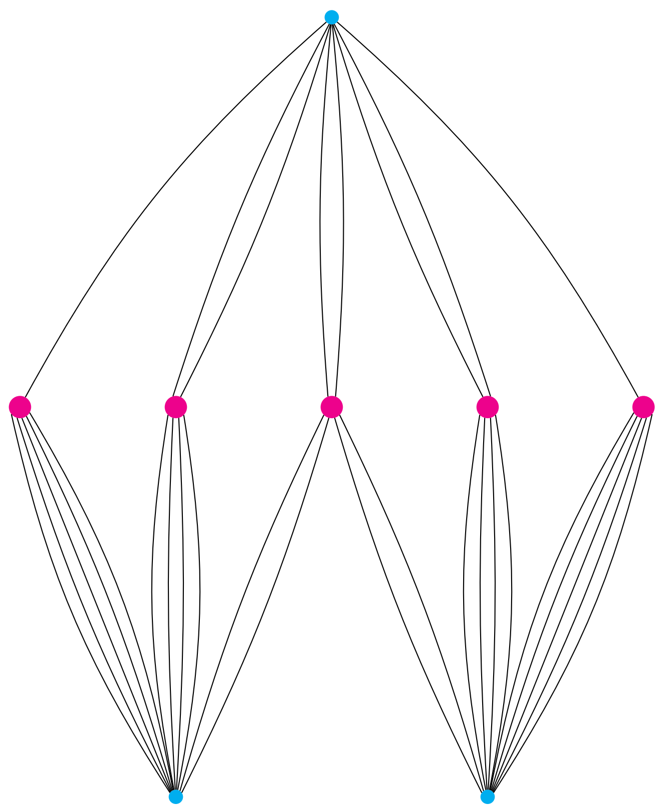

Graph 7: Thirty-Five Muffins, Thirteen Students, SC(35,13)-1

into SC(35,13)-1:

$$scott\left[\left(5,\frac{35}{13},6\right),\left(2,\frac{64}{13},11\right),\left(1,\frac{47}{13},8\right)\right].$$

We will transform SC(35,13)-1 to a new problem SC(35,13)-2. We will then find a solution to SC(35,13)-2 and use it to find and a solution to SC(35,13)-1, and use that to find a solution to SC(35,13)-0.

8.2.1 Transforming SC(35,13)-1 to SC(35,13)-2

Recall SC(35,13)-1:

$$scott\left[\left(5,\frac{35}{13},6\right),\left(2,\frac{64}{13},11\right),\left(1,\frac{47}{13},8\right)\right],$$

which is represented by Graph 7. Since there are no clusters, by Conjecture 8.4 we give all the major students equal amounts on all of their edges.

The two major students (the two student-vertices at the bottom of Graph 7) each want weighted degree $\frac{64}{13}$ and are of unweighted degree 11. Hence we give each of the edges coming out of it weight $\frac{64}{13}\times\frac{1}{11}=\frac{64}{143}$.

With these edges taken care of we will recurse into a smaller Scott-Muffin SC(35,13)-2. Before defining SC(35,13)-2 we look at the muffin vertices of SC(35,13)-1 that have gotten some of the way towards their weighted degree.

There are two kinds of muffin vertices in Graph 7. Note that the muffin vertices are the ones in the middle layer. SC(35,13)-1:

- The 3 muffin-vertices that have 2 edges to the 1 minor student. Since these muffins originally needed weighted degree $\frac{35}{13}$ and now have, from the edges to the major students, $4\times\frac{64}{143}=\frac{256}{143}$, they now need just $\frac{35}{13}-\frac{256}{143}=\frac{129}{143}$.
- The 2 muffin-vertices that have 1 edge to the 1 minor student. Since these muffins originally needed weighted degree $\frac{35}{13}$ and now have, from the edges to the major students, $5\times\frac{64}{143}=\frac{320}{143}$, they now need just $\frac{35}{13}-\frac{320}{143}=\frac{5}{11}$.

We will now define the SC(35,13)-2 problem.

