

An Improved Affine Equivalence Algorithm for Random Permutations

Itai Dinur*

Department of Computer Science, Ben-Gurion University, Israel

Abstract. In this paper we study the affine equivalence problem, where given two functions $\mathbf{F}, \mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, the goal is to determine whether there exist invertible affine transformations A_1, A_2 over $GF(2)^n$ such that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Algorithms for this problem have several well-known applications in the design and analysis of Sboxes, cryptanalysis of white-box ciphers and breaking a generalized Even-Mansour scheme. We describe a new algorithm for the affine equivalence problem and focus on the variant where \mathbf{F}, \mathbf{G} are permutations over n -bit words, as it has the widest applicability. The complexity of our algorithm is about $n^3 2^n$ bit operations with very high probability whenever \mathbf{F} (or \mathbf{G}) is a random permutation. This improves upon the best known algorithms for this problem (published by Biryukov et al. at EUROCRYPT 2003), where the first algorithm has time complexity of $n^3 2^{2n}$ and the second has time complexity of about $n^3 2^{3n/2}$ and roughly the same memory complexity.

Our algorithm is based on a new structure (called a *rank table*) which is used to analyze particular algebraic properties of a function that remain invariant under invertible affine transformations. Besides its standard application in our new algorithm, the rank table is of independent interest and we discuss several of its additional potential applications.

Keywords: Affine equivalence problem, block cipher, Even-Mansour cipher, cryptanalysis, rank table.

1 Introduction

In the affine equivalence problem, the input consists of two functions \mathbf{F}, \mathbf{G} and the goal is to determine whether they are affine equivalent, and if so, output the equivalence relations. More precisely, if there exist invertible affine transformations (over some field) A_1, A_2 such that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$, output A_1, A_2 . Otherwise, assert that \mathbf{F}, \mathbf{G} are not affine equivalent.

Variants of the affine equivalence problem have been studied in several branches of mathematics and are relevant to both asymmetric and symmetric cryptography. In the context of asymmetric cryptography, the problem was first formalized by Patarin [16] and referred to as isomorphism of polynomials. In this setting \mathbf{F}, \mathbf{G} are typically of low algebraic degree (mainly quadratic) over some field.

* The author was supported in part by the Israeli Science Foundation through grant No. 573/16.

The focus of this work is on the affine equivalence variant in which \mathbf{F}, \mathbf{G} map between n -bit words and the affine transformations A_1, A_2 are over $GF(2)^n$. This variant is mostly relevant in several contexts of symmetric-key cryptography. In particular, it is relevant to the classification and analysis of Sboxes (see [6, 13]) as affine equivalent Sboxes share differential, linear and several algebraic properties (refer to [7] for recent results on this subject). Moreover, algorithms for the affine equivalence problem were applied in [3] to generate equivalent representations of AES and other block ciphers. These algorithms also have cryptanalytic applications and were used to break white-box ciphers (e.g., in [14]). Additionally, solving the affine equivalence problem can be viewed as breaking a generalization of the Even-Mansour scheme [10], which has received substantial attention from the cryptographic community in recent years. The original scheme builds a block cipher from a public permutation \mathbf{F} using two n -bit keys k_1, k_2 and its encryption function is defined as $\mathbf{E}(p) = \mathbf{F}(p + k_1) + k_2$ (where addition is over $GF(2)^n$). The generalized Even-Mansour scheme replaces the key additions with secret affine mappings and breaking it reduces to solving the affine equivalence problem, as originally described in [3].

The best known algorithms for the affine equivalence problem were presented by Biryukov et al. at EUROCRYPT 2003 [3]. The main algorithm described in [3] has complexity of about $n^3 2^{2n}$ bit operations, while a secondary algorithm has time complexity of about $n^3 2^{3n/2}$, but also uses about the same amount of memory.¹ Besides its high memory consumption, another disadvantage of the secondary algorithm of [3] is that it cannot be used to prove that \mathbf{F} and \mathbf{G} are not affine equivalent.

In this paper we devise a new algorithm for the affine equivalence problem whose complexity is about $n^3 2^n$ bit operations with very high probability whenever \mathbf{F} (or \mathbf{G}) is chosen uniformly at random from the set of all permutations on n -bit words. Our algorithm is also applicable without any modification to arbitrary functions (rather than permutations) and seems to perform similarly on random functions. However we focus on permutations as almost all applications actually require solving the affine equivalence problem for permutations. Since our algorithm can be used to prove that \mathbf{F} and \mathbf{G} are not affine equivalent, it does not share the disadvantage of the secondary algorithm of [3].

