

Making Change and Finding Repfigits: Balancing a Knapsack

Daniel Lichtblau

Wolfram Research, Inc.
100 Trade Center Dr.
Champaign, IL 61820 USA
danl@wolfram.com

Abstract. We will discuss knapsack problems that arise in certain computational number theory settings. A common theme is that the search space for the standard real relaxation is large; in a sense this translates to a poor choice of variables. Lattice reduction methods have been developed in the past few years to improve handling of such problems. We show explicitly how they may be applied to computation of Frobenius instances, Keith numbers (also called “repfigits”), and as a first step in computation of Frobenius numbers.

Keywords: Frobenius instance solving, lattice reduction, integer linear programming, change-making problem, Frobenius numbers, Keith numbers, repfigits.

1 Introduction

Various problems in the realm of computational number theory have as key steps the solving of a linear equation or system over the integers, subject to some linear inequality constraints. The Frobenius instance problem (also known as the change-making problem) is a well-known example. The problem of finding what are called repfigits, to be described below, is another. These may be regarded as a class of knapsack problems wherein one is allowed to take only certain integer multiples of various items in forming a “valid” combination.. As such these fall into the category of integer linear programming (ILP).

Classical methods for solving such problems include branch-and-bound and cutting plane methods [8,16]. These approaches alone are often inadequate for certain classes of problems due to a phenomenon roughly described as “unbalanced bases”. A method for dealing with this deficiency was developed in [1,2]. In essence it involves working with a basis that is reduced (in the sense of [11]) and ordered by size, in conjunction with standard branch-and-bound. We will describe and illustrate this for the types of problem mentioned above. We remark that the handling of Frobenius instances is by no means new, having been discussed in the aforementioned references. We also show how the method is applied in a new algorithm for finding Frobenius numbers, a task that is substantially harder than solving Frobenius instances.

The algorithms described in this paper have been implemented in *Mathematica* [18]. Code is available from the author on request.

2 Frobenius Instances

Suppose we have a set of positive integers $A = (a_1, \dots, a_n)$ and a target M , a positive integer. We seek nonnegative integer multipliers $X = (x_1, \dots, x_n)$ such that $X \cdot A = M$. This is a standard problem in integer linear programming. A classical method [16] for solving this would be to solve the relaxed problem wherein we enforce all inequality constraints but all variables to be real rather than integer valued. If in the solution we encounter a variable a with value s that is not an integer then we spawn two subproblems where we enforce respectively that $a \leq \lfloor s \rfloor$ and $a \geq \lceil s \rceil$. We continue this process of solving relaxed subproblems, splitting when a variable has a noninteger value. It can be shown that eventually either we exhaust all possibilities or we obtain an integer valued solution [16]; in either case clearly the algorithm terminates.

A drawback to this approach is that the search space for the relaxed subproblems might appear to be “large”, in the sense of having many points with not all coordinates integer valued. In particular it may be the case that a standard LP solver will find real valued solutions to the restricted subproblems without making rapid progress to an (entirely) integer valued solution, because it might be possible to subdivide the (real valued) solution polytope in such a way that integer points do not readily appear at corners.

The method of [1,2] was developed as a way to improve on this situation. Roughly it proceeds as follows. First we find a description of a solution set to our equations that is a priori integer valued but possibly does not satisfy the required inequalities. We arrange that this solution set has “good” basis vectors, such that when we solve relaxed problems with these we more rapidly walk through our (real valued) solution polytope. More correctly, the polytope is likely to intersect fewer hyperplanes orthogonal to larger direction vectors and spaced by integral multiples of those vectors.

With respect to the Frobenius instance problem it goes as follows. First we find a solution vector X over \mathbb{Z}^n and a basis V for the integer null space (that is, $n-1$ independent vectors $V_j \in \mathbb{Z}^n$ with $V_j \cdot A = 0$; for this we use a method based on the Hermite normal form [5]). We use multiples of the basis vectors to find a “small” solution which we still call X . The tactic utilized to do this is sometimes called the embedding method. It apparently has been independently discovered several times; variants appear in [1,12,13,14]. In starting with a solution over \mathbb{Z}^n rather than the nonnegatives \mathbb{N}^n we are working with what is called an integer relaxation of the nonnegativity constraint. We will later enforce nonnegativity via the more common LP relaxations wherein integrality is not enforced.

