

Logic Minimization Techniques with Applications to Cryptology*

Joan Boyar[†]

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark
joan@imada.sdu.dk

Philip Matthews[‡]

Aarhus University, Aarhus, Denmark

René Peralta

Information Technology Laboratory, NIST, Gaithersburg, MD, USA
rene.peralta@nist.gov

Communicated by Kaisa Nyberg

Received 8 February 2011

Online publication 3 May 2012

Abstract. A new technique for combinational logic optimization is described. The technique is a two-step process. In the first step, the nonlinearity of a circuit—as measured by the number of nonlinear gates it contains—is reduced. The second step reduces the number of gates in the linear components of the already reduced circuit. The technique can be applied to arbitrary combinational logic problems, and often yields improvements even after optimization by standard methods has been performed. In this paper we show the results of our technique when applied to the S-box of the Advanced Encryption Standard (FIPS in Advanced Encryption Standard (AES), National Institute of Standards and Technology, 2001).

We also show that, in the second step, one is faced with an NP-hard problem, the Shortest Linear Program (SLP) problem, which is to minimize the number of linear operations necessary to compute a set of linear forms. In addition to showing that SLP is NP-hard, we show that a special case of the corresponding decision problem is MAX SNP-complete, implying limits to its approximability.

Previous algorithms for minimizing the number of gates in linear components produced cancellation-free straight-line programs, i.e., programs in which there is no cancellation of variables in GF(2). We show that such algorithms have approximation ratios of at least $3/2$ and therefore cannot be expected to yield optimal solutions to non-trivial inputs. The straight-line programs produced by our techniques are not always cancellation-free. We have experimentally verified that, for randomly chosen linear

* Some of this work appeared in [10] and some appeared in [7].

[†] Some of J. Boyar's work was done while visiting the University of California, Irvine, and some while visiting Aarhus University.

[‡] Work of P. Matthews was done while at Aarhus University.

transformations, they are significantly smaller than the circuits produced by previous algorithms.

Key words. Circuit complexity, Multiplicative complexity, Linear component minimization, Shortest Linear Program, Cancellation, AES, S-box.

1. Introduction

Constructing optimal combinational circuits is an intractable problem under almost any meaningful metric (gate count, depth, energy consumption, etc.). In practice, no known techniques can reliably find optimal circuits for functions with as few as eight Boolean inputs and one Boolean output (there are 2^{256} such functions). As an example of this, consider multiplicative complexity, the number of GF(2) multiplications (i.e., AND gates) necessary and sufficient to compute a function. The multiplicative complexity of the Boolean function E_4^8 , which is true if and only if exactly four of its eight input bits are true, is unknown [5].


In practice, we build circuit implementations of functions using a variety of heuristics. Many of these heuristics have exponential time complexity and thus can only be applied to small components of a circuit being built. This works reasonably well for functions that naturally decompose into repeated use of small components. Such functions include arithmetic functions (which we often build using full adders), matrix multiplication (which decomposes into multiplication of small submatrices), and more complex functions such as cryptographic functions (which are commonly based on multiple iterations of an algorithm containing linear steps and one or more nonlinear steps).

This work presents a new technique for logic synthesis and circuit optimization with the goal of minimizing the total number of gates. The reason for considering this minimization is that in an actual circuit implementation this would lead to smaller area and less power consumption, and in a program implementation this would lead to faster execution times. The technique can be applied to arbitrary functions, and yields improvements even on programs/circuits that have already been optimized by standard methods. We apply our technique to the S-box of the Advanced Encryption Standard (AES),¹ which, in addition to being used in AES, has been used in several proposals for a new hash function standard.² The result is, as far as we know, the smallest circuit yet constructed for this function. The circuit contains 32 AND gates and 83 XOR/XNOR gates for a total of 115 gates. We have also applied these techniques to the logic embedded in the nonlinear components of several candidates to the SHA-3 competition. The improvements in software performance were significant.

Finally, we note that there is another metric that is important in the area of secure multi-party computation. The computational cost of various applications in this area is proportional, not to the total number of gates, but rather to the number of nonlinear gates. For example, what is known as the “Free-XOR” technique [25] uses circuits with low multiplicative complexity to speed up multi-party computation. The work of [21] is currently the fastest implementation of a two-prover protocol for AES. That work uses

¹ Our circuit for the AES S-box has already been used as the basis of a software bit-sliced implementation of AES in counter mode [23].

² See <http://csrc.nist.gov/groups/ST/hash/sha-3/index.html>.



$t_1 = a \wedge b$	$t_2 = a \oplus b$	$t_3 = t_2 \wedge c$	$t_4 = t_1 \oplus t_3$
$u_1 = a \oplus b$	$u_2 = b \oplus c$	$u_3 = u_1 \wedge u_2$	$u_4 = u_3 \oplus b$

Fig. 1. Two circuits and corresponding straight-line programs for MAJ(a, b, c).

the circuit in [39] where the S-box contains 123 linear gates and 58 nonlinear gates. In contrast, our S-box circuit contains only 32 nonlinear gates. Thus, using our circuit immediately leads to a significant reduction in the cost of those protocols.

1.1. Preliminaries

The goal of this work is to minimize the total number of gates used in a circuit for a function. Our circuits are over the basis $\{\oplus, \wedge, 1\}$, where \oplus represents XOR, \wedge represents AND, and the 1 can be used with the \oplus for complementing bits. This basis is logically complete: any Boolean circuit can be transformed into this form using only local replacements. The circuit operations can be viewed either as performing Boolean logic or arithmetic modulo 2 (when viewing it the latter way, we will write outputs to be computed as polynomials with multiplication replacing \wedge and addition replacing \oplus). The number of \wedge gates is called the *multiplicative complexity* of the circuit. Connected components of the circuit containing \wedge gates are called *nonlinear*. Components free of \wedge gates are called *linear*. Circuits and programs for computing Boolean functions can be defined using straight-line programs, where each statement defines the operation of a gate or a line in a program. The examples in Fig. 1 define two different circuits and their corresponding straight-line programs for computing the majority function of three inputs, a, b , and c .

The *Boolean complexity* of a function is the minimal number of gates sufficient to compute that particular function, when any two-input one-output gates are allowed. A *Boolean predicate* is a Boolean function with only one output bit.

1.2. Combinational Circuit Optimization

The techniques described here would generally be applied to subcircuits of a larger circuit, such as an S-box in a cryptographic application, which have relatively few inputs and outputs connecting them to the remainder of the circuit. The key observation that

led us to our techniques is that circuits with low multiplicative complexity will naturally have large sections which are purely linear (i.e., contain only \oplus gates). Thus

it is plausible that a two-step process, which first reduces multiplicative complexity and then optimizes linear components, would usually lead to small circuits over the basis $\{\oplus, \wedge, 1\}$.

We have, of course, no way of proving this hypothesis. In fact, it seems unlikely that circuits with optimal Boolean complexity which are also optimal with respect to multiplicative complexity always exist. In practice, though, we conjecture that this two-step method will usually yield “good” circuits as compared with other methods, primarily because of the improved techniques presented here for optimizing linear circuits. As mentioned above, in this paper, we apply this method to the AES S-box and present experiments testing the techniques for optimizing linear circuits. Additionally, we (and others, see, e.g., [16]) have successfully applied the heuristics described in this paper to a number of circuit optimization problems of interest to cryptology. These include finite field arithmetic and binary multiplication. New records (i.e., circuits with fewer gates than previously known) are periodically posted at <http://cs-www.cs.yale.edu/homes/peralta/CircuitStuff/CMT.html>.

1.2.1. First Step

The first step of our technique consists of identifying nonlinear components of the sub-circuit to be optimized and reducing the number of \wedge gates. This reduction is not easy to do. For example, it is not obvious how to algorithmically transform one of the two equivalent circuits defined in Fig. 1 into the other.

Classic results by Shannon [35] and Lupanov [26] show that almost all predicates on n bits have a Boolean circuit complexity of about $\frac{2^n}{n}$. Analogous to the Shannon-Lupanov bound, it was shown in [9] that almost all Boolean predicates on n bits have a multiplicative complexity of about $2^{\frac{n}{2}}$. Strictly speaking, these theorems say nothing about the class of functions with polynomial circuit complexity. However, it is reasonable to expect that, in practice, the multiplicative complexity of functions is significantly smaller than their Boolean complexity.

Finding circuits with minimum multiplicative complexity is, in all likelihood, a highly intractable problem. However, recent work on multiplicative complexity contains an arsenal of reduction techniques that in practice yield circuits with small, and often optimal, multiplicative complexity [5]. That work focuses exclusively on symmetric functions (those whose value depends only on the Hamming weight of the input).

In this paper we use ad hoc heuristics to construct a circuit with low multiplicative complexity for inversion in $\text{GF}(2^4)$. ($\text{GF}(2^n)$ is the field with 2^n elements.) The technique is described in Sect. 2.1.

1.2.2. Second Step

The second step of our technique consists of finding maximal linear components of the circuit and then minimizing the number of XOR gates needed to compute the target functions computed in these linear components. A new heuristic for this computationally intractable problem is described in Sect. 3.2.

1.3. The Shortest Linear Program Problem

We argue below that minimizing the number of XOR gates in the second step is equivalent to solving the Shortest Linear Program (SLP) problem over $\text{GF}(2)$.

Let \mathbb{F} be an arbitrary field and let

$$\begin{aligned} &\alpha_{1,1}x_1 + \alpha_{1,2}x_2 + \cdots + \alpha_{1,n}x_n, \\ &\alpha_{2,1}x_1 + \alpha_{2,2}x_2 + \cdots + \alpha_{2,n}x_n, \\ &\vdots \\ &\alpha_{m,1}x_1 + \alpha_{m,2}x_2 + \cdots + \alpha_{m,n}x_n, \end{aligned}$$

be a set of linear forms where the $\alpha_{i,j}$'s are constants from \mathbb{F} and the x_i 's are variables over \mathbb{F} .

Suppose a subcircuit for a linear component in a circuit has x_i s as inputs and y_j s as outputs.³ The y_j s are linear functions of the x_i s in the field $\text{GF}(2)$, so the subcircuit is an algorithm for computing the linear forms (the functions the y_j s represent) given the x_i s as input, in the special case where $\mathbb{F} = \text{GF}(2)$.

We consider this question in the model of computation known as *linear straight-line programs*. A linear straight-line program is a variation on a straight-line program which does not allow multiplication of variables. That is, every line of the program is of the form $u := \eta v + \mu w$; where η, μ are in \mathbb{F} and v, w are variables. Some of the lines are output lines; these are the lines where the linear forms in the set are produced. For brevity, we will use the terms *linear programs* or simply *programs* to refer to linear straight-line programs. The *length* of the program is the number of lines it contains, and is equal to the number of XOR gates in a subcircuit computing these forms. A program is *optimal* if it is of minimum length.

The linear straight-line program model (see [11] for a discussion of linear complexity) has the advantage of being very structured, but is nevertheless optimal to within a constant factor as compared to arbitrary straight-line programs when the computation is over an infinite field. Over finite fields the optimality of linear straight-line programs is unknown,⁴ but we restrict our attention to this form and consider minimizing the length of the program.

The standard algorithm for computing the linear forms $A\mathbf{x}$, where A is an $m \times n$ matrix containing entries from a set of size r , requires $m(n - 1)$ operations. However, Savage [34] showed that $O(mn / \log_r m)$ operations are sufficient in many cases, including computations over $\text{GF}(2)$ if $m \geq 4$. Williams [37] improved this to $O(n^2 / \log^2 n)$ on a RAM with word length $\Theta(n)$ for n by n matrices over finite semirings. In contrast, Winograd [38] has shown that most sets of linear forms have a nonlinear complexity in the straight-line program model; in fact, for a “random” $m \times n$ matrix A the probability is high that its complexity is $\Omega(mn)$ (for infinite fields). However, there are nontrivial matrices which can be computed considerably faster than this.

³ We consider circuits without negations only. There is no loss of generality in doing so, because negations can be treated as standard XOR gates via $(-X = (X \oplus 1))$.