As a consequence of our improved algorithm, we solve within several minutes affine equivalence problem instances of size up to $n = 28$ on a single core. Optimizing our implementation and exploiting parallelism would most likely allow solving instances of size at least $n = 40$ using an academic budget. Such instances are out of reach of all previous algorithms for the problem.

Technically, the main algorithm devised in [3] for the affine equivalence problem is a guess-and-determine algorithm (which is related to the “to and fro” algorithm of [17], devised to solve the problem of isomorphism of polynomials) whereas the secondary algorithm is based on collision search (it generalizes

¹ Biryukov et al. also described a more efficient algorithm of complexity $n^3 2^n$ for the linear equivalence problem, which is a restricted variant of the affine equivalence problem.

Daemen’s attack on the original Even-Mansour cipher [8]). On the other hand, algorithms that use algebraic techniques (such as [5]) are mainly known for the asymmetric variant, in which \mathbf{F}, \mathbf{G} are functions of low degree, and it is not clear how to adapt them to arbitrary functions.

In contrast to previous algorithms, our approach involves analyzing algebraic properties of \mathbf{F}, \mathbf{G} which are of high algebraic degree. More specifically, we are interested in the polynomial representation (algebraic normal form or ANF) of each of the n output bits of \mathbf{F} (and \mathbf{G}) as a Boolean function in the n input bits. In fact, we are mainly interested in “truncated” polynomials that include only monomials of degree at least d (in particular, we choose $d = n - 2$). Each such polynomial can be viewed a vector in a vector space (with the standard basis of all monomials of degree at least $d = n - 2$). Therefore we can define the rank of the set of n truncated polynomials for each \mathbf{F}, \mathbf{G} as the rank of the matrix formed by arranging these polynomials as row vectors. In other words, we associate a rank value (which is an integer between 0 and n) to \mathbf{F} (and to \mathbf{G}) by computing the rank of its n truncated polynomials (derived from its n output bits) as vectors. We first show that if \mathbf{F}, \mathbf{G} are affine equivalent, their associated ranks are equal.²

To proceed, we analyze \mathbf{F}, \mathbf{G} independently. We derive from \mathbf{F} several functions, each one defined by restricting its 2^n inputs to an affine subspace of dimension $n - 1$. Since each such derived function (restricted to an affine subspace) has an associated rank, we assign to each possible $(n - 1)$ -dimensional subspace a corresponding rank. As there are 2^{n+1} possible affine subspaces (such a subspace can be characterized using its orthogonal subspace by a single linear expression over n variables and a free coefficient), we obtain 2^{n+1} rank values for \mathbf{F} . These values are collected in the *rank table* of \mathbf{F} , where a rank table entry r stores the set of all affine subspaces (more precisely, their compact representations as linear expressions) assigned to rank r .³

The main idea of the algorithm is to compute the rank tables of both \mathbf{F} and \mathbf{G} and then use these tables (and additional more complex structures derived from them) to recover the (unknown) affine transformation A_1 , assuming that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. In essence, the rank tables allow us to recover *matchings* between $(n - 1)$ -dimensional affine subspaces that are defined by A_1 : an affine subspace S in matched with S' if A_1 transforms S to the affine subspace S' . Each such matching between S and S' reveals information about A_1 in the form of linear equations. Hence we aim to use the rank tables to recover sufficiently many such matchings and compute A_1 using linear algebra. Once A_1 is derived, computing A_2 is trivial. The main property of the rank table that we prove and exploit to recover the matchings is that if S in matched with S' , then S appears in the rank table of \mathbf{G} in the same entry r as S' appears in the rank table of \mathbf{F} .

² The choice of $d = n - 2$ seems arbitrary at this stage. We only note that choosing a larger or smaller value for d typically results in a rank value which is constant for almost all functions, providing no information about them.

³ The formal definition of a rank table is slightly different.