We define variables $t = (t_1, \dots, t_{n-1})$ so that solution vectors are given by $X + tV$. For purposes of finding a valid solution to the problem at hand we need to impose two requirements. The first is that all components are nonnegative and the second is that all are integers. The first can be met by standard linear

programming. For the second, as in the classical approach, we will use branching on subproblems. Specifically we find solutions to relaxed LP problems wherein we now work over nonnegative reals. We then branch on noninteger values in those solutions. For example, if a solution has, say, $t_j = 5/4$, we create two subproblems identical to the one we just solved, but with the new constraints $t_j \leq \lfloor 5/4 \rfloor = 1$ and $t_j \geq \lceil 5/4 \rceil = 2$, respectively (though note that if other variables also had noninteger solutions then we need not have chosen this particular one). As observed above, this branching process will terminate eventually, with either a valid solution or the information that no such solution exists.

We explain again, in slightly different terms, why it is important to work with a small integer solution to the integer relaxation (that is, allowing negative values) and a lattice reduced basis for the null vectors. As our methodology is to take combinations of these null vectors, effectively they define directions in a solution polytope for the transformed problem. By working with a reduced basis we in effect change our coordinate system to one where the various search directions are roughly orthogonal. This helps us to avoid the possibility of taking many steps in similar directions in searching the polytope of nonnegative solutions for one that is integer valued. Thus we explore it far more efficiently. Moreover, in starting with a small solution we begin closer to the nonnegative orthant. Heuristically this seems to make the sought-for multipliers of the null vectors relatively small, and this is good for computational speed. This is discussed in section 2 of [1], with further explanation and illustrations found in [3].

A further efficiency, from [2], is to choose carefully the variable on which to branch. We order by increasing size the reduced lattice of null vectors. Branching will be done on the noninteger multiplier variable corresponding to the largest of these basis vectors. This has the effect of exploring the solution polytope in directions in which it is relatively thin, thus more quickly finding integer lattice points therein or exhausting the space. This refinement is important for handling pathological examples of the sort presented in [1] and [2].

As we have a constraint satisfaction problem we are also free to impose any linear integer objective function of our choosing. Thus an optimization is to obtain extremal values for linear forms with integer coefficients and use these results as simple cutting planes [16]. For example we can minimize or maximize the various coordinate values t_j , selecting one either in some specific order or at random for each subproblem. This process amounts to finding the width of the polytope along the directions of our lattice basis vectors, and enforcing integrality of the optimized variable helps to further restrict the search space.

As reported in [2] this method is very effective in solving Frobenius instances. We illustrate with an example from [17]

$$A = (10000000000, 10451674296, 18543816066, 27129592681, 27275963647, \\ 29754323979, 31437595145, 34219677075, 36727009883, 43226644830, \\ 47122613303, 57481379652, 73514433751, 74355454078, 78522678316, \\ 86905143028, 89114826334, 91314621669, 92498011383, 93095723941)$$

We let the target be 862323776. In several seconds the instance solver returns the empty set. This has implications for bounding the Frobenius number of A ; we will discuss that in a later section.

3 Keith Numbers

Keith numbers, also known as repfigits, were introduced in 1987 by Michael Keith [10] as a sort of computational novelty that relates a Fibonacci-like sequence to a linear equation involving its seed. They are defined as follows. Suppose we are given a number s of n digits (we work in base 10, but these can be defined with respect to arbitrary bases). We may form a sequence in Fibonacci style as follows. The first n elements are the digits themselves. The $(n + 1)^{\text{th}}$ element is the sum of the first n digits. Subsequent elements are the sums of the preceding n elements. Then s is called a Keith number provided it appears in this sequence. As an example, the sequence for 197 is $\{1, 9, 7, 17, 33, 57, 107, 197, \dots\}$ and so 197 is a Keith number. Keith originally referred to these as repfigits, for “replicating Fibonacci digits”.