⁴ It is not known if multiplication of variables can ever be used to reduce program length when the program outputs only linear functions.

Over $\text{GF}(2)$, finding the shortest linear straight-line program is equivalent to our original goal of finding a circuit with only XOR gates and minimizing the number used. Linear forms have many applications, especially to problems in scientific computation, and there has been considerable success in finding efficient algorithms for computing them in special cases. The best known example is the fast Fourier transform, an $O(n \log n)$ algorithm, discovered by Cooley and Tukey in 1965 [15].

In Sect. 3.1.1 we show that finding the shortest linear straight-line program is NP-hard. This can be seen in relation to Håstad's result [20] showing that tensor rank is NP-hard and thus finding the minimum bilinear program for computing bilinear forms is NP-hard.

In Sect. 3.1.2 the NP-hardness result is used to prove a special case of the problem MAX SNP-complete [32] (and also APX-complete). This means there are no ϵ -approximation algorithms for the problem unless $\text{P} = \text{NP}$ [1].

A linear straight-line program over $\text{GF}(2)$ is said to be a *cancellation-free straight-line program* if, for every line of the program $u := v + w$, none of the variables in the expression for v are also present in the expression for w ; i.e., there is no cancellation of variables in the computation. A small example showing that the optimal linear program is not always cancellation-free over $\text{GF}(2)$ is:

$$x_1 + x_2; \quad x_1 + x_2 + x_3; \quad x_1 + x_2 + x_3 + x_4; \quad x_2 + x_3 + x_4.$$

It is not hard to see, by exhaustive search, that the optimum cancellation-free straight-line program has length 5. A solution of length 4 which allows cancellations is

$$v_1 = x_1 + x_2; \quad v_2 = v_1 + x_3; \quad v_3 = v_2 + x_4; \quad v_4 = v_3 + x_1.$$

In Sect. 3.1.3 we show that the approximation ratio for cancellation-free techniques is at least $3/2$. This discovery led us to create the heuristic in Sect. 3.2, allowing cancellations, for minimizing linear straight-line programs and the corresponding circuits.

2. First Step

We will illustrate the first step of the circuit minimization using AES's S-box as an example. The nonlinear operation in AES's S-box is to compute an inverse in the field $\text{GF}(2^8)$. A recursive method for building a circuit for inverses in $\text{GF}(2^m)$, given a circuit for inverses in $\text{GF}(2^m)$, is due to Itoh and Tsujii [22]. The circuits produced by this method are said to have a *tower fields architecture*. Since there are multiple possible representations for Galois fields, several authors have concentrated on finding representations that yield efficient circuits under the tower fields architecture. We use the same general technique for the reduction from inversion in $\text{GF}(2^8)$ to $\text{GF}(2^4)$ inversion, but we use a completely different technique for computing the inversion in $\text{GF}(2^4)$. We then place the optimized circuit for $\text{GF}(2^4)$ inversion in its appropriate place in AES's S-box and, in the second step, apply a novel optimization technique to the linear parts of the resulting circuit.

2.1. $\text{GF}(2^4)$ Inversion: A Nonlinear Component

The tower fields architecture for inversion in $\text{GF}(2^8)$ has (nontrivial) easily identifiable nonlinear components corresponding to inversion in subfields. The first step in our method is to focus on one of these components and derive a circuit that uses few \wedge gates. The component for inversion in $\text{GF}(2^2)$ is too small for us to benefit significantly from optimizing it. Instead we focus on inversion in $\text{GF}(2^4)$. There are many representations of $\text{GF}(2^4)$. Following Canright [13], who compared 432 different representations as a tower of fields [12] and found this one optimal using his techniques for reducing circuit size, we construct

- $\text{GF}(2^2)$ by adjoining a root W of $x^2 + x + 1$ over $\text{GF}(2)$;
- $\text{GF}(2^4)$ by adjoining a root Z of $x^2 + x + W^2$ over $\text{GF}(2^2)$.

$\text{GF}(2^2)$ is represented using the basis (W, W^2) , and $\text{GF}(2^4)$ using the basis (Z^2, Z^8) . Thus, an element $\delta \in \text{GF}(2^4)$ is written as $\delta_1 Z^2 + \delta_2 Z^8$, where $\delta_1, \delta_2 \in \text{GF}(2^2)$. Similarly, an element γ in $\text{GF}(2^2)$ is written as $\gamma_1 W + \gamma_2 W^2$, where $\gamma_1, \gamma_2 \in \text{GF}(2)$. Since Z satisfies $x^2 + x + W^2 = 0$ and W satisfies $x^2 + x + 1 = 0$, one can calculate that $Z^4 = Z^2 + W$, $Z^8 = Z^2 + 1$ ($1 = Z^8 + Z^2$), $Z^{10} = Z^4 + Z^2$, $Z^{16} = Z^8 + W$, $W^3 = W^2 + W$, $W^4 = W$, and $W^5 = W^2$. These equations can be used to reduce expressions to check equalities.

Using this representation, an element of $\text{GF}(2^4)$ can be written as $\Delta = (x_1 W + x_2 W^2)Z^2 + (x_3 W + x_4 W^2)Z^8$, where $x_1, x_2, x_3, x_4 \in \text{GF}(2)$. The inverse of this element, $\Delta' = (y_1 W + y_2 W^2)Z^2 + (y_3 W + y_4 W^2)Z^8$, can then be calculated using the following polynomials over $\text{GF}(2)$:

- $y_1 = x_2 x_3 x_4 + x_1 x_3 + x_2 x_3 + x_3 + x_4$
- $y_2 = x_1 x_3 x_4 + x_1 x_3 + x_2 x_3 + x_2 x_4 + x_4$
- $y_3 = x_1 x_2 x_4 + x_1 x_3 + x_1 x_4 + x_1 + x_2$
- $y_4 = x_1 x_2 x_3 + x_1 x_3 + x_1 x_4 + x_2 x_4 + x_2$

The fact that Δ' is the inverse of Δ can be verified by multiplying the two elements together and reducing using the equations mentioned above (along with $x^2 = x$ and $x + x = 0$). The symbolic result is $(QW + QW^2)Z^2 + (QW + QW^2)Z^8$, where $Q = x_1 x_2 x_3 x_4 + x_1 x_2 x_3 + x_1 x_2 x_4 + x_1 x_3 x_4 + x_2 x_3 x_4 + x_1 x_2 + x_1 x_3 + x_1 x_4 + x_2 x_3 + x_2 x_4 + x_3 x_4 + x_1 + x_2 + x_3 + x_4$. The fact that the value of Q is 1 unless all four variables have the value 0, when it is 0, can be seen by observing that it is the symmetric function $\Sigma_4^4 + \Sigma_3^4 + \Sigma_2^4 + \Sigma_1^4$. If exactly four variables are set, then the first term gives the value 1 (and the others 0); if three are set, then the second, third, and fourth terms give the value 1; if exactly two are set, then only the third gives the value 1; and if only one is set, then only the last gives the value 1. Hence, the result is 1, except for the zero input.⁵

Thus the task at hand is to construct a circuit with four inputs and four outputs that calculates the above system of equations using as few \wedge gates as possible. Currently, our

⁵ A circuit for finite field inversion must have some output for the noninvertible zero element. In the following constructions we follow the AES convention that the output on input zero is zero.

heuristic search programs can handle functions with one output and up to eight inputs. (Since they are heuristics, one is not certain that an output is optimal, so they cannot be used, for example, to determine a tight lower bound for the multiplicative complexity of E_4^8 .) This means that we can directly construct optimal circuits for each of the four equations individually, but not for the system itself. For the full system we took the following approach:

1. pick an equation and construct an efficient circuit for it;
2. store intermediate functions computed in the previous steps for possible use in constructing a circuit for the next equation to be tackled;
3. iterate until all equations have been computed.

The first step is nontrivial even for predicates on few inputs. The heuristic we used is inspired by methods from automatic theorem proving [6]: consider an arbitrary predicate f on n inputs. We refer to the last column of the truth table for f as the *signal* of f . The columns in the truth table corresponding to each of the inputs to f are *known* signals. A search for a circuit for f starts with this set S of known signals. If u, v are known signals for functions g, h respectively, then the bit-wise XOR (AND) of u and v is the signal for the predicate $g \oplus h$ ($g \wedge h$). We can *grow* the set S by adding the XOR of randomly chosen signals. We call this step an *XOR round*. The analogous step where the AND of signals is added to S is called an *AND round*. Each round is parameterized by the number of new signals added and the maximum number of AND gates allowed. In either an XOR round or an AND round, two signals are not combined if doing so creates a signal with more AND gates than is allowed. The heuristic alternates between XOR and AND rounds until the target signal is found or the set S becomes too large. In the latter case, since this is a randomized procedure, we start again. Various enhancements and optimizations have been implemented. Their description is outside the scope of this paper. We can report, however, that we succeeded in determining the multiplicative complexity of all 2^{16} predicates on four bits. It turns out that 3 multiplications are enough to compute any predicate on four variables.⁶ This is of interest to designers of cryptographic functions, since many constructions have been proposed which use 4×4 S-boxes. We have not yet been able to do the same for all predicates on 5 bits.

We performed the three steps above for each of the 24 orderings of $\{y_1, y_2, y_3, y_4\}$. The ordering (y_4, y_2, y_1, y_3) gave the best results. The resulting circuit, expressed as a straight-line program over $\text{GF}(2)$, is shown in Fig. 2 (outputs are indicated by an (*)).

This circuit contains 5 \wedge gates and 11 \oplus gates. It is a significant improvement over previous constructions; e.g., Paar's construction [30] has a gate count of 10 \wedge gates and 15 \oplus gates for the same function. It is harder to compare to Canright's construction [13]. In his original, he had 9 \wedge gates (and NAND gates) and 14 \oplus gates (and XNOR gates),

⁶ Lest the reader think this trivial, he/she may attempt to compute the function $f(x_1, x_2, x_3, x_4) = x_1x_2x_3x_4 + x_1x_2x_3 + x_1x_2x_4 + x_2x_3x_4 + x_1x_2 + x_1x_3 + x_1x_4 + x_2x_3 + x_3x_4$ using only three multiplications.

$t_1 = x_1 + x_2$	$t_2 = x_1 \times x_3$	$t_3 = x_4 + t_2$
$t_4 = t_1 \times t_3$	$y_4 = x_2 + t_4$ (*)	$t_5 = x_3 + x_4$
$t_6 = x_2 + t_2$	$t_7 = t_6 \times t_5$	$y_2 = x_4 + t_7$ (*)
$t_8 = x_3 + y_2$	$t_9 = t_3 + y_2$	$t_{10} = x_4 \times t_9$
$y_1 = t_{10} + t_8$ (*)	$t_{11} = t_3 + t_{10}$	$t_{12} = y_4 \times t_{11}$
$y_3 = t_{12} + t_1$ (*)		

Fig. 2. Inversion in $\text{GF}(2^4)$.

but he optimized, allowing NOR gates. After this, he had 8 NAND gates, 2 NOR gates, and 9 XOR/XNOR gates.

Under the given representation for $\text{GF}(2^4)$, the multiplicative complexity of inversion is 5. This can be argued as follows: the upper bound is given by the construction. The four outputs that have to be computed all have degree 3. One \wedge is needed to compute a polynomial of degree 2. Then, an additional \wedge is necessary to produce each of the four linearly independent polynomials, since each is of degree 3.