1. There is 1 muffin of value $\frac{47}{13}$ and degree 8. We denote this as $(1,\frac{47}{13},8)$. (This used to be the 1 minor student, which is the top most student in Graph 7.)
2. There are 3 students of value $\frac{129}{143}$ and degree 2. We denote this as $(3,\frac{129}{143},2)$. Note that $\frac{\texttt{degree}}{\texttt{value}}=\frac{2}{129/143}\sim 2.22$. These students have a larger $\frac{\texttt{degree}}{\texttt{value}}$ than those in the next item so these are the major students. (These used to be the muffins that had 2 edges to the minor student.)
3. There are 2 students of value $\frac{5}{11}$ and degree 1. We denote this as $(2,\frac{5}{11},1)$. Note that $\frac{\texttt{degree}}{\texttt{value}}=\frac{1}{5/11}\sim 2.20$. These students have a smaller $\frac{\texttt{degree}}{\texttt{value}}$ than those in the prior item so these are the minor students. (These used to be the muffins that had 1 edges to the minor student.)

Hence we have

$$scott\left[\left(1,\frac{47}{13},8\right),\left(3,\frac{129}{143},2\right),\left(2,\frac{5}{11},1\right)\right].$$

We call this problem SC(35,13)-2. It is partially captured by Graph 8.

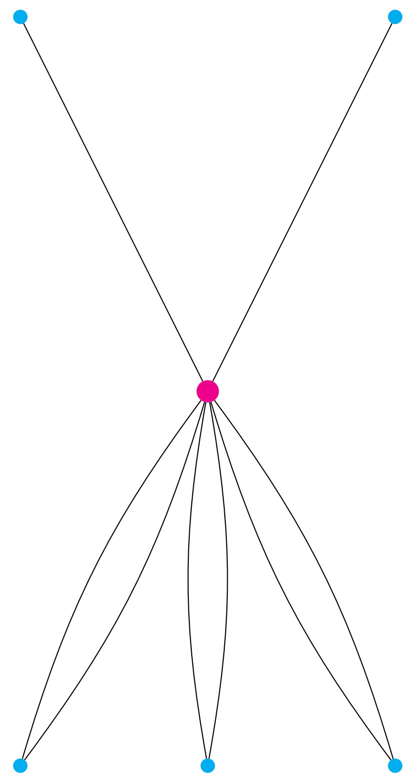

Graph 8: Thirty Five Muffins, Thirteen Students, SC(35,13)-2

This graph has no clusters.

We use Conjecture 8.4 and give all the major students equal amounts on all of their edges. Since each major student has degree 2 and must get $\frac{129}{143}$ all of those edges get weight $\frac{129}{143}\times\frac{1}{2}=\frac{129}{286}$. The minor students need $\frac{5}{11}$ and are of degree 1 so the edge to each one must be $\frac{5}{11}$. This completes the solution, though we need to check that the muffins worked out (the first triple in SC(35,13)-2). For the solution see Graph 9.

We need to check that the first part of SC(35,13)-2 works: $(1,\frac{47}{13},8)$. There are 2 edges going into the muffin node of weight $\frac{5}{11}$ and 6 of weight $\frac{129}{286}$. Hence the total weight going into the muffin is

$$2\times\frac{5}{11}+6\times\frac{129}{286}=\frac{47}{13}.$$

Do not be surprised at this. The way we set it up it had to be that way.

8.2.2 Using the Solution to SC(35,13)-2 to Solve SC(35,13)-1

We use the solution to SC(35,13)-2 as expressed in Graph 9 to solve SC(35,13)-1. Actually, this is quite easy, since when we went from SC(35,13)-1 to SC(35,13)-2, we had already assigned weights