Since the number of $(n-1)$ -dimensional affine subspaces is 2^{n+1} , each containing 2^{n-1} elements, a naive approach to computing the rank table (which works independently on each subspace) has complexity of at least $2^{n+1} \cdot 2^{n-1} = 2^{2n}$. However, using symbolic computation of polynomials, we show how to reduce this complexity to about $n^3 2^n$ bit operations. While this computational step is easy to analyze, this is not the case for the overall algorithm's performance. Indeed, its success probability and complexity depend on the monomials of degree at least $n-2$ of \mathbf{F} and \mathbf{G} . In particular, if all n output bits of \mathbf{F} and \mathbf{G} are functions of degree $n-3$ or lower, they do not contain any such monomials. As a result, all affine subspaces for \mathbf{F} and \mathbf{G} are assigned rank zero and the rank tables of these functions contain no useful information, leading to failure of the algorithm.⁴

When \mathbf{F} (or \mathbf{G}) is chosen uniformly at random from the set of all possible n -bit permutations (or n -bit functions in general), the case that its algebraic degree is less than $n-2$ is extremely unlikely for $n \geq 8$. Nevertheless, rigorous analysis of the algorithm seems challenging as its performance depends on subtle algebraic properties of random permutations. To deal with this situation, we make a heuristic assumption about the distribution of high degree monomials in random permutations which enables us to use well-known results regarding the rank distribution of random Boolean matrices. Consequently, we derive the distribution of the sizes of the rank table entries for a random permutation. This distribution and additional properties enable us to show that asymptotically the algorithm succeeds with probability close to 1 in complexity of about $n^3 2^n$ bit operations. This heuristic analysis is backed up by thousands of experiments on various problem instances of different sizes. Rigorously analyzing the algorithm and extending it to succeed on all functions (or permutations) with probability 1 in the same complexity remain open problems.

The properties of the rank table and the algorithm for computing it are of independent interest. In particular, we propose methods to build experimental distinguishers for block ciphers based on the rank table and a method to efficiently detect high-order differential distinguishers based on the algorithm for its computation. Furthermore, our techniques are relevant to decomposition attacks on the white-box ASASA block cipher instances proposed by Biryukov et al. [2]. In this application, we adapt the algorithm for computing the rank table in order to improve the complexity of the integral attack on ASASA published in [9] from $2^{3n/2}$ to about 2^n (where n is the block size of the instance).

The rest of the paper is organized as follows. In Section 2 we describe some preliminaries and give an overview of the new affine equivalence algorithm in Section 3. In Section 4 we prove the basic property of rank equality for affine equivalent functions, while in Section 5 we define and analyze the matching between $(n-1)$ -dimensional affine subspaces that we use to recover A_1 . In Section 6 we define the rank table and additional objects used in our algorithm, and describe the relation between these objects for affine equivalent functions.

⁴ In case the algorithm fails, one can try to apply it to \mathbf{F}^{-1} and \mathbf{G}^{-1} which may be of higher algebraic degree.

In Section 7 we analyze properties of rank tables for random permutations under our heuristic assumption. Then, we describe and analyze the new affine equivalence algorithm in Section 8. Next, in Section 9, we describe applications of the new algorithm and the rank table structure. Finally, we conclude the paper in Section 10.

2 Preliminaries

For a finite set R , denote by $|R|$ its size. Given a vector $u = (u[1], \dots, u[n]) \in GF(2)^n$, let $wt(u)$ denote its Hamming weight. Throughout this paper, addition between vectors $u_1, u_2 \in GF(2)^n$ is performed bit-wise over $GF(2)^n$.

Multivariate Polynomials Any Boolean function $F : \{0, 1\}^n \rightarrow \{0, 1\}$ can be represented as a multivariate polynomial whose algebraic normal form (ANF) is unique and given as $F(x[1], \dots, x[n]) = \sum_{u=(u[1], \dots, u[n]) \in \{0, 1\}^n} \alpha_u M_u$, where $\alpha_u \in \{0, 1\}$ is the coefficient of the monomial $M_u = \prod_{i=1}^n x[i]^{u[i]}$, and the sum is over $GF(2)$. The algebraic degree of the function F is defined as $deg(F) = \max\{wt(u) \mid \alpha_u \neq 0\}$.

In several cases it will be more convenient to directly manipulate the representation of F as a multivariate polynomial $P(x[1], \dots, x[n]) = \sum_{u \in \{0, 1\}^n} \alpha_u M_u$. Note that unlike F , the polynomial P is not treated as a function but rather as a symbolic object. $P(x[1], \dots, x[n])$ can be viewed as a vector in the vector space spanned by the set of all monomials $\{M_u \mid u \in \{0, 1\}^n\}$.