2.2. A View of the Structure of AES's S-Box

In the previous section, using the tower fields architecture, we identified and optimized (with respect to multiplicative complexity) a major nonlinear component in an implementation of the AES S-box. The multiplications in $\text{GF}(2^4)$ are also nonlinear, but we have used the same circuit as Canright for these components. That completes the first step of our technique for circuit optimization, but in other circuits, one may be able to identify more nonlinear components with few enough inputs that they can also be optimized before continuing. At this point, we replaced the $\text{GF}(2^4)$ inversion subcircuit, in Canright's [13] (already optimized) circuit, with the subcircuit in Fig. 2. As expected, the resulting circuit contained large linear connected components. In fact, from a cryptanalyst's point of view, the topology of the resulting circuit is potentially of interest: the S-box of AES consists of an initial linear expansion U from 8 to 22 bits, followed by a nonlinear contraction F from 22 to 18 bits, and ending with a linear contraction B from 18 to 8 bits. The U and B matrices are given below. AES's S-box is $S(\mathbf{x}) = B \cdot F(U \cdot \mathbf{x}) + [11000110]^T$, where \cdot is matrix multiplication and \mathbf{x} is the 8-bit S-box input. We do not know if there are any cryptanalytic implications to the structure of these matrices. The first row and last columns of U should raise an eyebrow, as should the 12th and the last three columns of B . Note that the initial linear expansion and the linear contraction were defined to contain as much of the circuit as possible while still being linear, increasing the portion of the circuit which could be further optimized by concentrating on the linear components. Thus, the portion of the circuit defined by U , for example, overlaps with the $\text{GF}(2^8)$ inversion. Also included in these linear components is the linear transformation to change bases, before computing the inverse in $\text{GF}(2^8)$, plus the linear transformation to change back to the original basis, followed by the affine transformation which is the final operation in the

S-box.

$$U = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 & 0 & 0 \end{bmatrix},$$

$$B = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}.$$

3. Second Step

The second step is to optimize the linear components in the circuit. One method of finding the nonlinear components to be optimized in the first step was to find maximal linear components of the circuit, remove them, and look at the remaining nonlinear components. Whether this was done or not, after the optimized nonlinear components are inserted into their appropriate places in the circuit, the beginning of the second step should be to find maximal linear components in this new circuit (since after optimization, some of the nonlinear portions may contain \oplus gates which can be included in the “old” linear parts, as in the case of the U and B matrices from AES’s S-box).

These maximal components define linear components of the circuit which should be minimized in Step 2. In the case of the AES S-box, the top linear component corresponds to the matrix U , and the bottom linear component corresponds to the matrix B .

No other significant linear components were found. After finding these, the next step was to minimize the circuits for computing U and B .

3.1. Hardness of Minimizing Linear Components

First, we show that the problem of linear circuit minimization, or equivalently, the Shortest Linear Program (SLP), is NP-hard.

3.1.1. NP-Hardness

The problem SHORTEST LINEAR PROGRAM (SLP) is as follows: Given a set of linear forms E over a field \mathbb{F} , find a shortest linear program to compute E .

In order to prove NP-hardness, we consider the corresponding decision problem, SLPd: Given a set of linear forms E over a field \mathbb{F} and a positive integer k , determine if there exists a straight-line linear program with at most k lines which computes E .

We will prove SLPd NP-hard, even if the constants in the set of linear forms to be computed are only zeros and ones. Furthermore, if the field \mathbb{F} is finite, then SLPd is easily seen to be in NP, so SLPd is NP-complete over finite fields.⁷

The interest of this section is not just in the final result that SLP is NP-hard, but also in the method used to prove it. In particular, most of this section is devoted to the proof of Lemma 1, which gives the exact complexity for sets of linear forms of a certain simple type. This proof is *algorithmic* in form, and its algorithmic nature can be exploited to prove a further result in Sect. 3.1.2.

In order to show NP-hardness, we reduce from VERTEX COVER. A *vertex cover* of a graph $G = (V, E)$ is a subset V' of V such that every edge of E is incident with at least one vertex of V' . VERTEX COVER is defined as follows: Given a graph $G = (V, E)$ and an integer k , determine if there exists a vertex cover of size at most k .

The following polynomial-time reduction f transforms an arbitrary graph, $G = (V, E)$, and a bound, k , to a set of linear forms with another bound, \bar{k} . The input variables are $X = V \cup \{z\}$, where z is a distinguished variable not occurring in V . The linear forms are $\bar{E} = \{z + a + b \mid (a, b) \in E\}$, and the program length we ask about is $\bar{k} = k + |\bar{E}|$. This is an instance of SLPd, and it is clear that $f(G, k) = (\bar{E}, X, \bar{k})$ can be produced in polynomial time. We call a set of linear expressions in this restricted form, $z + x_i + x_j$, a set of *z-expressions*. Note that three distinct z-expressions are linearly independent over any field.

Before we proceed, we illustrate with an example. The graph, G , in Fig. 3 has a vertex cover of size $k = 3$: $\{a, c, e\}$. The corresponding instance of SLPd, $f(G, 3)$ is $\bar{E} = \{z + a + b, z + b + c, z + c + d, z + d + e, z + e + f, z + a + f, z + c + g, z + e + g\}$, $X = \{z, a, b, c, d, e, f, g\}$, and $\bar{k} = 3 + 8$. A linear program for this of size 11 is

$$\begin{array}{llll} v_1 := z + a; & v_2 := z + c; & v_3 := z + e; & v_4 := v_1 + b; \\ v_5 := v_2 + b; & v_6 := v_2 + d; & v_7 := v_3 + d; & v_8 := v_3 + f; \\ v_9 := v_1 + f; & v_{10} := v_2 + g; & v_{11} := v_3 + g; & \end{array}$$

⁷ We avoid the discussion of models for dealing with infinite fields, such as in [36] or [4], by proving NP-hardness when the constants in the forms are only zeros and ones and showing that a shortest linear straight-line program for the forms considered can be created with only zeros and ones as constants.




Fig. 3. Graph with 8 edges and cover size 3.

where the computation of $v_1, v_2,$ and v_3 corresponds to the vertex cover in the graph G , and the remaining operations produce the eight forms in \bar{E} . The variables v_4, \dots, v_{11} are called *output variables*.

A *cover* for a set \bar{E} of z -expressions is a subset W of $X \setminus \{z\}$ such that every expression in \bar{E} contains at least one variable in W . Note that if $(\bar{E}, X, \bar{k}) = f(G, k)$, a cover for \bar{E} trivially defines a vertex cover for the graph G and vice versa.

Lemma 1. *Let (\bar{E}, X) be a set of z -expressions without repetitions; that is, \bar{E} is a set of expressions of the form $z + x_i + x_j$, where x_i, x_j are distinct variables in $X \setminus \{z\}$, z is a distinguished variable in X , and no two of these z -expressions contain exactly the same variables. There is a cover of \bar{E} of size at most k if and only if there is a linear straight-line program P for \bar{E} of length $\bar{k} = k + |\bar{E}|$. In addition, given a linear straight-line program P for \bar{E} , a cover for \bar{E} of size at most $|P| - |\bar{E}|$ can be computed in polynomial time.*

Proof. We will refer to the elements of $X \setminus \{z\}$ as “the variables” and z as “the symbol,” although as an element of a linear program, z is also an input variable.

Given a cover W of size k for \bar{E} , a (cancellation-free) linear straight-line program for \bar{E} can be created consisting of $z + w_i$ for each $w_i \in W$, followed by linear expressions computing each output, created by adding a second variable to the appropriate $z + w_i$. This program has length $k + |\bar{E}|$.

It remains to be shown that, given a linear straight-line program P for \bar{E} , we can efficiently find a cover, W , for \bar{E} of size no more than $|P| - |\bar{E}|$. This cover is computed by associating elements of $X \setminus \{z\}$ with some non-output lines of the program— W will then be the union of all the variables so associated. Since we will assign at most one element of $X \setminus \{z\}$ to each non-output line, the cover is of size at most $|P| - |\bar{E}|$.

Let $F^{(i)}$ be the linear function computed at line i ; the result there is assigned to v_i . It will be convenient to use the notation $F^{(i)}$ to refer both to the function and to the set of variables (not including z) in $F^{(i)}$. The association of variables with lines of the program will be denoted by a mapping $m : \mathbb{N} \rightarrow X \setminus \{z\} \cup \{\lambda\}$. Initially, we set $m(i) = \lambda$ for all lines i . When line i is first processed, $m(i)$ will, in some cases, be set to a variable in $F^{(i)}$. We define

$$W^{(i)} = \{x \in X \mid x = m(j) \text{ for some } j \leq i\}.$$

```

Compute-Cover()
for  $i = 1$  to  $|P|$  do
     $m(i) \leftarrow \lambda$ 
    if  $F^{(i)}$  is not an output then
        if  $\exists$  a variable  $x$  in  $F^{(i)} \setminus W^{(i-1)}$  then
             $m(i) \leftarrow x$ 
        else {  $F^{(i)}$  is an output }
            if  $|F^{(i)} \setminus W^{(i)}| = 2$  then
                Fix-up( $i, i$ )      (see Fig. 6)

```

Fig. 4. Computing the cover W .

At the end, the cover W will be $W^{(|P|)}$, the set of all variables that are assigned to some $m(i)$.

The algorithm **Compute-Cover** works as follows (see Fig. 4): Starting at the first line of the linear straight-line program P , it associates with each non-output line i a variable in $X \setminus \{z\}$ which occurs in $F^{(i)}$ and which is currently unassociated (if there is no such variable, the line is assigned the null symbol, λ). When an output is reached, **Compute-Cover** checks if the set of all variables currently assigned to earlier lines covers that output, i.e., if there is some variable in $W^{(i-1)}$ which occurs in $F^{(i)}$. If this is not the case, then a fix-up procedure is invoked (see Fig. 6). This fix-up procedure changes some of the associations until all the output expressions up to that point are covered. After **Compute-Cover** has terminated, all the output expressions will be covered, so W is the desired cover, and $|P| \geq |W| + |\bar{E}|$. If the straight-line program P is restricted to being cancellation-free, the fix-up procedure will never be necessary; it is only called if an output line was produced as a linear combination of two lines, where at one of those lines a cancelled variable was added to the cover, W .

The remainder of the proof first establishes the precise conditions under which the fix-up procedure is called, and then describes the action taken. We first define the two properties that **Compute-Cover** seeks to establish for each line l of the program.

Property 1. If line l is not an output, either

- all variables in $F^{(l)}$ are in $W^{(l)}$, or
- $m(l) = x \in F^{(l)}$ and $m(i) \neq x$ for all $i < l$.

Property 2. There is at most one variable in $F^{(l)}$ which is not in $W^{(l)}$.

In terms of these two properties, **Compute-Cover** (Fig. 4) can be described as follows. Given that Properties 1 and 2 hold for lines 1 to $i - 1$, establish Property 1 for line i and check if Property 2 holds for line i . If not, the fix-up procedure will be called. For all j , the size of $W^{(j)}$ will never decrease.

It will be convenient to assume each line of the program is of the form

$$v_i := \eta \cdot v_{i'} + \mu \cdot v_{i''} \quad (i > 0, \eta, \mu \neq 0).$$

To do this, we define $v_0 = z$, $v_{-i} = x_i$ for each variable x_i , and $m(i) = \lambda$ for all $i \leq 0$. We note that Property 2 holds for these initial assignments which occur before the first

$v_{-4} := d$	$F^{(-4)} = d$	$m(-4) = \lambda$
$v_{-3} := c$	$F^{(-3)} = c$	$m(-3) = \lambda$
$v_{-2} := b$	$F^{(-2)} = b$	$m(-2) = \lambda$
$v_{-1} := a$	$F^{(-1)} = a$	$m(-1) = \lambda$
$v_0 := z$	$F^{(0)} = z$	$m(0) = \lambda$
$v_1 := v_{-1} + v_{-2}$	$F^{(1)} = a + b$	$m(1) = b$
$v_2 := v_{-3} + v_{-4}$	$F^{(2)} = c + d$	$m(2) = c$
$v_3 := v_{-2} + v_{-3}$	$F^{(3)} = b + c$	$m(3) = \lambda$
$v_4 := v_0 + v_1$	$F^{(4)} = z + a + b$	$m(4) = \lambda$
$v_5 := v_3 + v_4$	$F^{(5)} = z + a + c$	$m(5) = \lambda$
$v_6 := v_2 + v_5$	$F^{(6)} = z + a + d$	$m(6) = \lambda$

Fig. 5. The last three lines are outputs. When the last line is reached, neither a nor d is in $W^{(6)}$, so fix-up is called.