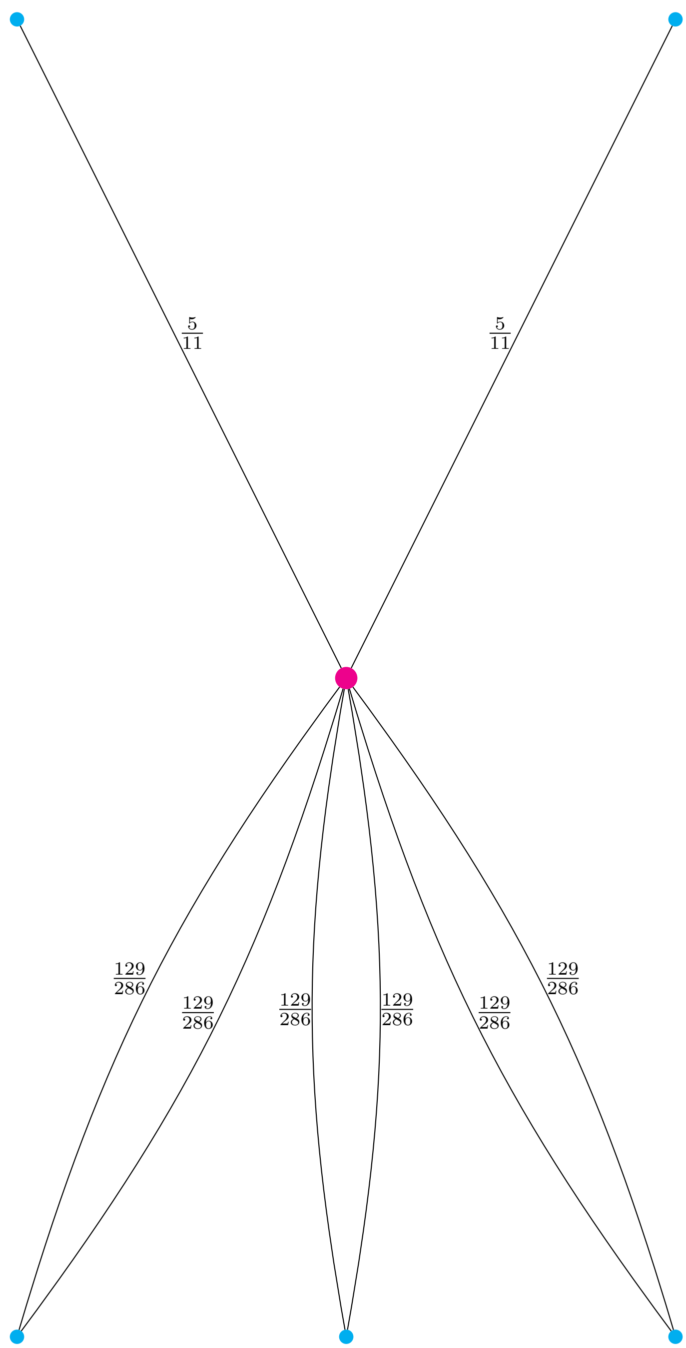

Graph 9: Thirty Five Muffins, Thirteen Students, Solution to SC(35,13)-2

to edges and then removed them. Now all we need to do is put them back. Graph 10 shows a solution to SC(35,13)-1.