Given a multivariate polynomial $P(x[1], \dots, x[n]) = \sum_{u \in \{0, 1\}^n} \alpha_u M_u$ and a positive integer d , define $P_{(\geq d)}$ by taking all the monomials of P of degree at least d , namely, $P_{(\geq d)}(x[1], \dots, x[n]) = \sum_{u \in \{0, 1\}^n \wedge wt(u) \geq d} \alpha_u M_u$. Note that $P_{(\geq d)}$ can be represented using at most $\sum_{i=d}^n \binom{n}{i}$ non-zero coefficients.

Given a function $F : \{0, 1\}^n \rightarrow \{0, 1\}$ represented by a polynomial $P(x[1], \dots, x[n])$, define $F_{(\geq d)} : \{0, 1\}^n \rightarrow \{0, 1\}$ as the function represented by $P_{(\geq d)}$.

Vectorial Functions and Polynomials Given a vectorial Boolean function $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, let $F^{(i)} : \{0, 1\}^n \rightarrow \{0, 1\}$ denote the Boolean function of its i 'th output bit.

We say that a sequence of m polynomials $\mathbf{P} = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^m$ represents \mathbf{F} if for each $i \in \{1, 2, \dots, m\}$, the i 'th polynomial $P^{(i)}$ represents $F^{(i)}$.

Given a positive integer d , denote $\mathbf{P}_{(\geq d)} = \{P_{(\geq d)}^{(i)}(x[1], \dots, x[n])\}_{i=1}^m$. The vectorial function $\mathbf{F}_{(\geq d)} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is defined analogously.

The algebraic degree $deg(\mathbf{P})$ of \mathbf{P} is defined as the maximal degree of its m polynomials. The algebraic degree $deg(\mathbf{F})$ is defined analogously.

As each $P^{(i)}$ can be viewed as a vector in a vector space, we define the *symbolic rank* of \mathbf{P} as the rank of the m vectors $\{P^{(i)}\}_{i=1}^m$. We denote the symbolic rank of \mathbf{P} as $SR(\mathbf{P})$. Note that $SR(\mathbf{P}) \in \mathbb{Z}_{m+1}$.

Affine Transformations and Affine Equivalence An affine transformation $A : \{0, 1\}^m \rightarrow \{0, 1\}^n$ over $GF(2)^m$ is defined using a Boolean matrix $L_{n \times m}$ and a word $a \in \{0, 1\}^n$ as $A(x) = L(x) + a$ ($L(x)$ is simply matrix multiplication). The transformation is invertible if $m = n$ and L is an invertible matrix. If $a = 0$, then the A is called a linear transformation (such functions are a subclass of affine functions).

Two functions $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ are *affine equivalent* if there exist two invertible affine transformations $A_1 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $A_2 : \{0, 1\}^m \rightarrow \{0, 1\}^m$ such that $G = A_2 \circ F \circ A_1$. It is easy to show that the affine equivalence relation partitions the set of all functions into (affine) equivalence classes. We denote $F \equiv G$ if F is affine equivalent to G .

Symbolic Composition Given $P = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^m$, and an affine function $A_1 : \{0, 1\}^{n'} \rightarrow \{0, 1\}^n$, the composition $P \circ A_1$ is a sequence of m polynomials in n' variables. For $i \in \{1, 2, \dots, m\}$, the i 'th polynomial in this composition is $P^{(i)} \circ A_1$. It can be computed by substituting each variable $x[j]$ in $P^{(i)}$ with the (affine) symbolic representation of the j 'th output bit of A_1 (and simplifying the outcome to obtain the ANF).

For example, given $P(x[1], x[2], x[3]) = x[1]x[2] + x[1]x[3] + x[2] + 1$ and $A_1 : \{0, 1\}^2 \rightarrow \{0, 1\}^3$ defined by the relations $x[1] = y[1] + y[2] + 1, x[2] = y[2], x[3] = y[1] + y[2]$, then

$$\begin{aligned} P \circ A_1 &= (y[1] + y[2] + 1)(y[2]) + (y[1] + y[2] + 1)(y[1] + y[2]) + y[2] + 1 = \\ &= (y[1]y[2] + y[2] + y[2]) + (y[1] + y[1]y[2] + y[1]y[2] + y[2] + y[1] + y[2]) + y[2] + 1 = \\ &= y[1]y[2] + y[2] + 1. \end{aligned}$$

Thus, we compose each monomial M_u with coefficient 1 in P with A_1 to obtain a polynomial expression, add all the expressions and simplify the result. Formally, if we denote P 's M_u coefficient by α_u , then

$$P \circ A_1 = \sum_{u \in \{0, 1\}^n} \alpha_u \cdot (M_u \circ A_1).$$

Note that composition with an affine function does not increase the algebraic degree of the composed polynomial, namely $\deg(P \circ A_1) \leq \deg(P)$.