“line” of the program P . Line 1 of P contains at most two variables and cannot be an output. Thus, Compute-Cover assigns one of these variables to $m(1)$. So, the first time Compute-Cover processes line 1, it establishes Properties 1 and 2 for line 1.

Claim 1. *Let line l be a non-output line and assume Property 2 holds for all lines before l . If Property 1 holds for line l , then Property 2 also holds for line l .*

Proof. If all variables in line l are in $W^{(l-1)}$, both properties hold for line l . Otherwise, let non-output line l be $v_l := \eta \cdot v_{l'} + \mu \cdot v_{l''}$. By assumption, Property 2 holds for $F^{(l')}$ and $F^{(l'')}$. Since $W^{(l')} \cup W^{(l'')} \subseteq W^{(l-1)}$, there are at most two variables in $F^{(l)}$ but not in $W^{(l-1)}$. By Property 1, $m(l) \in F^{(l)} \setminus W^{(l-1)}$. Thus Property 2 holds for line l . \square

Claim 1 implies that, prior to the first call to $\text{Fix-up}(i, i)$, establishing Property 1 also establishes Property 2. Compute-Cover explicitly establishes Property 1 at non-output lines. We have already argued that Properties 1 and 2 hold immediately after Compute-Cover processes line 1. It follows that the first time $\text{Fix-up}(i, i)$ is called, Property 1 holds for lines 1 through i (vacuously for line i) and Property 2 for all lines prior to i .

We now consider the way the algorithm processes output lines. It is not obvious that the fix-up procedure can ever be called. Figure 5 shows an example where this happens.

Claim 2. *Before a call to $\text{Fix-up}(s, t)$, either directly from Compute-Cover or recursively from fix-up, Property 1 holds for lines 1 through t and Property 2 holds for all lines before line s .*

Proof. Consider a call to $\text{Fix-up}(s, t)$. We have already shown the claim holds if this is the first call to fix-up. If this is not the first call to fix-up, we may assume, inductively, that the claim holds for all previous calls. Suppose the previous call was to $\text{Fix-up}(i, l)$ for an output line $i \leq l$:

$$v_i := \eta \cdot v_{i'} + \mu \cdot v_{i''} \quad (\eta, \mu \neq 0)$$

which produces the expression $z + a + b$. This call is caused by neither a nor b being in $W^{(i)}$.

Without loss of generality, assume that $i' < i''$. If both a and b are in $F^{(i')}$ or both are in $F^{(i'')}$, then, by the induction hypothesis, at least one of a or b is in $W^{(i)}$, a contradiction. Thus, neither line contains both a and b . We may assume, without loss of generality, that a is not present in line i' and b is not present in line i'' .

Since

$$\eta \cdot F^{(i')} + \mu \cdot F^{(i'')} = z + a + b$$

we have that

$$\begin{aligned} F^{(i')} &= v_1 z + (b/\eta) - (H/\eta), \\ F^{(i'')} &= v_2 z + (a/\mu) + (H/\mu), \end{aligned}$$

where $\eta v_1 + \mu v_2 = 1$ and H is a polynomial with variables in $X \setminus \{a, b, z\}$.

We now show that line i' is a non-output line and line i'' is an output line. Since $b \notin W^{(i')}$, Property 2 implies that all variables in H are in $W^{(i')}$. Therefore, all variables in H are in $W^{(i''-1)}$ and, by assumption, $a \notin W^{(i''-1)}$. If line i'' is not an output, Property 1 implies $m(i'') = a$. This contradicts $a \notin W^{(i)}$. Thus line i'' is an output line. Since no three distinct z expressions are linearly dependent, and lines i and i'' are output lines, it follows that line i' is a non-output line.

Thus, $F^{(i'')} = z + a + c$ where c is a variable distinct from a and b . It follows that

$$v_2 = 1; \quad \mu = 1; \quad H = c; \quad v_1 = 0; \quad i' > 0.$$

Therefore,

$$\begin{aligned} F^{(i')} &= (b - c)/\eta, \\ F^{(i'')} &= z + a + c. \end{aligned}$$

Inductively, Property 1 holds for lines 1 through l , so $m(i') = c$.

The fix-up procedure, as defined in Fig. 6, backs up to line i' , changes the mapping m so $m(i') = b$, not c , and then calls Correct to scan forward from i to check if there is a later line j where $m(j) = b$. Such an occurrence is changed, since no variable can be assigned to two different lines without violating Property 1 at the second such line. If such a line j is found, there are two cases to consider: (1) there is another variable x in $F^{(j)}$ which is not in $W^{(j)}$, and (2) all variables in $F^{(j)}$ are already covered. In the first case, the variable x is assigned to $m(j)$, which could cause another violation of Property 1 if x is also assigned at a later line. Thus, Correct is called recursively from line j to check for x being assigned later. This x could only be assigned to one later line, so Correct eventually terminates, ensuring that no two lines are assigned the same variable. In the second case, the variable c is assigned to line j ; it cannot have been assigned elsewhere, since it had only been assigned to line i' before the call to Fix-up(i, l).

The next loop in the fix-up procedure ensures that Property 1 still holds up through line l . If this property does not hold at some line, it is because some variable x (either

Fix-up(i, l)

{ i is the current line being fixed; l is original line being fixed }

{ line i , $v_i := \eta v_{i'} + v_{i''}$, produces the expression $z + a + b$,
 a is not present in line i' and b is not present in i'' }

{ line i' is not an output, $m(i') = c$, and $F(i') = (b - c)/\eta$ }

{ line i'' is an output and $F(i'') = z + a + c$ }

$m(i') \leftarrow b$

{ Check if b has been assigned to a later line }

Correct(b, i, l, c)

{ Check that Property 1 holds everywhere }

for $j \leftarrow i'$ **to** l **do**

if $F^{(j)}$ is not an output and $\exists x \in F^{(j)} \setminus W^{(j)}$ **then**

if ($m(j) = \lambda$) or ($m(j) \notin F^{(j)}$) **then**

$m(j) \leftarrow x$

for $k = j + 1$ **to** l **do**

if $m(k) = x$ **then** $m(k) \leftarrow \lambda$

 { Check that Property 2 holds everywhere }

for $j \leftarrow 1$ **to** l **do**

if $F^{(j)}$ is an output **then**

if $|F^{(j)} \setminus W^{(j)}| = 2$ **then** **Fix-up**(j, l); **return**;

Correct(y, i, l, c)

{ The variable y has just been assigned to a line. }

{ Check if y is assigned to between lines $i + 1$ and l . }

{ If so, assign the variable c or some variable in that line. }

$j \leftarrow i + 1$

while $j \leq l$ and $m(j) \neq y$ **do** $j \leftarrow j + 1$

if $j \leq l$ { i.e. $m(j) = y$ } **then**

if \exists a variable $x \in F^{(j)} \setminus W^{(j)}$ **then**

$m(j) \leftarrow x$

if $x \neq c$ **then** **Correct**(x, j, l, c)

else

$m(j) \leftarrow c$

Fig. 6. The fix-up procedure.

c or some other variable replaced in **Correct**) has been removed from some $W^{(k)}$, but $x \in F^{(k)}$. There are two possible cases here: (1) $m(k) = \lambda$, and (2) $m(k) = d$, where $d \notin F^{(k)}$. The first case is easy, and $m(k)$ is set to x . The second case could only have arisen from an earlier call to the fix-up procedure, at a point where $m(k)$ was set to the variable d because all of its variables were covered by $W^{(k)}$ and the variable d had been removed from an earlier line. In this case, we switch the assignment from d to x . If x was assigned to a later line (when correcting for a b being assigned to a later line and finding a line where all the variables were already covered), that assignment is removed.

(Note that this does not decrease the size of any $W^{(j)}$ since x is added to them when $m(k)$ gets the value x .) Thus, Property 1 holds up through line l .

The removal of c from $W^{(i')}$ may also cause Property 2 to fail. By Claim 1, the corrections for Property 1 in the previous loop ensure that the first failure for Property 2 is not at a non-output line. Some of the failures at output lines may be rectified by the adjustments fixing Property 1. “Fix-up” is called recursively to fix the others. If $\text{Fix-up}(s, t)$ is the first recursive call within $\text{Fix-up}(i, l)$ then

- the first loop in Fix-up ensures Property 1 holds through line $l = t$;
- the first failure of Property 2 must be at an output line, and therefore Property 2 holds through line $s - 1$.

Otherwise, $\text{Fix-up}(i, l)$ terminates before the call to $\text{Fix-up}(s, t)$. In this case Properties 1 and 2 hold up through line l (hence the **return** statement after one recursive call to Fix-up). Thus the call to $\text{Fix-up}(s, t)$ occurs in the code of Fig. 4 and $s = t > l$. From lines $l + 1$ through t there is no call to Fix-up . Hence, by Claim 1, the steps in the main loop (of Fig. 4) ensuring Property 1 also ensure Property 2 up through line $t - 1$. Property 1 holds vacuously at line t . \square

Finally, we note that the last call to Fix-up , and the remaining iterations of the loop in Fig. 4 ensure that Properties 1 and 2 hold everywhere. Thus, if the algorithm terminates, Property 2 will hold for all lines of P , and therefore $W = W^{(|P|)}$ is a cover of size at most $|P| - |\bar{E}|$.

We now turn to the proof of termination.

Claim 3. *A call to Fix-up never decreases the size of any $W^{(j)}$.*

Proof. There are two ways Fix-up appears to decrease the size of some $W^{(j)}$. The first is by swapping c for b at line i' where b is already assigned to some other line $j > i$. If the procedure terminates at this point, this would decrease the size of $W^{(j')}$ for $j' \geq j$. However, the procedure $\text{Correct}(b, i, l, c)$ checks for this and assigns some other variable, x or c , to line j and corrects for x recursively. The starting line for the search by Correct is larger for each recursive call, so eventually it terminates, adding a new variable to the cover which has not been assigned to a later line. Since c was originally assigned to line i' , it was not assigned to any other line, so this corrects the temporary decrease in $W^{(j')}$ for $j' \geq j$.

The second apparent decrease in the size of some $W^{(j)}$ does not actually ever create a decrease. During the check for Property 1 still holding, if $m(j)$ is set to x , but x is assigned to some later line k , then $m(k)$ is set to λ . (Note that this x may be the variable c which was removed at line i' , but it could be some other variable if some $m(j)$ had recently been set to λ .) However, adding x to $W^{(j)}$ also added it to $W^{(k)}$, so there is no actual decrease in the size of any $W^{(k')}$. \square

Let k_1, k_2, \dots be the sequence of line numbers for output lines which require a call to the fix-up procedure, and let $W_i^{(j)}$ denote the set $W^{(j)}$ at the point just before the fix-up procedure is called for line k_i .

Note that no two adjacent members of k_1, k_2, \dots are equal.

Let j be an index for which $k_j < k_{j+1}$ (if no such index exists, the sequence is clearly finite and this terminates). We claim that $|W_j^{(k_j)}| < |W_{j+1}^{(k_j)}|$. By the previous claim, the size of the cover never decreases. Thus, the claim follows if we show that a variable is added to the cover by the fix-up procedure when going from line k_j to k_{j+1} .

Consider how the fix-up procedure operates between the calls at lines k_j and k_{j+1} . Suppose that line k_j is

$$v_{k_j} := \eta \cdot v_{k'_j} + \mu \cdot v_{k''_j}.$$

We know that $k'_j < k''_j < k_j < k_{j+1}$. Suppose the formal expressions computed at these lines are

$$\begin{aligned} F^{(k'_j)} &= (b - c)/\eta; & F^{(k''_j)} &= z + a + c, \\ F^{(k_j)} &= z + a + b; & F^{(k_{j+1})} &= \dots \end{aligned}$$

For line k_j to have caused a call to “Fix-up”, neither a nor b could have been in the cover $W_j^{(k_j)}$. Thus the algorithm first visited line k'_j and changed the mapping $m(k'_j)$ from c to b , then executed the first “for” loop, correcting lines not satisfying Property 1, and finally moved down the program, checking each line for Property 2, until reaching line k_{j+1} . But this means that Property 2 held at line k''_j , and this could only have happened if a or c was in the cover. Since neither of them was in the cover immediately after the swap of b for c at line k'_j , one of them must have been added by the fix-up procedure at one of the lines in between. Thus $|W_j^{(k_j)}| < |W_{j+1}^{(k_j)}|$.