Keith numbers tend to be quite rare (there are only 71 of them below 10^{19}). Prior methods for finding them involved clever segmentation of an enumeration. While flawless (in the sense that they find all of them), these are limited in range due to algorithmic complexity and memory requirements. At the time the present work was begun the state of the art, from [10], was that all such numbers up to 19 digits had been found but no larger ones were known. We will take this substantially further.

To begin we must find equations to describe these things. If the digits are $\{d_0, d_1, \dots, d_{n-1}\}$ then the number in question is $\sum_{j=0}^{n-1} d_j 10^{n-1-j}$. We form the appropriate sequence using a Fibonacci matrix of dimension n . This is simply a matrix that, when operating on a vector, replaces each element up to the last by its successor, and replaces the last by the sum of the elements. For example,

for $n = 3$ it is $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$.

If we multiply this matrix by itself $k - 1$ times then the dot product of the bottom row with the digit sequence will give the $(n + k)^{\text{th}}$ term in the sequence. Some simple inequality considerations will give fairly tight bounds on how many such multiples can possibly work for a given number of digits n . We will use each possibility to form a homogeneous linear diophantine equation.

We demonstrate with a short example. We start by obtaining the set of candidate equation vectors for 5 digit examples. One of them is $(-8207, 1705, 3069, 3395, 3524)$. That is, we will seek a solution to the system

$$-8207x_1 + 1705x_2 + 3069x_3 + 3395x_4 + 3524x_5 = 0$$

with each $x_j \in 0, 1, \dots, 9$ and $x_1 \geq 1$.

As in the last section, the first step in the process of [1] and [2] is to find a full set of integer solutions to such a system. Since these are homogeneous

equations we require only the integer null space. This can be obtained readily from the Hermite normal form for the matrix comprised of the vector for the homogeneous equation, augmented by an identity matrix. Again we want to work with vectors that are small and close to orthogonal so we apply lattice reduction to get a “good” set of vectors spanning the same solution set. We obtain $(-3, -1, -3, -3, -1)$, $(-2, -4, -3, 3, -3)$, $(1, 6, -5, 6, -2)$, and $(7, -3, -15, 5, 26)$.

Notice that for any solution vector, its negative is also a solution vector. Looking at the first vector in our solution basis we thus see that 31331 is a Keith number of five digits. That was too easy; there is no guarantee we will have a solution vector with all components in the desired range. We look at a slightly larger example to see this. For six digits one candidate equation-defining vector is $(-96160, -4224, 5752, 7144, 7482, 7616)$. This time we get the following small null vectors: $(0, 3, -3, 0, 4, 0)$, $(-1, 1, -3, -2, -4, -4)$, $(0, -6, -3, 1, 0, -2)$, $(0, -1, 1, 5, 0, -6)$, and $(0, 4, -12, 15, -12, 9)$. So we have a generating set of small null vectors none of which have entirely nonnegative or entirely nonpositive values (with the first being nonzero, in order that they give a legitimate six digit number). We now need a way to recombine these so that the first component is positive and the rest are nonnegative.

Again we have something we can tackle with standard branch-and-bound iterations [8,16]. Let $\{v_1, \dots, v_n\}$ be our null space basis (here n is one less than the number of digits). We seek an integer vector of the form $a_1v_1 + \dots + a_nv_n$ such that all components lie in the range $\{0, 1, \dots, 9\}$ with the first component strictly positive. For this we create variables $\{a_1, \dots, a_n\}$ which will ultimately be required to take on integer values. In order to effect this we will solve relaxed linear programming problems with appropriate inequality constraints. As noted earlier, these are simply constraint satisfaction problems, so we can use arbitrary linear objective functions as a means of obtaining integer cuts cheaply. When some variables do not take on integer values in the solution we choose one such on which to branch and spawn a pair of subproblems. Our choice again is from [2]; we branch on that variable of noninteger solution value whose corresponding basis vector is largest.

To summarize, we first find the appropriate sets of integer equations. For each we find spanning sets of solutions that do not in general satisfy the digit inequality constraints. We lattice reduce these. We use ILP methods to find all possible solutions subject to the usual inequality constraints on digits.

We used this method to find all Keith numbers up through 29 digits. We show all repfigits between 20 and 29 digits below.