Analogously, given an affine function $A_2 : \{0, 1\}^m \rightarrow \{0, 1\}^{m'}$, the composition $A_2 \circ P$ is a sequence of m' polynomials in n variables. It can be computed by substituting each variable $x[j]$ in the (affine) symbolic representation of A_2 with $P^{(j)}$. Equivalently, if $A_2(x) = L(x) + a$, then $A_2 \circ P$ can be computed by symbolic matrix multiplication (and addition of a) as $L(P) + a$. In particular, if $m' = 1$ and $a = 0$, then A_2 reduces to a vector $v = (v[1], v[2], \dots, v[m]) \in \{0, 1\}^m$ and $v(P) = \sum_{i=1}^m v[i]P^{(i)}$ is a symbolic inner product.

By rules of composition, if F is represented by P , then $P \circ A_1$ represents $F \circ A_1$ (which is a standard composition of functions) and $A_2 \circ P$ represents $A_2 \circ F$.

Half-Space Masks and Coefficients Let $A : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^n$ be an affine transformation such that $A(x) = L(x) + a$ for a matrix $L_{n \times n-1}$ with linearly independent columns. Then the (affine) range of A is an $(n - 1)$ -dimensional affine subspace spanned by the columns of L with the addition of a . The subspace orthogonal to the range of A is of dimension 1 and hence spanned by a single non-zero vector $h \in \{0, 1\}^n$. Namely, a vector $v \in \{0, 1\}^n$ is in the range of A if and only if $h(v + a) = 0$, i.e., v satisfies the linear equation $h(v) + h(a) = 0$.

Since h partitions the space of $\{0, 1\}^n$ into two halves, we call h the *half-space mask* (HSM) of A and call the bit $h(a)$ the *half-space free coefficient* (HSC) of A .

We call the linear subspace spanned by the columns of L the *linear range* of A . A vector $v \in \{0, 1\}^n$ is in the linear range of A if and only if $h(v) = 0$.

Canonical Affine Transformations Given non-zero $h \in \{0, 1\}^n$ and $c \in \{0, 1\}$, there exist many affine transformations whose HSM and HSC are equal to h, c , respectively. We will use the fact (stated formally below) that all affine transformations with an identical affine range are related by composition on the right with an invertible affine transformation.

Fact 1 *The affine transformations $A_1 : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^n$ and $A_2 : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^n$ have the same affine range if and only if there exists an invertible affine transformation $A' : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^{n-1}$ such that $A_1 = A_2 \circ A'$.*

Given A_1, A_2 the matrix A' above can be computed using basic linear algebra.

We now define the *canonical affine transformation* $C_{|h,c} : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^n$ with respect to h, c . Let ℓ denote the index of the first non-zero bit of $h = (h[1], \dots, h[n])$. Write $C_{|h,c}(x) = L(x) + a$. We define $a = c \cdot e_\ell$ (where e_ℓ is the ℓ 'th unit vector) and define $L[i]$ (the i 'th column of L) using h and the unit vectors as follows:

$$L[i] = \begin{cases} e_i & \text{if } i < \ell \\ e_{i+1} + h[i+1]e_\ell & \text{otherwise } (\ell \leq i \leq n-1) \end{cases}$$

Thus, on input $(y[1], \dots, y[n-1])$, the transformation $C_{|h,c}$ is defined by the symbolic form

$$(x[1], x[2], \dots, x[n]) = (y[1], \dots, y[\ell-1], \sum_{i=\ell}^{n-1} h[i+1]y[i] + c, y[\ell], \dots, y[n-1]).$$

The motivation behind the definition of $C_{|h,c}$ is that it allows very simple symbolic composition when applied on the right: its main action is to replace the variable $x[\ell]$ with the affine combination that is specified by the coefficients of h and by c . Other variables are just renamed: variables with index $i < \ell$ remain with the same index, while for each variable with index $i > \ell$, its index is reduced by 1.