Hence for each j where $k_j < k_{j+1}$, the size of the cover at some line increases. Let n be the length of the program. Since all k_j are positive, there can be at most n calls to the fix-up procedure before some $W^{(k_j)}$ increases in size. Other than the time required for a possible recursive call, each call to the fix-up procedure is linear in n . From the bound $|W^{(k_j)}| < |X| \leq n$, for $1 \leq j \leq n$, it follows that Compute-Cover requires at most $O(n^4)$ time. (The fact that the execution time is polynomial is irrelevant for the purposes of showing NP-hardness, but will be important later.) This completes the proof of Lemma 1. \square

The following theorem follows immediately, since we have given a polynomial time reduction from VERTEX COVER, which is NP-complete.

Theorem 1. *For any field \mathbb{F} , SHORTEST LINEAR PROGRAM is NP-hard.*

For finite fields, it is easy to see that SLPd \in NP. Thus we have the following.

Theorem 2. *For any finite field \mathbb{F} , the decision version of SHORTEST LINEAR PROGRAM is NP-complete.*

Note that in the proof of Lemma 1, if the straight-line program P had been restricted to be cancellation-free, the proof would have been easier, because the fix-up procedure would never be necessary; it is only called if an output line was produced as a linear

combination of two lines, where at one of those lines a cancelled variable was added to the cover, W . This immediately gives us the following.

Theorem 3. *For any finite field \mathbb{F} , SHORTEST LINEAR PROGRAM is NP-complete even if the programs produced are restricted to being cancellation-free.*

3.1.2. Limits to Approximation

The major result of the previous subsection is that it is NP-hard to find an optimal linear program for computing a set of linear forms. Thus, it is natural to turn our attention to approximation algorithms for this problem. Here we concentrate entirely on polynomial-time approximation algorithms with provable performance guarantees.

We show that SHORTEST LINEAR PROGRAM has no ϵ -approximation scheme unless $P = NP$. Recall that these are families of algorithms, one for each $\epsilon > 0$, which are polynomial time and achieve an approximation ratio of $1 + \epsilon$. We use a concept called MAX SNP-completeness, which was introduced by Papadimitriou and Yannakakis [32]. Arora et al. [1] have shown that no MAX SNP-complete problem has an ϵ -approximation scheme unless $P = NP$. We show that BOUNDED Z-EXPONENT (defined below), is MAX SNP-complete, showing that there is no ϵ -approximation scheme for SHORTEST LINEAR PROGRAM unless $P = NP$, since it is a generalization of BOUNDED Z-EXPONENT.

MAX SNP is a complexity class of optimization problems. It is contained within NP in the sense that the decision versions of the problems are all in NP. Papadimitriou and Yannakakis [32] proved that many problems are MAX SNP-complete, including the following: BOUNDED VERTEX COVER: Given a graph with maximum vertex degree bounded by a constant b , find the smallest vertex cover.

To talk about completeness for this class, we need a notion of reduction. The reductions Papadimitriou and Yannakakis defined, called L -reductions, preserve the existence of ϵ -approximation schemes. The following definitions and propositions are taken directly from the original paper.

Let Π and Π' be two optimization (maximization or minimization) problems, and let f be a polynomial-time transformation from problem Π to problem Π' . We say that f is an L -reduction if there are constants $\alpha, \beta > 0$ such that, for each instance I of Π , the following two properties are satisfied:

- (a) The optima of I and $f(I)$, written $\text{OPT}(I)$ and $\text{OPT}(f(I))$ respectively, satisfy the relation $\text{OPT}(f(I)) \leq \alpha \text{OPT}(I)$.
- (b) For any solution of $f(I)$ with cost c' , we can find in polynomial time a solution of I with cost c such that $|c - \text{OPT}(I)| \leq \beta |c' - \text{OPT}(f(I))|$.

The constant β will usually be 1. The following two propositions, stated in [32], follow easily from the definition.

Proposition 1. *L -reductions compose.*

Proposition 2. *If Π L -reduces to Π' and if there is a polynomial-time approximation algorithm for Π' with worst-case error ϵ , then there is a polynomial-time approximation algorithm for Π with worst-case error $\alpha\beta\epsilon$.*

BOUNDED Z-EXPN is the following problem: Given a set of z -expressions (as defined in Theorem 1) in which each non- z variable appears at most b times (b is a fixed constant), generate an optimal linear program for computing the expressions (over some fixed field \mathbb{F}).

Theorem 4. *BOUNDED Z-EXPN is MAX SNP-complete.*

Proof. First, we will show that BOUNDED Z-EXPN is in MAX SNP. To show membership in MAX SNP, we will exhibit an L-reduction of BOUNDED Z-EXPN to Bounded Vertex Cover, a problem in MAX SNP.

For every non- z variable x_i , we associate a vertex \bar{x}_i . The L-reduction f maps z -expressions to edges as follows: $f("z + x_i + x_j") = \text{"edge } (i, j)\text{"}$. Since variable occurrences are bounded by b in BOUNDED Z-EXPN, the vertex degrees will be bounded by b in the graph.

We proved in the previous section that a set of z -expressions can be optimally computed by first computing $z + x_i$ for those x_i which are in the minimum vertex cover, and then using these intermediate results to compute the z -expressions. Thus $\text{OPT}(f(I)) + |E| = \text{OPT}(I)$ where $|E|$ is both the number of z -expressions and the number of edges in the graph.

We claim that this reduction is an L-reduction. Property (a) is satisfied because the equation above implies that $\text{OPT}(f(I)) \leq \text{OPT}(I)$. Property (b) is satisfied because, from a vertex cover, we can build a linear program which computes the z -expressions in the manner described above. This gives $c = \text{OPT}(I) + |c' - \text{OPT}(f(I))|$.

To show that the problem is MAX SNP-hard we reverse the reduction so that it goes from Bounded Vertex Cover to Bounded Z-EXPN. The function f now maps "edge (i, j) " into " $z + x_i + x_j$ ".

Proof of Property (a): By Lemma 1 we have that $\text{OPT}(I) + |E| = \text{OPT}(f(I))$. Since the maximum degree in the graph is bounded by b and every edge must be adjacent to at least one vertex of the cover, there can be at most $b \cdot \text{OPT}(I)$ edges, of the cover. Thus $\text{OPT}(f(I)) \leq (b + 1)\text{OPT}(I)$.

Proof of Property (b): The proof of Lemma 1 gave a polynomial-time procedure for converting any linear program computing a set of z -expressions into a vertex cover for the corresponding graph. By inspecting this procedure, one sees that $c = \text{OPT}(I) + |c' - \text{OPT}(f(I))|$. □

The fact that BOUNDED Z-EXPN is complete for the class MAX SNP implies that there is no ϵ -approximation scheme for it unless $P = NP$. In fact, Clementi and Trevisan [14] have shown that BOUNDED VERTEX COVER is not approximable within $16/15 - \epsilon$ for sufficiently large maximum degree. By Proposition 2, this means that there is no $1 + (1/15 - \epsilon)/\alpha\beta = 1 + (1/15 - \epsilon)/(1 + b)$ -approximation algorithm for SLP unless $P = NP$. We also mention that, assuming the Unique Games Conjecture [24], Austrin et al. have improved this inapproximability bound to a function that approaches 2 as the bound on the degree gets large [2].

The fact that BOUNDED Z-EXPN is in the class MAX SNP means that there is an approximation algorithm for it with a constant approximation ratio. In fact, it is obvious

that Z-EXPN, even without the boundedness constraint, has an approximation algorithm with a constant approximation. The straightforward linear straight-line program for computing the $|E|$ forms only requires $2|E|$ lines, and every straight-line program for E must contain at least $|E|$ lines (assuming no repetitions within the set E). Thus, the straightforward algorithm comes within a factor of 2 of optimal. Moreover, since there is an approximation algorithm for vertex cover which comes within a factor of 2 of optimal, we can do even better for Z-EXPN. Since the optimal linear program contains $|W| + |E|$ steps, where W is the minimum vertex cover, by Lemma 1, there is an algorithm which takes $2|W| + |E|$ steps. Since $|W| < |E|$, the ratio $(2|W| + |E|)/(|W| + |E|)$ is at most $3/2$, so there is a $(3/2)$ -approximation algorithm for Z-EXPN. There are, however, no known approximation algorithms which obtain a constant ratio for the general SLP problem.

3.1.3. Cancellation Can Yield Smaller Circuits

Thus, unless $P = NP$, this problem does not even have efficient ϵ -approximation schemes, so our goal in this research is restricted to improving on known heuristics. As far as we know, the most successful heuristics are variations on a greedy algorithm due to Paar [31]. We report significant improvements over the latter methods. Paar's algorithm gives non-cancelling results. It keeps a list of variables computed, which is initially only the inputs. Then it repeatedly determines which two variables, XORed together, occur in most outputs. One such pair is selected and XORed together. This result is added as a new variable which appears in all outputs where both variables previously appeared. This can be repeated until everything has been computed. One possible variant of this was presented in the same article [31]: When there is more than one most frequently occurring pair, instead of selecting one, try all possibilities, using recursion. The original algorithm is very fast; the variant is not.

A different technique is due to Bernstein [3]. Bernstein's algorithm has the advantages of using less storage and functioning better on two-operand platforms, i.e., where $a := a \oplus b$ is an allowed operation, but $a := b \oplus c$ is not. However, experiments mentioned in [3] indicate that Bernstein's algorithm usually produces results with more gates than Paar's.

Previous work on circuit minimization for AES S-boxes (e.g., [13,30,33]) only consider cancellation-free straight-line programs for producing a set of linear forms over $GF(2)$. Canright [13] even does an exhaustive search to find an optimal cancellation-free straight-line program. This does not, however, necessarily imply that Canright has found the optimal linear straight-line program. Some authors appear to make the incorrect assumption that there always exists a cancellation-free optimal linear program over $GF(2)$.

As mentioned in the introduction, restricting the search for optimal straight-line programs for computing linear forms over $GF(2)$ to cancellation-free programs can lead to suboptimal solutions. In our counter-example, the optimal cancellation-free program has length $\frac{5}{4}$ times that of the true shortest program. It is natural to ask how close to optimal cancellation-free programs can get as the number of variables increases. In this subsection we show that the best cancellation-free straight-line programs are not guaranteed to even have length within a factor $3/2$ that of the shortest straight-line linear program.

```

for  $i = 1$  to  $k(n - 1)$  do
     $u_i := x_i + x_{i+k}$ 
for  $i = 0$  to  $n - 2$  do
     $s_i := u_{ik+1} + u_{ik+2}$ 
    for  $j = 1$  to  $k - 2$  do
         $s_i := s_i + u_{ik+j+2}$ 
for  $i = 0$  to  $n - 3$  do
     $p_i := s_i + s_{i+1}$ 

```

Fig. 7. Straight-line program with cancellations.

The following construction uses two integer parameters k and n , which can be made large to make the $3/2$ inapproximability result hold asymptotically. The parameter k is the number of variables in a *block*, and n is the number of distinct blocks. Blocks have disjoint sets of variables: Block i , where $0 \leq i \leq n - 1$, is the linear form $b_i = x_{ik+1} + x_{ik+2} + \dots + x_{(i+1)k}$. The construction produces a linear straight-line program which is not cancellation-free. All intermediate linear forms (the linear forms produced at each line of the program) computed by this straight-line linear program will belong to the set of required outputs. The first part of the linear straight-line program will produce sums of consecutive pairs of blocks $s_i = b_i + b_{i+1}$, for $0 \leq i \leq n - 2$, mixing the variables in the two blocks in such a way that also producing a single block alone would require extra additions compared to the program here. Then, pairs of these consecutive sums are computed, $p_i = s_i + s_{i+1}$, for $0 \leq i \leq n - 3$. Each p_i is computed with only one further addition, but the two s_i s added share a common block which is cancelled, so $p_i = b_i + b_{i+2}$. We express this linear program, denoted P , using **for** loops in Fig. 7, but for any fixed k and n it is a straight-line program of length $k(n - 1) + (k - 1)(n - 1) + n - 2 = 2kn - 2k - 1$.