$\{1, 2, 7, 6, 3, 3, 1, 4, 4, 7, 9, 4, 6, 1, 3, 8, 4, 2, 7, 9\}$
 $\{2, 7, 8, 4, 7, 6, 5, 2, 5, 7, 7, 9, 0, 5, 7, 9, 3, 4, 1, 3\}$
 $\{4, 5, 4, 1, 9, 2, 6, 6, 4, 1, 4, 4, 9, 5, 6, 0, 1, 9, 0, 3\}$
 $\{8, 5, 5, 1, 9, 1, 3, 2, 4, 3, 3, 0, 8, 0, 2, 3, 9, 7, 9, 8, 9\}$
 $\{7, 6, 5, 7, 2, 3, 0, 8, 8, 2, 2, 5, 9, 5, 4, 8, 7, 2, 3, 5, 9, 3\}$
 $\{2, 6, 8, 4, 2, 9, 9, 4, 4, 2, 2, 6, 3, 7, 1, 1, 2, 5, 2, 3, 3, 3, 7\}$
 $\{3, 6, 8, 9, 9, 2, 7, 7, 5, 9, 3, 8, 5, 2, 6, 0, 9, 9, 9, 7, 4, 0, 3\}$
 $\{6, 1, 3, 3, 3, 8, 5, 3, 6, 0, 2, 1, 2, 9, 8, 1, 9, 1, 8, 9, 6, 6, 8\}$

- {2, 2, 9, 1, 4, 6, 4, 1, 3, 1, 3, 6, 5, 8, 5, 5, 5, 8, 4, 6, 1, 2, 2, 7}
- {9, 8, 3, 8, 6, 7, 8, 6, 8, 7, 9, 1, 5, 1, 9, 8, 5, 9, 9, 2, 0, 0, 6, 0, 4}
- {1, 8, 3, 5, 4, 9, 7, 2, 5, 8, 5, 2, 2, 5, 3, 5, 8, 0, 6, 7, 7, 1, 8, 2, 6, 6}
- {1, 9, 8, 7, 6, 2, 3, 4, 9, 2, 6, 4, 5, 7, 2, 8, 8, 5, 1, 1, 9, 4, 7, 9, 4, 5}
- {9, 8, 9, 3, 8, 1, 9, 1, 2, 1, 4, 2, 2, 0, 7, 1, 8, 0, 5, 0, 3, 0, 1, 3, 1, 2}
- {1, 5, 3, 6, 6, 9, 3, 5, 4, 4, 5, 5, 4, 8, 2, 5, 6, 0, 9, 8, 7, 1, 7, 8, 3, 4, 2}
- {1, 5, 4, 6, 7, 7, 8, 8, 1, 4, 0, 1, 0, 0, 7, 7, 9, 9, 9, 7, 4, 5, 6, 4, 3, 3, 6}
- {1, 3, 3, 1, 1, 8, 4, 1, 1, 1, 7, 4, 0, 5, 9, 6, 8, 8, 3, 9, 1, 0, 4, 5, 9, 5, 5}
- {1, 5, 4, 1, 4, 0, 2, 7, 5, 4, 2, 8, 3, 3, 9, 9, 4, 9, 8, 9, 9, 9, 2, 2, 6, 5, 0}
- {2, 9, 5, 7, 6, 8, 2, 3, 7, 3, 6, 1, 2, 9, 1, 7, 0, 8, 6, 4, 5, 2, 2, 7, 4, 7, 4}
- {9, 5, 6, 6, 3, 3, 7, 2, 0, 4, 6, 4, 1, 1, 4, 5, 1, 5, 8, 9, 0, 3, 1, 8, 4, 1, 0}
- {9, 8, 8, 2, 4, 2, 3, 1, 0, 3, 9, 3, 8, 6, 0, 3, 9, 0, 0, 6, 6, 9, 1, 1, 4, 1, 4}
- {9, 4, 9, 3, 9, 7, 6, 8, 4, 0, 3, 9, 0, 2, 6, 5, 8, 6, 8, 5, 2, 2, 0, 6, 7, 2, 0, 0}
- {4, 1, 7, 9, 6, 2, 0, 5, 7, 6, 5, 1, 4, 7, 4, 2, 6, 9, 7, 4, 7, 0, 4, 7, 9, 1, 5, 2, 8}
- {7, 0, 2, 6, 7, 3, 7, 5, 5, 1, 0, 2, 0, 7, 8, 8, 5, 2, 4, 2, 2, 1, 8, 8, 3, 7, 4, 0, 4}