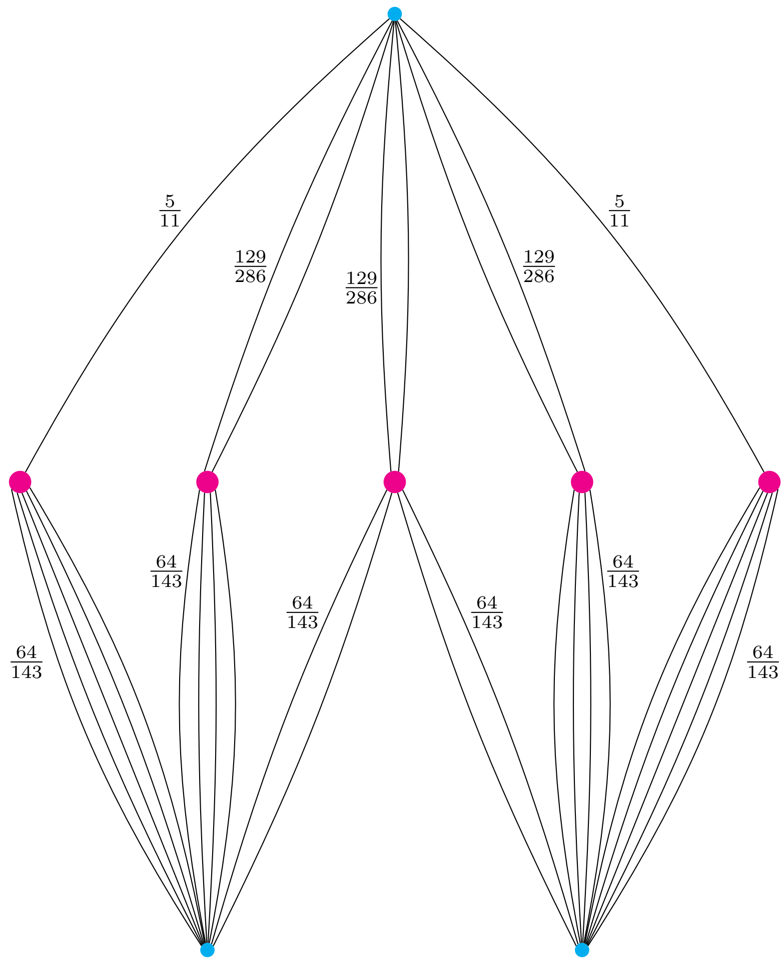

Graph 10: Thirty Five Muffins, Thirteen Students, Solution to SC(35,13)-1

8.2.3 Using the Solution of SC(35,13)-1 to Solve SC(35,13)-0

We explain how to take a solution to SC(35,13)-1 and use it to obtain a solution to SC(35,13)-0. This will be a case where clusters become vertices and some thought is needed to convert the solution. We will be asking you to flip back and fourth between (1) the problem SC(35,13)-0 which is Graph 6, (2) the solution to SC(35,13)-1, which is Graph 10, and (3) the solution to SC(35,13)–, which is Graph 11.

The left bottom student-vertex in Graph 10 corresponds to the left bottom cluster of Graph 6. The 11 edges coming out of the left bottom student-vertex in Graph 10 correspond to the 11 2-paths (student-muffin-student) coming out of the left bottom cluster of Graph 6. We think of the students in the left bottom cluster as *giving away* $\frac{64}{143}$ and *keeping* $1 - \frac{64}{143}$. Hence we get part of Graph 11.

Look at the left most student-vertex in the cluster (the same will hold for the right most) who we call Alice. Alice keeps for herself $4 \times \frac{79}{143} = \frac{316}{143}$ and needs (which will come from the internal muffin) $\frac{35}{13} - \frac{316}{143} = \frac{69}{143}$. Hence the muffin between left and middle is split $\frac{69}{143}$ for left and $1 - \frac{69}{143} = \frac{74}{143}$ for middle.

The same happens for the right most student in the cluster. So now the left and right both have weighted degree $\frac{35}{13}$. What about the middle? He has

$$3 \times \frac{79}{143} + 2 \times \frac{74}{143} = \frac{35}{13}.$$

This should not surprise you. We set it up this way.

The rest of the edges are mostly forced. Look at the 5 student-vertices on the third level from the bottom in Graph 11 (don't look at the edge from those nodes going up). Look at the left most student vertex. We already know that the 5 edges coming into it from the bottom contribute $5 \times \frac{64}{143} = \frac{320}{143}$. Since the node needs weighted degree $\frac{35}{13}$, the edge coming out of it going upwards must have weight

$$\frac{35}{13} - \frac{320}{143} = \frac{5}{11}.$$

This edge of weight $\frac{5}{11}$ goes into a muffin-vertex. Since this muffin vertex has weighted degree the other edge coming out of it has weight

$$1 - \frac{5}{11} = \frac{6}{11}.$$

In a similar manner we can find the weights of all of the edges. It will all work out. See Graph 11 for the full solution.