We claim that an optimal cancellation-free program (for computing all the linear forms which are the result of some line in this program) does at least enough additional operations to compute each of the blocks, and this would require at least $n(k - 1)$ additional lines. Let F denote the set consisting of the first $(2k - 1)(n - 1)$ lines of P , and let L denote the set of the last $n - 2$ lines. All of the $2kn - 2k - 1$ lines output by the above straight-line program are linear forms which must be output. The lines in L are the only ones with cancellations. None of the results from the lines in F can be used to compute the lines in P , because, for any two lines $f \in F$ and $l \in L$, f contains at least one variable which is not present in the form calculated by l . It is conceivable that some of the non-output results computed in the process of producing the outputs in L could be used in computing those in F , but, since they are all outputs, at least one extra operation is needed to produce each output from F . Thus, we can consider computing the outputs in L independently from those in F .

Blocks b_2 through b_{n-3} each appear in two of the outputs from L , but there is no other overlap between the outputs in L . Thus, the only reuse of forms computed which is possible is within the blocks. An optimal way to compute the forms in L is to first compute each of the n blocks, using $k - 1$ additions for each. After this, each form in L can be created by adding two blocks together, using one addition for each, as in P . The

computation of the blocks gives $n(k - 1)$ extra additions, for a total of $3kn - 2k - n - 1$ additions. Asymptotically, the ratio $\frac{3kn - 2k - n - 1}{2kn - 2k - 1}$ is $3/2$ for large n and k .

Theorem 5. *Any algorithm for computing short straight-line linear programs, which only produces cancellation-free straight-line programs, has an approximation ratio of at least $3/2$.*

Thus, even optimal cancellation-free circuits can be far from optimal in the unrestricted model. The heuristic we present below is not restricted to producing cancellation-free circuits. Furthermore, there appears to be little reason for restricting the search to cancellation-free circuits, as we have shown that finding an optimal cancellation-free circuit is also NP-hard in Sect. 3.1.1.

3.2. A New Heuristic

Let S be a set of linear functions. For any linear predicate f , we define the distance $\delta(S, f)$ as the minimum number of additions of elements from S necessary to obtain f .

The problem is to find a short linear program that computes $f(\mathbf{x}) = M\mathbf{x}$ where M is an $m \times n$ matrix over $\text{GF}(2)$. The heuristic is as follows. We keep a “base” S of “known” functions. Initially S is just the set of variables x_1, \dots, x_n . We maintain the vector $\text{Dist}[]$ of distances from S to the linear functions given by the rows of M . That is, $\text{Dist}[i] = \delta(S, f_i)$ where f_i is the i th row of M multiplied by the input vector \mathbf{x} . Initially, $\text{Dist}[i]$ is just one less than the Hamming weight of row i . We then perform the following loop:

- pick a new base element by adding two existing base elements;
- update $\text{Dist}[]$;

until $\text{Dist}[i] = 0$ for all i .

The current criterion for picking the new base element is

- pick one that minimizes the sum of new distances;
- resolve ties by *maximizing* the Euclidean norm of the vector of new distances.

This tie resolution criterion, which we term “Norm”, may seem counter-intuitive. The basic idea is that we prefer a distance vector like $0, 0, 3, 1$ to one like $1, 1, 1, 1$. In the latter case, we would need 4 more gates to finish. In the former, 3 might do it.

The bulk of the time of the heuristic is spent on picking the new base element. Our experiments show that the following “pre-emptive” choice usually improves running time without increasing the size of the output circuit:

- if any two bases $S[i], S[j]$ are such that $S[i] \oplus S[j]$ is a row in M , then pick this sum as the new base element.

The tie resolution criterion is a critical part of the heuristic. It does well on most matrices we have tried, but we have found specific matrices for which other decision rules do better. Intuitively, no one simple rule should work for all matrices. The effectiveness of the heuristic most likely depends on the topology of the digraph represented by the input matrix. We have not pursued this line of inquiry. We have, however, tested our

$$\begin{array}{l}
 y_0 = x_0 + x_1 + x_2 \\
 y_1 = x_1 + x_3 + x_4 \\
 y_2 = x_0 + x_2 + x_3 + x_4 \\
 y_3 = x_1 + x_2 + x_3 \\
 y_4 = x_0 + x_1 + x_3 \\
 y_5 = x_1 + x_2 + x_3 + x_4
 \end{array}
 \quad
 M =
 \begin{bmatrix}
 1 & 1 & 1 & 0 & 0 \\
 0 & 1 & 0 & 1 & 1 \\
 1 & 0 & 1 & 1 & 1 \\
 0 & 1 & 1 & 1 & 0 \\
 1 & 1 & 0 & 1 & 0 \\
 0 & 1 & 1 & 1 & 1
 \end{bmatrix}$$

Fig. 8. Example sequence of equations and corresponding matrix.

heuristic with various tie resolution methods against Paar’s algorithm [31]. On random matrices, our heuristic gives significant improvements under Norm as well as under three other tie-breaking rules (see Sect. 5).

The distance vector in our heuristics is computed by exhaustive search. The reason the heuristic is practical for moderate-size matrices is that the distance can only decrease. In fact, it can only decrease by 1. So when a new base is being considered, if a distance is d , then only combinations of exactly $d - 1$ old base elements and the new base element need to be considered.

3.3. A Small Example Using the Heuristic

Suppose we need a circuit that computes the system of equations defined in Fig. 8, which is equivalent to finding a circuit for multiplication by the 6×5 matrix, M , given in the figure.

The *target signals* to be computed are simply the rows of M . The initial base is $\{x_0, x_1, x_2, x_3, x_4\}$, which corresponds to

$$S = \{ [1 \ 0 \ 0 \ 0 \ 0], [0 \ 1 \ 0 \ 0 \ 0], [0 \ 0 \ 1 \ 0 \ 0], \\
 [0 \ 0 \ 0 \ 1 \ 0], [0 \ 0 \ 0 \ 0 \ 1] \}.$$

The initial distance vector is

$$D = [2 \ 2 \ 3 \ 2 \ 2 \ 3].$$

The heuristic must find two base vectors whose sum, when added to the base, minimizes the sum of the new distances. It turns out that the right choice is to calculate $x_1 + x_3$. So the new base S is expanded to contain the signal

$$[0 \ 1 \ 0 \ 1 \ 0] = [0 \ 1 \ 0 \ 0 \ 0] + [0 \ 0 \ 0 \ 1 \ 0].$$

The new distance vector is

$$D = [2 \ 1 \ 3 \ 1 \ 1 \ 2].$$

The full run of the program is shown in Fig. 9. The tie-breaking criteria is used in Step 3. If one had chosen $x_0 + x_1$ instead of $x_4 + t_6$, the new distance vector would be $[1 \ 1 \ 2 \ 0 \ 1 \ 1]$, which has norm $\sqrt{8}$, while the one found has norm $\sqrt{10}$. Note that there is cancellation in Steps 6 and 8.

Step 1: $t_5 = x_1 + x_3$ (found signal = [0 1 0 1 0]). New D : [2 1 3 1 1 2]
 Step 2: $t_6 = x_2 + t_5$ (found target signal $y_3 = [0 1 1 1 0]$). New D : [2 1 3 0 1 1]
 Step 3: $t_7 = x_4 + t_6$ (found target signal $y_5 = [0 1 1 1 1]$). New D : [2 1 2 0 1 0]
 Step 4: $t_8 = x_0 + x_1$ (found signal = [1 1 0 0 0]). New D : [1 1 1 0 1 0]
 Step 5: $t_9 = x_0 + t_5$ (found target signal $y_4 = [1 1 0 1 0]$). New D : [1 1 1 0 0 0]
 Step 6: $t_{10} = x_2 + t_7$ (found target signal $y_1 = [0 1 0 1 1]$). New D : [1 0 1 0 0 0]
 Step 7: $t_{11} = x_2 + t_8$ (found target signal $y_0 = [1 1 1 0 0]$). New D : [0 0 1 0 0 0]
 Step 8: $t_{12} = t_7 + t_8$ (found target signal $y_2 = [1 0 1 1 1]$). New D : [0 0 0 0 0 0]
 (DONE!)

Fig. 9. Example running heuristic for minimizing linear components.

$y_{14} = x_3 + x_5$	$y_{13} = x_0 + x_6$	$y_9 = x_0 + x_3$
$y_8 = x_0 + x_5$	$t_0 = x_1 + x_2$	$y_1 = t_0 + x_7$
$y_4 = y_1 + x_3$	$y_{12} = y_{13} + y_{14}$	$y_2 = y_1 + x_0$
$y_5 = y_1 + x_6$	$y_3 = y_5 + y_8$	$t_1 = x_4 + y_{12}$
$y_{15} = t_1 + x_5$	$y_{20} = t_1 + x_1$	$y_6 = y_{15} + x_7$
$y_{10} = y_{15} + t_0$	$y_{11} = y_{20} + y_9$	$y_7 = x_7 + y_{11}$
$y_{17} = y_{10} + y_{11}$	$y_{19} = y_{10} + y_8$	$y_{16} = t_0 + y_{11}$
$y_{21} = y_{13} + y_{16}$	$y_{18} = x_0 + y_{16}$	

Fig. 10. Top linear transformation: Inputs are x_0, x_1, \dots, x_7 . Outputs to the next level are $x_7, y_1, y_2, \dots, y_{21}$.

Thus, after the x_i , which may be nonlinear functions of other variables, are computed, the y_i are computed by following the algorithm produced and, in this case, letting $y_0 = t_{11}, y_1 = t_{10}, y_2 = t_{12}, y_3 = t_6, y_4 = t_9, y_5 = t_7$.

4. A Circuit for the S-Box of AES

Our techniques yield a circuit for the AES S-box composed of 115 gates in three parts: a “top” linear transformation, U ; a middle nonlinear part; and a “bottom” linear transformation, B . The linear transformations are defined by the matrices U and B of Sect. 2.2.

For the matrix U , the smallest circuits we found had $23 \oplus$ gates. Among the many such circuits, the shortest ones have depth 7. It is worthwhile to note that if $24 \oplus$ gates are allowed, circuits with depth 4 exist for U . Figure 10 shows a circuit of size 23 and depth 7. The circuit maps inputs x_0, \dots, x_7 to outputs x_7, y_1, \dots, y_{21} .

Figure 11 shows the nonlinear middle part of the S-box circuit. It is a function from 22 to 18 bits. The circuit contains $32 \wedge$ gates and $30 \oplus$ gates. It maps inputs x_7, y_1, \dots, y_{21} to outputs z_0, \dots, z_{17} .

For matrix B , the randomized version of our heuristic yields many circuits with $30 \oplus$ gates. The heuristic is fast enough that we are able to pick a circuit which is both small and short. Figure 12 shows a circuit of depth 6. The circuit maps inputs z_0, \dots, z_{17} to outputs s_0, \dots, s_7 .