We remark that lattice methods alone can find sporadic large Keith numbers. One approach, from [15], improves the chances of getting a valid result from the lattice reduction step. The idea is to augment each null vector with a zero, and augment the lattice with a row consisting of some nonzero value (typically one) in the new column of zeros, and $\frac{9}{2}$ everywhere else. Thus if there is a valid solution then this augmented lattice contains the vector consisting of that nonzero value (or its negative) and the remaining entries in the range $\{-\frac{9}{2}, \frac{9}{2}\}$. As this would be a fairly "small" vector, one can hope that it will appear in the reduced basis (this is essentially the idea used by Schnorr and Euchner, in a binary setting, to raise the density at which one can typically solve subset sum problems). In practice we get a few Keith numbers this way as well as several more near misses.

It might be effective to combine this different lattice formulation with the branch-and-bound regimen described above. This is not entirely trivial as that required that the vectors span a solution space for a homogeneous equation, whereas the vectors in this modified lattice need not satisfy that equation.

Observe that the Keith numbers beginning at 24 digits but smaller than 29 digits all have leading digit in the set $\{1, 2, 9\}$. One might well wonder if there is a deep reason for this, and whether the trend returns after 29 digits. Also all known Keith numbers from 25 digits onward have a final digit that is either even or 5, and hence they cannot be prime. Again, one might wonder whether this trend continues, and, if so, whether there is an interesting reason behind it.

We will say a bit about the practical complexity of the method we have described for solving ILPs. While in principle branching can have very bad performance, in practice we find that the complexity scales reasonably well with problem size. Although we cannot claim that the methods described in this paper are polynomial time even in fixed dimension, in practice they do seem to scale that way. The behavior with respect to dimension is of course not so nice. For Keith number computations we find that, on average, the time spent for handling $n + 1$ digits is roughly twice the time needed for n digits. Considering that each additional digit multiplies the search space by a factor of 10 this is still

not so bad. To give an indication of computational speed, our implementation was able to handle 29 digits in around three days on a 3.0 GHz machine. We emphasize that this is but a crude measure both of complexity and actual performance, as various sorts of optimizations could have a significant impact on each. For example, preliminary experiments with the software in [6] indicate room for improvement in the handling of the ILP solving after the lattice reduction phase.

4 Frobenius Numbers

We are given a set $A = (a_1, \dots, a_n)$ of positive integers with $\gcd(A) = 1$. For later purposes we assume that the set is in ascending order. It can be shown that there are at most finitely many numbers not representable as a nonnegative integer combination of elements in A . The largest such nonrepresentable is called the Frobenius number of the set. The “Frobenius number problem” is to find it. Generally speaking this tends to be a different and usually harder problem than the Frobenius instance solving discussed earlier. We will give a much abbreviated discussion here in order to indicate how the sort of knapsack solving under discussion plays a role in computation of these numbers. References may be found in [4] and [17].

It turns out that this is trivial for $n = 2$ (this was shown by Sylvester in the late 1800’s, even before the problem was popularized by Frobenius early in the twentieth century). Moreover in the 1980’s some very good methods appeared for the case $n = 3$. For larger n , if one orders the elements by increasing size and restricts a_1 to be less than around 10^7 then there are effective algorithms to find the Frobenius number for A . This is roughly independent of the size of n ; they can handle $n = 100$, for example. Our interest is in handling the case where a_1 , the smallest element in A , is large (say, up to 10^{100}). As the problem is known to be intrinsically difficult we cannot hope to have both a_1 and n large. Hence we limit the latter to 10 or so.