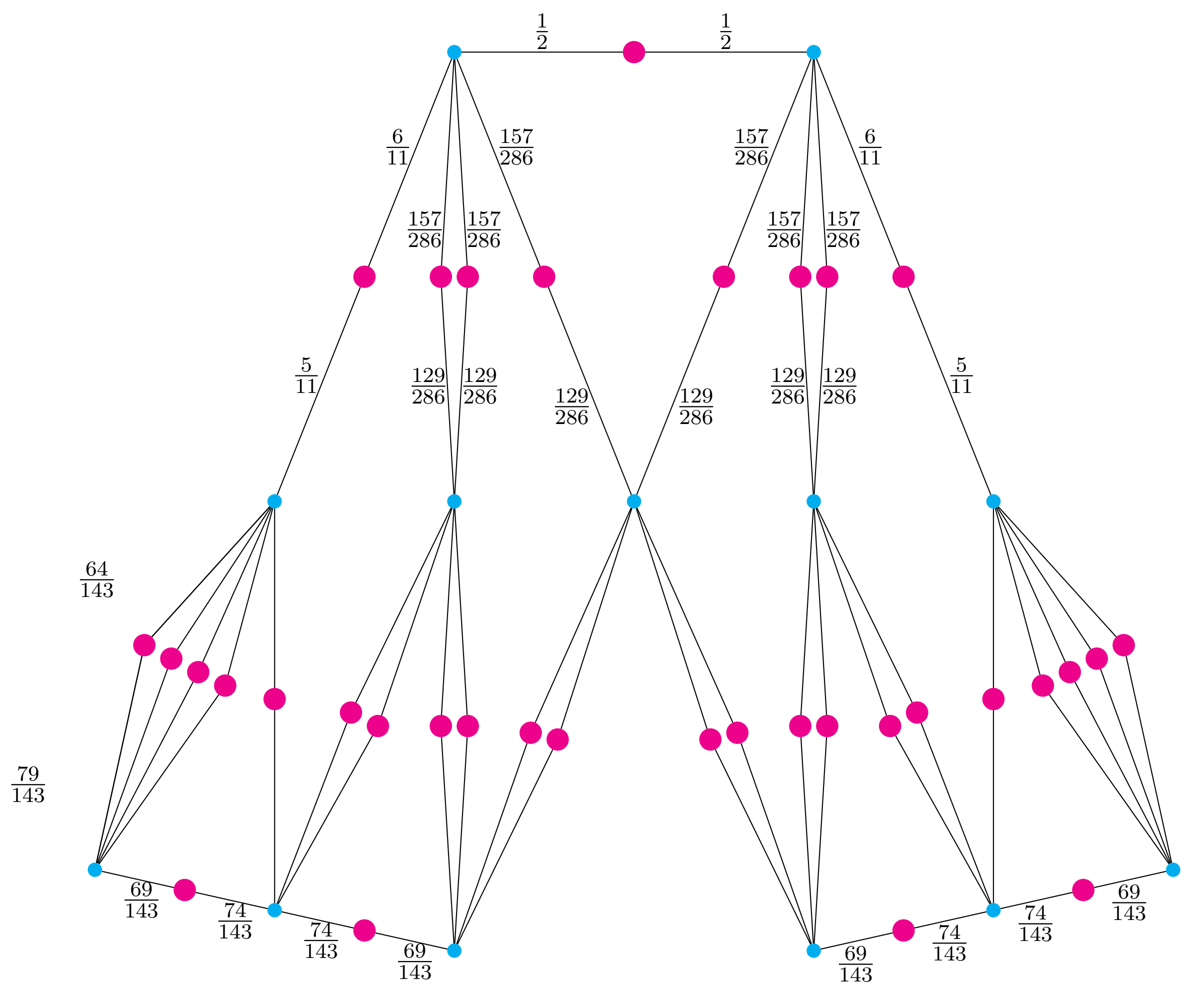

Graph 11: Thirty Five Muffins, Thirteen Students, Solution to SC(35,13)-0

8.3 Reflections on What We Have Done

We have demonstrated a way to, given Scott-Muffin A, find a smaller Scott-Muffin problem B, such that a solution to B gives a solution to A. We do not know that an optimal solution to B gives an optimal solution to A but we believe this to be true. This has held for every single case we have tried. It will be one of our conjectures.