$t_2 = y_{12} \times y_{15}$	$t_3 = y_3 \times y_6$	$t_4 = t_3 + t_2$
$t_5 = y_4 \times x_7$	$t_6 = t_5 + t_2$	$t_7 = y_{13} \times y_{16}$
$t_8 = y_5 \times y_1$	$t_9 = t_8 + t_7$	$t_{10} = y_2 \times y_7$
$t_{11} = t_{10} + t_7$	$t_{12} = y_9 \times y_{11}$	$t_{13} = y_{14} \times y_{17}$
$t_{14} = t_{13} + t_{12}$	$t_{15} = y_8 \times y_{10}$	$t_{16} = t_{15} + t_{12}$
$t_{17} = t_4 + t_{14}$	$t_{18} = t_6 + t_{16}$	$t_{19} = t_9 + t_{14}$
$t_{20} = t_{11} + t_{16}$	$t_{21} = t_{17} + y_{20}$	$t_{22} = t_{18} + y_{19}$
$t_{23} = t_{19} + y_{21}$	$t_{24} = t_{20} + y_{18}$	
$t_{25} = t_{21} + t_{22}$	$t_{26} = t_{21} \times t_{23}$	$t_{27} = t_{24} + t_{26}$
$t_{28} = t_{25} \times t_{27}$	$t_{29} = t_{28} + t_{22}$	$t_{30} = t_{23} + t_{24}$
$t_{31} = t_{22} + t_{26}$	$t_{32} = t_{31} \times t_{30}$	$t_{33} = t_{32} + t_{24}$
$t_{34} = t_{23} + t_{33}$	$t_{35} = t_{27} + t_{33}$	$t_{36} = t_{24} \times t_{35}$
$t_{37} = t_{36} + t_{34}$	$t_{38} = t_{27} + t_{36}$	$t_{39} = t_{29} \times t_{38}$
$t_{40} = t_{25} + t_{39}$		
$t_{41} = t_{40} + t_{37}$	$t_{42} = t_{29} + t_{33}$	$t_{43} = t_{29} + t_{40}$
$t_{44} = t_{33} + t_{37}$	$t_{45} = t_{42} + t_{41}$	$z_0 = t_{44} \times y_{15}$
$z_1 = t_{37} \times y_6$	$z_2 = t_{33} \times x_7$	$z_3 = t_{43} \times y_{16}$
$z_4 = t_{40} \times y_1$	$z_5 = t_{29} \times y_7$	$z_6 = t_{42} \times y_{11}$
$z_7 = t_{45} \times y_{17}$	$z_8 = t_{41} \times y_{10}$	$z_9 = t_{44} \times y_{12}$
$z_{10} = t_{37} \times y_3$	$z_{11} = t_{33} \times y_4$	$z_{12} = t_{43} \times y_{13}$
$z_{13} = t_{40} \times y_5$	$z_{14} = t_{29} \times y_2$	$z_{15} = t_{42} \times y_9$
$z_{16} = t_{45} \times y_{14}$	$z_{17} = t_{41} \times y_8$	

Fig. 11. The middle nonlinear section: inputs are $x_7, y_1, y_2, \dots, y_{21}$. Outputs to the next level are z_0, z_1, \dots, z_{17} . Note that the computation of t_{25} through t_{40} is the inversion in $GF(2^4)$.

$t_{46} = z_{15} + z_{16}$	$t_{47} = z_{10} + z_{11}$	$t_{48} = z_5 + z_{13}$
$t_{49} = z_9 + z_{10}$	$t_{50} = z_2 + z_{12}$	$t_{51} = z_2 + z_5$
$t_{52} = z_7 + z_8$	$t_{53} = z_0 + z_3$	$t_{54} = z_6 + z_7$
$t_{55} = z_{16} + z_{17}$	$t_{56} = z_{12} + t_{48}$	$t_{57} = t_{50} + t_{53}$
$t_{58} = z_4 + t_{46}$	$t_{59} = z_3 + t_{54}$	$t_{60} = t_{46} + t_{57}$
$t_{61} = z_{14} + t_{57}$	$t_{62} = t_{52} + t_{58}$	$t_{63} = t_{49} + t_{58}$
$t_{64} = z_4 + t_{59}$	$t_{65} = t_{61} + t_{62}$	$t_{66} = z_1 + t_{63}$
$s_0 = t_{59} + t_{63}$	$s_6 = t_{56}$ XNOR t_{62}	$s_7 = t_{48}$ XNOR t_{60}
$t_{67} = t_{64} + t_{65}$	$s_3 = t_{53} + t_{66}$	$s_4 = t_{51} + t_{66}$
$s_5 = t_{47} + t_{65}$	$s_1 = t_{64}$ XNOR s_3	$s_2 = t_{55}$ XNOR t_{67}

Fig. 12. Bottom linear transformation: Inputs are z_0, z_1, \dots, z_{17} . Outputs are s_0, s_1, \dots, s_7 .

As mentioned earlier, our circuit was based on Canright's [13]. Our nonlinear middle part corresponds fairly closely to his, except that his subcircuit for inversion in $\text{GF}(2^4)$ was replaced by ours. He does not consider all of the top linear transformation as one unit. However, since this middle part of his circuit corresponds to ours, with the same inputs and outputs, he computes those inputs using linear operations. The number of XOR/XNOR gates he uses to compute this top linear transformation is 29. Similarly, he uses 31 XOR/XNOR gates to compute what corresponds to our bottom linear transformation. After optimizations, his circuit has a total of 80 XOR/XNOR gates, 34 NANDs, and 6 NORs. We did not attempt to use NOR gates to further reduce the size of our circuit.

A more direct comparison was also made comparing our techniques for minimizing linear circuits and Canright's. In [13], he presents a factorization of two 16 by 8 matrices, $(\frac{X^{-1}}{(MX)^{-1}})$ and $(\frac{(MX)}{X})$, showing that these can be computed using 20 and 18 gates, respectively. Our heuristic produces circuits with 18 and 17 gates, respectively.⁸

5. Experiments with Different Tie-Breaking Methods

In order to compare the effects of using different tie-breakers, we tested our heuristics on matrices generated as follows.

- We first chose a size (for example, 10×20 matrices, which represent 10 linear forms on 20 distinct variables).
- We then picked a *bias* ρ between 0 and 1.
- For each entry of the matrix, we set the bit to 1 with probability ρ and to 0 with probability $1 - \rho$. Thus ρ is the expected fraction of variables that appears in each linear form.
- Matrices with rows which are all zeros were discarded, as were matrices containing duplicate rows.

The testing was performed with a C++ program, compiled with `g++ -O3`, on a quad-core `x86_64`, running Ubuntu 9.10, with Intel Xenon 5150 processors running at 2.66 GHz, with 8 GB memory. There were no other users on the machine. The programs and matrices used can be found at www.imada.sdu.dk/~joan/xor/, though minor changes are necessary to run the programs with different files as input or to change the matrix size and bias for the matrix generator. We compared the different heuristics on sets of 100 random matrices with different sizes and densities. The experiments showed that the heuristics were slower when the bias was larger. This was expected, since the initial "distances" (number of operations on the base vectors to obtain the target vectors) were then larger on average when there were more ones in the matrices.

The tie-breakers we compared were the following:

- *Norm*: maximizing the Euclidean norm
- *Norm-largest*: maximizing the square of the Euclidean norm minus the largest distance

⁸ Using the Improved2.cc program with the matrix `canmat`, available with the other programs and matrices described in the next section, gives these results.

- *Norm-diff*: maximizing the square of the Euclidean norm minus the difference of the largest two distances
- *Random*: In processing the possible new base vectors, if the current possible new base vector has the same sum of distances as the previous best (current choice), then flip an unbiased coin. If heads, then keep the current choice. If tails, then apply the Norm criterion. This heuristic may end up choosing a pair with non-maximum Euclidean norm. On the other hand, it allows substitution of one optimum (by sum-of-distances and Euclidean norm) pair by another found later in the search.

In all cases, except the “Random” one, when there were still ties after applying the “tie-breaker,” the first pair with both the minimum sum of distances and the optimal value for the tie-breaker was chosen. This was the base pair with lexicographically minimum indices (i, j) . The exception to this is when there is a target with distance 1, meaning that using one extra gate will produce a target. A check is made for this case by scanning the distances and choosing the first with distance 1 when such exists. This check is efficient, and when there is a target of distance 1, it saves lengthy computations of new distances for each possible pair of bases.

Randomized tie-breaking allows running the heuristic several times and picking the best result. In our tests we ran the heuristic with “Random” tie-breaking three times.

We also compared these heuristics to Paar’s heuristic [31] on the same matrices. Paar’s heuristic repeatedly finds the most frequently occurring base pair and adds that as the next base pair. It is significantly faster than our heuristic, but it produces only cancellation-free circuits. Its performance, relative to the heuristics proposed here, decreases as the bias increases, using more than 30 % extra gates when the bias is $3/4$ (when the number of rows is at least 15) and 40 % extra when the bias is $9/10$.

Among the biases tried, the number of gates in the circuits found by our heuristics is similar with biases $1/2$ and $3/4$. It is not a strictly increasing function of the bias, since when nearly all of the variables are used in nearly all of the forms, the outputs from many of the gates can be reused for many targets. Thus, circuits with fewer gates were found when the bias was $9/10$ than when it was $1/2$ or $3/4$. This was also true for Paar’s heuristic, but less dramatically so.

All the tie resolution criteria performed fairly similarly, producing circuits of nearly the same size, with Random apparently doing slightly better (more often producing smaller circuits), presumably because it tries three different circuits and uses the best. Random also runs for about three times as long as the others. The results of these tests are presented in tables in the [Appendix](#). In the tables, the column headings specify the matrix size and the bias. For each heuristic, and all matrix sizes and biases, 100 randomly chosen matrices were tested.

For each tie-breaker rule and Paar’s heuristic, for each matrix size and bias, the average number of gates in the circuits found and the number of matrices where that heuristic did not obtain the minimum value of all of the heuristics was computed, along with the running time in seconds. The Paar heuristic was beaten by at least one of the other heuristics on all 700 matrices except for 17 of the 100 with bias $1/4$ (and there was only one matrix on which Paar’s heuristic beat any of the other heuristics). In fact, for the tests with bias larger than $1/4$, Paar’s heuristic did worse than any of the other heuristic on every one of the matrices; usually the values obtained for the newer heuris-

tics were similar, with Random possibly being marginally better, but with the value for Paar’s heuristic being significantly larger.

Paar’s heuristic (and, for matrices between size 4 and 10, a variant which does at most one gate better on average in the data presented) was tested [31] on square matrices of sizes 4×4 through 16×16 , and the average number of XOR gates is presented, along with the relative improvement over the straightforward implementation. These square matrices came from applying Mastrovito’s [27] matrix description of multiplication in $\text{GF}(2^n)$ to constant multiplication. Paar tries all possible constants in $\text{GF}(2^n)$ for n between 4 and 16, giving these square matrices. Since our heuristics are so much slower and the matrices in the cryptographic applications we are interested in do not necessarily have this form, we have not tested on all of these restricted matrices of those sizes, but rather on random matrices with different biases. For 15×15 matrices, Paar gets an average of 52.9 gates. This is similar to our results for Paar’s algorithm with 15×15 matrices with biases $1/2$ and $3/4$, where the Paar heuristic gets averages of 51.7 and 53.3 gates, respectively. For bias $1/2$, our deterministic heuristics get average gate counts between 44.21 and 44.28, while Random gets 43.81. For bias $3/4$, our deterministic heuristics all get average count 40.82, while Random gets 40.38. Thus, our relative improvement over the Paar heuristic is between 17 % and 32 % for these types of matrices. Paar’s result of 52.9 gates for 15×15 matrices is a relative improvement of 45.5 % over the straightforward approach.

The last row in each table in the [Appendix](#) shows the average of the values which are the minimum of those calculated by the different heuristics for each matrix. The fact that this number is always strictly smaller than the average for any specific tie-breaker shows that, for each of the tie-breakers, there are cases where it gets a worse result than at least one of the others. This is also shown by the column headed “Not min” which shows the number of matrices for which that tie-breaker did not achieve the lowest value found by the tie-breakers.

6. Conclusions and Work in Progress

We developed and tested new techniques for decreasing circuit size. The techniques were applied to the extensively studied AES S-box. We obtained the smallest circuit yet constructed for this function. The circuit contains 32 AND gates and 83 XOR/XNOR gates for a total of 115 gates. As by-products of the experiment, we obtained very small circuits for inversion in $\text{GF}(2^4)$ and $\text{GF}(2^8)$.