In [17] one finds a definition of a “fundamental domain” which is a generalization that subsumes both lattice diagrams in earlier literature and a graph description from [4]. Similar ideas, expressed in the terminology of Minimal Distance Diagrams, appear in [9]. The Frobenius number will be determined by the furthest corner from the origin, in a suitably weighted l_1 norm. As we will see, an important domain feature is what we term “elbows”. We give a quick idea of what is this fundamental domain for our set $A = (a_1, \dots, a_n)$.

We start with the lattice of integer combinations of $\{a_2, \dots, a_n\}$ that are zero modulo a_1 . This is a full dimensional lattice in \mathbb{Z}^{n-1} . The set of residues of \mathbb{Z}^{n-1} modulo this lattice gives rise to the fundamental domain, which lives in a space of dimension one less than the size of our set. It is not hard to see that there are a_1 distinct residue classes, so we know the cardinality of this domain. We now define the “weight” of a vector $v \in \mathbb{Z}^{n-1}$ as $v \cdot (a_2, \dots, a_n)$. It can be shown that every residue class has at least one element with all nonnegative entries. From those we choose one of minimal weight. In case of a tie we choose the one that

is lexicographically last. This uniquely defines the set of residues that we take to comprise the fundamental domain.

This domain has several interesting properties. (1) It is a staircase: if it contains a lattice element then it contains all nonnegative vectors with any coordinate strictly smaller. (2) It tiles \mathbb{Z}^{n-1} . (3) It is a cyclic group $\mathbb{Z}/a_1\mathbb{Z}$. (4) It can be given a circulant graph structure. It is this structure that was utilized in various shortest-path graph methods. Old and new approaches using such methods are discussed at length in [4]. In contrast, the method put forth in [17] primarily makes use of the staircase structure.


From the staircase property the fundamental domain has turning points we refer to as elbows. It has extremal points called corners. Specifically, a corner is a point c in the domain, such that $c + e_j$ is not in the domain, where e_j is the j^{th} coordinate vector. An elbow is a point x that is not in the domain, but is such that, for each j , either $x_j = 0$ or $x - e_j$ is in the domain. An elbow with all but one coordinate zero is called an “axial” elbow. It indicates how far one can go along a given axis and still remain inside the domain. There are two other definitions that play a role in the algorithm. We will not describe them too carefully but, roughly, there are as follows.

(1) Protoelbows. These have both positive and negative coordinates and correspond to certain “minimal” equivalences (that is, reducing relations) in the lattice.


(2) Praelbows. These are the positive parts of the protoelbows. Elbows are minimal elements in the partially ordered (ascending by inclusion) set of preelbows.

With respect to these domains the Frobenius number corresponds to the farthest corner from the origin where distance is an l_1 metric weighted by element sizes. In brief one sees this as follows. Recall that each element in the domain corresponds to a residue modulo a_1 that satisfies a minimal nonnegativity property. Thus values in the same residue class but of smaller weight cannot be attained using nonnegative combinations of elements of A , whereas values of equal or larger weight are attained as such combinations. We conclude that the largest nonattainable value for the residue class of the element $X = (x_2, \dots, x_n)$ is $a_2x_2 + \dots + a_nx_n - a_1$. From this and the staircase property of the fundamental domain it is clear that the largest nonattainable value overall is a_1 less than the largest weight of a corner element.

Below are pictures (provided courtesy of Stan Wagon) of fundamental domains corresponding to $n = 3$ and $n = 4$ respectively (recall that the fundamental domain lives in a space of dimension one less than n). In the planar diagram the elbows are the lattice points on the axes that bound the diagram, and the lattice point in the interior just outside the “ell”. The corners are the two extremal points reached by intersecting vertical and horizontal lines through the elbows. This picture tells the entire story as regards the $n = 3$ case because it can be shown that there is at most one interior elbow and two corners, and finding them is easy. The three elbows (two axial, one internal) are denoted by circles.



In the three dimensional diagram the elbows are again the axial bounding lattice points as well as bounding points where the staircase goes up in the coordinate planes and in the interior. They are demarcated by tetrahedra.