For the Problem of finding a (35,13)-procedure we did the following:

Recast it as finding a solution to SC(35,13)-0:

$$\mathit{scott}\left[(35,1,2),\left(5,\frac{35}{13},6\right),\left(8,\frac{35}{13},5\right)\right].$$

Reduced SC(35,13)-0 to SC(35,13)-1:

$$\mathit{scott}\left[\left(5,\frac{35}{13},6\right),\left(2,\frac{64}{13},11\right),\left(1,\frac{47}{13},8\right)\right]$$

Reduced SC(35,13)-1 to SC(35,13)-2:

$$\mathit{scott}\left[\left(1,\frac{47}{13},8\right),\left(3,\frac{129}{143},2\right),\left(2,\frac{5}{11},1\right)\right].$$

This last problem, SC(35,13)-2, is easy (if it wasn't we would have found a SC(35,13)-3). We got the optimal solution to SC(35,13)-2 and used it to get a solution to SC(35,13)-1. We took this solution to SC(35,13)-1 and used it to get a solution to SC(35,13)-0, our original problem. By conjecture we have the optimal solution for SC(35,13)-0.

Carrying out the reduction of Scott problems to smaller scott problems is very fast. Using a solution for Scott problem B to get a solution for Scott problem A is also very fast. Note that the reduction in size is often large as well so there are not that many iterations.

9 Open Problems

We restate our conjectures and add one more.

Conjecture 9.1 *For all m, s, k, $f(km, ks) = f(m, s)$. We* **REALLY** *believe this one as does* **EVERYONE** *who has ever tasted muffin mathematics. Seems hard to prove.*

Conjecture 9.2 *Let $m \geq s$. Let $V = \lceil \frac{2m}{s} \rceil$. Assume $f(m,s) < \mathrm{FC}(m,s)$. There is an optimal (m,s)-procedure in which everyone is either a V-student or a $(V-1)$-student. We strongly believe that this is true, but at an earlier time we thought it was true without the condition $f(m,s) < \mathrm{FC}(m,s)$, so we could get fooled again. (See Note 3.3.)*

Conjecture 9.3 *Let $m \geq s$. Let $a, b \in \mathbb{N}$. Let $d = gcd(s,b)$. If $f(m,s) \geq \frac{a}{b}$ then there is an (m,s)-procedure where all of the shares are multiples of $\frac{1}{d}$. We* **really** *believe this one. Seems hard to prove.*

Conjecture 9.4 *The algorithm in Section 7 always produces $f(m,s)$. As discussed in Section 7 we don't really believe this one.*

We strongly believe that Scott's algorithm (1) always outputs the correct answer and (2) is fast, both in practice (recall that it computed a lower bound on $f(m,s)$ for all m,s with $m+s \le 100,000$ in a few minutes) and in theory (polynomial in m,s).

Normally in problems with numbers as inputs (like factoring) we think of the input n as being of length $\lg n$. So we should be asking if there is an algorithm for $f(m,s)$ that runs in time poly in $\lg m, \lg s$. If we want to actually output the (m,s)-procedure then this is impossible since a procedure itself takes roughly $O(m+s)$ to describe.

There is a way to modify Scott's algorithm so that it only outputs the lower bound on $f(m,s)$ and not the (m,s)-procedure. It is plausible that this runs in time polynomial in $\lg s, \lg m$. However, while we think it is fast, we do not think it is that fast *in the worst case*. In the average case (this would need to be defined carefully) we think it is polynomial in $\lg s, \lg m$. We summarize:

Conjecture 9.5

1. *Scott's algorithm outputs an optimal (m,s)-procedure and does so in time poly in m,s.*
2. *Modified-Scott computes $f(m,s)$ in time polynomial in $\lg m, \lg s$, on average (this needs to be defined rigorously).*

Conjecture 9.5 implies Conjectures 9.1, 9.2, and 9.3.

10 Acknowledgments

We thank Nathan Grammel for all the pictures and graphs. In particular, the graphs in Section 8 greatly improved the exposition. We would also like to thank Doug Chen and Lane Hemaspaandra for extensive proofreading and discussion, which also greatly improved the exposition. And of course we thank Lane Hemaspaandra for allowing us to present our work in his column.