The result that SHORTEST LINEAR PROGRAM is NP-hard indicates that using heuristic techniques is more realistic than expecting to find the smallest subcircuits for linear parts of a Boolean circuit. The result that a special case of SHORTEST LINEAR PROGRAM is MAX SNP-complete indicates that there is a limit to how well these heuristic techniques can be guaranteed to perform.

Since cancellation-free techniques can produce linear straight-line programs which are a factor $3/2$ larger than the optimal, the heuristic developed here (in Step 2) is not restricted to cancellation-free operations.

The experiments with linear circuit optimization indicate that our techniques are likely to be superior to previous techniques which produced only cancellation-free circuits. We expect this to be particularly useful for cryptographic applications, both for

hardware and software implementations, where many XOR operations are used, along with some AND operations to introduce nonlinearity.

It would be interesting to determine how close to optimal the circuits found by these techniques usually are and how much better they are than the optimal cancellation-free circuits. Finding even better techniques which are not restricted to finding cancellation-free circuits would also be very interesting.

Work on finding exact solutions using SAT-solvers has developed a technique which will quickly find a circuit with 23 gates, the same size we report here for our techniques, for the top linear transformation [18,19]. They also prove that this cannot be achieved with 22 gates, so the number of gates used here for the top linear transformation is optimal.

Recent work has shown that the lower bound of $3/2$ for the approximation ratio of cancellation-free straight-line programs can be improved to 2, using a generalization of the example at the end of Sect. 1.3.

In practice, one would like to construct small low-depth circuits. This paper has discussed size only. However, it is plausible that a short circuit can be obtained by first minimizing size and then shortening the circuit along critical paths using balancing and other simple techniques. Preliminary results using this general approach are highly encouraging. An application to the AES S-box yields a circuit of depth 16 with only 128 gates [8]. This size is larger than that given here, where the depth is 28, but comparable to other results which have significantly more depth [13,28]. Previous attempts at reducing depth without too much expansion in size were only able to produce depth 22 and size 148 [29].

Acknowledgements

The authors would like to thank the anonymous referees for their suggestions. One of the referees also discovered a problem with an earlier proof of Lemma 1, which we correct here. The following summer interns at NIST also contributed to some of the experimental work reported here: Holman Gao and Michael Bartock.

Research of J. Boyar was partially supported by the Danish Council for Independent Research, Natural Sciences.

Appendix A. Experimental Results on Samples of 100 Random Matrices

Heuristic	15×15 matrices, Bias = $\frac{1}{4}$			15×15 matrices, Bias = $\frac{1}{2}$		
	Average	Not min	Seconds	Average	Not min	Seconds
Norm	29.65	16	12	44.21	48	125
Norm-largest	29.63	14	12	44.23	49	121
Norm-diff	29.65	15	11	44.28	51	119
Random	29.59	10	29	43.81	23	322
Paar	31.07	83	0.01	51.70	100	0.02
Minimum	29.48	0	–	43.50	0	–

Heuristic	15 × 15 matrices, Bias = $\frac{3}{4}$			15 × 15 matrices, Bias = $\frac{9}{10}$		
	Average	Not min	Seconds	Average	Not min	Seconds
Norm	40.82	47	291	30.28	31	388
Norm-largest	40.82	46	290	30.28	31	428
Norm-diff	40.82	46	292	30.29	32	388
Random	40.39	23	838	30.01	14	1145
Paar	53.27	100	0.03	43.11	100	0.02
Minimum	40.11	0	–	29.86	0	–

Heuristic	20 × 20 matrices, Bias = $\frac{3}{4}$		
	Average	Not min	Seconds
Norm	67.47	62	86,465
Norm-largest	67.43	60	82,597
Norm-diff	67.40	58	82,780
Random	66.87	30	234,815
Paar	90.86	100	0.11
Minimum	66.43	0	–

Heuristic	20 × 10 matrices, Bias = $\frac{3}{4}$			10 × 20 matrices, Bias = $\frac{3}{4}$		
	Average	Not min	Seconds	Average	Not min	Seconds
Norm	31.44	25	1.35	42.04	44	30,626
Norm-largest	31.43	24	1.38	42.08	44	30,490
Norm-diff	31.44	25	1.34	42.12	44	30,740
Random	31.23	11	4.08	41.76	22	84,540
Paar	43.32	100	0.02	50.02	100	0.02
Minimum	31.12	0	–	41.50	0	–

References

- [1] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy, Proof verification and the hardness of approximation problems. *J. Assoc. Comput. Mach.* **45**, 501–555 (1998)
- [2] P. Austrin, S. Khot, M. Safra, Inapproximability of vertex cover and independent set in bounded degree graphs, in *IEEE Conference on Computational Complexity* (IEEE Computer Society, Los Alamitos, 2009), pp. 74–80
- [3] D.J. Bernstein, Optimizing linear maps modulo 2, in *Workshop Record of SPEED-CC: Software Performance Enhancement for Encryption and Decryption and Cryptographic Compilers*. <http://cr.yp.to/papers.html#linearmod2>
- [4] L. Blum, M. Shub, S. Smale, On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bull. Am. Math. Soc.* **21**, 1–46 (1989)
- [5] J. Boyar, R. Peralta, Tight bounds for the multiplicative complexity of symmetric functions. *Theor. Comput. Sci.* **396**(1–3), 223–246 (2008)

- [6] J. Boyar, R. Peralta, Patent application number 61089998 filed with the U.S. Patent and Trademark Office. A new technique for combinational circuit optimization and a new circuit for the S-Box for AES, 2009
- [7] J. Boyar, R. Peralta, A new combinational logic minimization technique with applications to cryptology, in *9th International Symposium on Experimental Algorithms, SEA 2010*. Lecture Notes in Computer Science, vol. 6049 (Springer, Berlin, 2010), pp. 178–189
- [8] J. Boyar, R. Peralta, A depth-16 circuit for the AES S-box. Cryptology ePrint archive, report 2011/332, 2011. <http://eprint.iacr.org/>
- [9] J. Boyar, R. Peralta, D. Pochuev, On the multiplicative complexity of Boolean functions over the basis $(\wedge, \oplus, 1)$. *Theor. Comput. Sci.* **235**, 43–57 (2000)
- [10] J. Boyar, P. Matthews, R. Peralta, On the shortest linear straight-line program for computing linear forms, in *33rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2008*. Lecture Notes in Computer Science, vol. 5162 (Springer, Berlin, 2008), pp. 168–179
- [11] P. Bürgisser, M. Clausen, M.A. Shokrollahi, *Algebraic Complexity Theory* (Springer, Berlin, 1997), Chap. 13
- [12] D. Canright, A very compact Rijndael S-box. Technical report NPS-MA-05-001, Naval Postgraduate School, 2005
- [13] D. Canright, A very compact Rijndael S-box, in *7th International Workshop on Cryptographic Hardware and Embedded Systems, CHES 2005*. Lecture Notes in Computer Science, vol. 3659 (Springer, Berlin, 2005), pp. 441–455
- [14] A.E.F. Clementi, L. Trevisan, Improved non-approximability results for vertex cover with density constraints, in *Computing and Combinatorics* (1996), pp. 333–342
- [15] J.W. Cooley, J.W. Tukey, An algorithm for the machine calculation of complex Fourier series. *Math. Comput.* **19**, 297–301 (1965)
- [16] N. Courtois, D. Hulme, T. Mourouzis, Solving circuit optimisation problems in cryptography and cryptanalysis. *IACR Cryptology ePrint Archive*, 2011:475, 2011
- [17] FIPS, *Advanced Encryption Standard (AES)* (National Institute of Standards and Technology, Gaithersburg, 2001)
- [18] C. Fuhs, P. Schneider-Kamp, Synthesizing shortest linear straight-line programs over GF(2) using SAT, in *13th International Conference on Theory and Applications of Satisfiability Testing*. Lecture Notes in Computer Science, vol. 6175 (Springer, Berlin, 2010), pp. 71–84
- [19] C. Fuhs, P. Schneider-Kamp, Optimizing the AES S-Box using SAT, in *Proceedings of the 8th International Workshop on the Implementation of Logics* (2010)
- [20] J. Håstad, Tensor rank is NP-Complete. *J. Algorithms* **11**(4), 644–654 (1990)
- [21] Y. Huang, D. Evans, J. Katz, L. Malka, Faster secure two-party computation using garbled circuits, in *Proceedings of the 20th USENIX Security Symposium*, San Francisco, CA, August 2011
- [22] T. Itoh, S. Tsujii, A fast algorithm for computing multiplicative inverses in GF(2^m) using normal bases. *Inf. Comput.* **78**(3), 171–177 (1988)
- [23] E. Käsper, P. Schwabe, Faster and timing-attack resistant AES-GCM, in *11th International Workshop on Cryptographic Hardware and Embedded Systems, CHES 2009*. Lecture Notes in Computer Science, vol. 5747 (Springer, Berlin, 2009), pp. 1–17
- [24] S. Khot, On the power of unique 2-prover 1-round games, in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing, STOC '02*, New York, NY, USA (ACM, New York, 2002), pp. 767–775
- [25] V. Kolesnikov, T. Schneider, Improved garbled circuit: free XOR gates and applications, in *Proceedings of Automata, Languages and Programming, 35th International Colloquium, ICALP 2008*. Lecture Notes in Computer Science, vol. 5126 (Springer, Berlin, 2008), pp. 486–498
- [26] O.B. Lupanov, A method of circuit synthesis. *Izv. Vysš. Učebn. Zaved., Radiofiz.* **1**, 120–140 (1958)
- [27] E. Mastrovito, VLSI architectures for computation in Galois fields. Ph.D. thesis, Linköping University, Dept. Electr. Eng., Sweden, 1991
- [28] S. Morioka, A. Satoh, An optimized S-Box circuit architecture for low power AES design, in *Revised Papers from the 4th International Workshop on Cryptographic Hardware and Embedded Systems, CHES 2002*. Lecture Notes in Computer Science, vol. 2523 (Springer, Berlin, 2003), pp. 172–186
- [29] Y. Nogami, K. Nekado, T. Toyota, N. Hongo, Y. Morikawa, Mixed bases for efficient inversion in $f(((2^2)^2)^2)$ and conversion matrices of subbytes of AES, in *12th International Workshop on Crypto-*

- graphic Hardware and Embedded Systems, CHES 2010*. Lecture Notes in Computer Science, vol. 6225 (Springer, Berlin, 2010), pp. 234–247
- [30] C. Paar, Some remarks on efficient inversion in finite fields, in *1995 IEEE International Symposium on Information Theory*, Whistler, BC, Canada (1995), p. 58
 - [31] C. Paar, Optimized arithmetic for Reed-Solomon encoders, in *IEEE International Symposium on Information Theory* (1997), p. 250
 - [32] C. Papadimitriou, M. Yannakakis, Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.* **43**, 425–440 (1991)
 - [33] A. Satoh, S. Morioka, K. Takano, S. Munetoh, A compact Rijndael hardware architecture with S-Box optimization, in *Advances in Cryptology—Proceedings of ASIACRYPT 01*. Lecture Notes in Computer Science, vol. 2248 (Springer, Berlin, 2001), pp. 239–254
 - [34] J.E. Savage, An algorithm for the computation of linear forms. *SICOMP* **3**(2), 150–158 (1974)
 - [35] C. Shannon, The synthesis of two-terminal switching circuits. *Bell Syst. Tech. J.* **28**, 59–98 (1949)
 - [36] L.G. Valiant, Completeness classes in algebra, in *Proceedings of the 11th Annual ACM Symposium on the Theory of Computing* (1979), pp. 249–261
 - [37] R. Williams, Matrix-vector multiplication in sub-quadratic time (some preprocessing required), in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms* (2007), pp. 995–1001
 - [38] S. Winograd, On the number of multiplications necessary to compute certain functions. *Commun. Pure Appl. Math.* **23**, 165–179 (1970)
 - [39] J. Wolkerstorfer, E. Oswald, M. Lamberger, An ASIC implementation of AES SBoxes, in *Topics in Cryptology—CT-RSA 2002*. Lecture Notes in Computer Science, vol. 2271 (Springer, Berlin, 2002), pp. 67–78