With this brief background we now say a bit about how knapsack solving can play a role in the computation of Frobenius numbers. First, it turns out that axial elbows are defined by an integer programming problem (see [17]) which, in complexity if not details of definition, is similar to the Frobenius instance problem. From the axial elbows we immediately get good lower and upper bounds on the Frobenius number (each elbow less one is in the domain, and the farthest corner is bounded by the vector with all components given by corresponding axial elbows less one). More importantly for our purposes is that they give a search space from which to find all elbows. They are particular integer points in a polyhedron that satisfy certain inequality conditions. Full details, including a method for finding them, are provided in [17]. From the elbows one can find all corners and in particular the one that gives the Frobenius number of the set.

Another tactic presented in [17] makes direct use of Frobenius instance solving to find axial elbows. One uses a bisection approach, working down from an a priori bound on the axial elbow values. The goal is to find the smallest value for which a certain set of Frobenius instances have no solution.

Still another point of overlap between Frobenius instances and numbers is the obvious fact that whenever a Frobenius instance solver returns an empty solution we automatically have a lower bound on the Frobenius number. One can test random values that are, say, an order of magnitude below the heuristic approximation for the Frobenius number presented in [4]. If any such test gives no solution we thereby establish a lower bound that is often better than a priori bounds to be found in the literature.

There is also a heuristic method in [17] for more efficiently “guessing” the likely Frobenius number from a restricted set of elbows. It gives an a priori upper

bound, and a single Frobenius instance invocation can then verify whether it is in fact the actual value. In random examples this appears always to be the case.

We sketched above how efficient ILP knapsack solving, of the sort used for Frobenius instances, also may be applied to the (generally much harder) problem of finding Frobenius numbers. From the fundamental domain pictures one realizes a possible alternative approach. Working an axis at a time, a branch-and-bound strategy might be directly applied to get to extremal vertices (corners) in the domain. Thus we could have a bilevel branching algorithm, with the outer level iterating over these extrema, and the inner one using relaxed LPs to solve the ILPs needed to move to new vertices. This might become an alternative to the algorithm described in [17].

We remark that there is a connection between another knapsack solving technique and the problem of computing Frobenius numbers. It is well known that integer programming e.g. for knapsack problems can be done with toric Gröbner bases [7]. What is not so obvious is that they may also be used to deduce the stairway structure of the fundamental domain. An algorithm for this purpose is presented in [17]. The basic idea is to formulate a term ordering so that the staircase structure of the fundamental domain is captured by the staircase of the Gröbner basis lead monomials. This has the added virtue of finding all elbows at once, so no elaborate method is needed to search a bounding box defined by axial elbows. An implementation by the author has handled Frobenius number problems involving as many as 7 numbers of 40 digits. While it is not competitive with the main approach in [17] (which has handled sets of up to 11 numbers), Frobenius number problems of this size are apparently larger than what can be handled by other methods from the published literature.

Another link between Frobenius numbers and methods from ideal theory appears implicitly in [9] as well as other literature concerning what are called “multiple-loop networks”. There, as in [4], one works with a circulant directed graph based on residues modulo a positive integer. Now, however, the modulus is the largest rather than smallest element of the given set. While the authors did not explicitly consider the connection to Frobenius numbers, some of the ideas are quite similar. In particular the maximum diameter of this graph, which is similar to the Frobenius number (though using an unweighted metric), plays an important role in their work. They discuss several aspects of the related domain (actually family of domains, as they do not impose uniqueness conditions) in terms of monomial ideals. They define a generalized ell shape and prove their domains are always of such a shape; this corresponds to the uniqueness of the interior elbow as shown in [17]. It would be interesting to understand better how their ideas, in particular regarding use of monomial ideals, relate to the toric ideal construction of the fundamental domain given in [17].

5 Summary

We have investigated several examples of integer linear programs with the common feature that a straightforward branch-and-bound approach, working with

real relaxations, will tend to bog down in searching a large polytope. We utilize a method based on solving of an integer relaxation to the set of equality constraints. We reformulate the problem as one of adding combinations of null vectors to a specific solution. The vectors in question tend to be well suited to the problem at hand because we use lattice reduction to make them close to orthogonal. We then enforce inequality constraints via branch-and-bound on real relaxations of the new problem. We use a branching choice that tends to make the polytope thin in the search direction and thus helps to exhaust it efficiently.

We reviewed this approach as it was applied to the change-making problem [1,2]. We then used it to find all Keith numbers through 29 digits; previous methods had gone only through 19 digits. This brought us to a range where we could observe curious patterns in leading and trailing digits, something not present in the smaller Keith numbers. We also gave a brief idea of how this method is used in a new algorithm to compute Frobenius numbers.

A further direction would be to incorporate effective cutting planes. The code we used only attempts a very naive sort of cut (by using random coordinate variables as objective function). Preliminary experiments with an external library [6] indicate that more serious cutting plane efforts can give substantial speed improvement; we have seen ILP examples that improve by an order of magnitude. We emphasize that this still requires preprocessing with lattice reduction in the manner described in this paper.

Acknowledgments. I thank Stan Wagon for providing the diagrams of fundamental domains and for his careful reading and many suggestions that improved the exposition. I thank Victor Moll for inviting me to attend the 2005 Clifford Lectures conference at Tulane, where I presented an earlier version of this work. I thank the two anonymous referees for their several suggestions which improved readability, and the second referee in particular for pointing out the relevance of [9] and related literature on multiple-loop networks.

References

1. K. Aardal, C. A. J. Hurkens, and A. K. Lenstra. Solving a system of linear diophantine equations with lower and upper bounds on the variables. *Mathematics of Operations Research*, 25(3):427–442, 2000.
2. K. Aardal and A. K. Lenstra. Hard equality constrained integer knapsacks. In *Proceedings of the 9th Conference on Integer Programming and Combinatorial Optimization (IPCO 2002)*, volume 2337 of *Lecture Notes in Computer Science*, pages 350–366. Springer-Verlag, 2002.
3. K. Aardal, R. Weismantel, and L. A. Wolsey. Non-standard approaches to integer programming. *Discrete Appl. Math.*, 123(1-3):5–74, 2002.
4. D. Beihoffer, J. Hendry, A. Nijenhuis, and S. Wagon. Faster algorithms for frobenius numbers. *Elec. J. Combinatorics*, 12, 2005.
5. W. Blankenship. Algorithm 288: Solution of simultaneous linear diophantine equations. *Communications of the ACM*, 9(7):514, 1966.

6. COmputational Infrastructure for Operations Research (COIN-OR). <http://www.coin-or.org/documentation.html>.
7. P. Conti and C. Traverso. Buchberger algorithm and integer programming. In *AAECC-9: Proceedings of the 9th International Symposium on Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, volume 539 of *Lecture Notes in Computer Science*, pages 130–139. Springer-Verlag, 1991.
8. G. B. Dantzig. *Linear Programming - 1: Introduction*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, Princeton, NJ, USA, 1963.
9. D. Gómez-Perez, J. Gutierrez, and Á. Ibeas. Circulant digraphs and monomial ideals. In *Proceedings of the 8th International Workshop on Computer Algebra in Scientific Computing (CASC 2005)*, volume 3718 of *Lecture Notes in Computer Science*, pages 196–207, 2005.
10. M. Keith. Determination of all keith numbers up to 10^{19} . Available from <http://users.aol.com/s6sj7gt/keithnum.htm>, 1998.
11. A. Lenstra, H. L. Jr., and L. Lovasz. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
12. D. Lichtblau. Revisiting strong Grbner bases over Euclidean domains. Manuscript, 2003.
13. K. Matthews. Short solutions of a $x=b$ using a l_1 -based hermite normal form algorithm. Manuscript, 2001.
14. P. Nguyen. Cryptanalysis of the goldreich-goldwasser-halevi cryptosystem from crypto '97. In *Proceedings of Advances in Cryptology - CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 288–304. Springer-Verlag, 1999.
15. C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. In *Proceedings of the 8th International Conference on Fundamentals of Computation Theory*, volume 529 of *Lecture Notes in Computer Science*, pages 68–85. Springer-Verlag, 1991.
16. A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986.
17. S. Wagon, D. Einstein, D. Lichtblau, and A. Strzebonski. Frobenius numbers by lattice enumeration. Submitted, 2006.
18. S. Wolfram. *The Mathematica Book*. Wolfram Media, Champaign, IL, USA, 5th edition, 2003.