Remark 1. Note that we have to show that the definition of $C_{|h,c}$ is valid. First, the $n - 1$ columns of L are clearly linearly independent. It remains to prove

that h, c are indeed the HSM and HSC of $C|_{h,c}$. For this purpose, it suffices to show that for each column $L[i]$, the vector $L[i] + a$ satisfies the equation $h(L[i] + a) + c = 0$. Since $h(a) = h(c \cdot e_\ell) = c \cdot h(e_\ell) = c$ (as $h_\ell = 1$), it remains to show that $h(L[i]) = 0$. Indeed, if $0 \leq i < \ell$, then $h(L[i]) = h_i = 0$ (as ℓ is the index of the first non-zero bit of h). Otherwise, $\ell \leq i \leq n - 1$, then $h(L[i]) = h[i + 1] + h[i + 1]h[\ell] = 0$ (as $h[\ell] = 1$).

3 Overview of the New Affine Equivalence Algorithm

We demonstrate the new algorithm using an example. Although it is oversimplified, this example is sufficient to convey the main ideas of our algorithm.

Definition of Functions We define the function $\mathbf{F} : \{0, 1\}^3 \rightarrow \{0, 1\}^3$ using its symbolic representation $\mathbf{P} = \{P^{(i)}(x[1], x[2], x[3])\}_{i=1}^3$,

$$\begin{aligned} P^{(1)}(x[1], x[2], x[3]) &= x[1]x[2] + x[1]x[3] + x[2] + 1 \\ P^{(2)}(x[1], x[2], x[3]) &= x[1]x[2] + x[1] + x[2] \\ P^{(3)}(x[1], x[2], x[3]) &= x[1]x[3] + x[3]. \end{aligned}$$

We define $\mathbf{G} : \{0, 1\}^3 \rightarrow \{0, 1\}^3$ using 2 affine transformations as $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$, where A_2 is simply the identity and A_1 is defined using the relations

$$x[1] = y[1] + y[3] + 1, \quad x[2] = y[1] + y[2], \quad x[3] = y[2].$$

Composing $A_2 \circ \mathbf{P} \circ A_1$ and simplifying the resultant ANFs, gives the symbolic representation of \mathbf{G} as $\mathbf{Q} = \{Q^{(i)}(y[1], y[2], y[3])\}_{i=1}^3$, where

$$\begin{aligned} Q^{(1)}(y[1], y[2], y[3]) &= y[1]y[3] + y[1] + y[2] + 1 \\ Q^{(2)}(y[1], y[2], y[3]) &= y[1]y[2] + y[1]y[3] + y[2]y[3] + y[3] + 1 \\ Q^{(3)}(y[1], y[2], y[3]) &= y[1]y[2] + y[2]y[3]. \end{aligned}$$

The input to our affine equivalence algorithm is \mathbf{F}, \mathbf{G} defined above and its goal is to recover the (presumably) unknown affine transformation A_1 . The first step of the algorithm is to interpolate \mathbf{F}, \mathbf{G} and obtain \mathbf{P}, \mathbf{Q} , respectively.

Rank Tables and Histograms The most basic property that we prove in Theorem 1 is that since \mathbf{F} and \mathbf{G} are affine equivalent, the symbolic ranks of \mathbf{P} and \mathbf{Q} (as vectors) are equal. Indeed, it is easy to verify that both \mathbf{P} and \mathbf{Q} have symbolic rank of 3. More significantly, Theorem 1 is stronger and asserts that $SR(\mathbf{P}_{(\geq d)}) = SR(\mathbf{Q}_{(\geq d)})$ for every $d \geq 1$. Indeed, if we take $d = 2$, we get $\mathbf{P}_{(\geq 2)} = \{x[1]x[2] + x[1]x[3], x[1]x[2], x[1]x[3]\}$, which has symbolic rank 2. This is also the symbolic rank of $\mathbf{Q}_{(\geq 2)} = \{y[1]y[3], y[1]y[2] + y[1]y[3] + y[2]y[3], y[1]y[2] + y[2]y[3]\}$.

We would like to use this property to recover A_1 . Let us examine the 2-dimensional affine subspace defined by the 3-bit HSM $h' = 100$ (whose bits are $h[1] = 1, h[2] = 0, h[3] = 0$) and the single bit HSC $c = 0$. We calculate the symbolic form of $(\mathbf{F} \circ C_{|h',0})_{(\geq 1)}$ by evaluating $(\mathbf{P} \circ C_{|h',0})_{(\geq 1)}$ (i.e., plugging $x[1] = 0$ into $\mathbf{P}_{(\geq 1)}$) and obtain $\{x[2], x[2], x[3]\}$ which has symbolic rank 2. Similarly, for $c = 1$ we calculate $(\mathbf{P} \circ C_{|h',1})_{(\geq 1)}$ (i.e., plug $x[1] = 1$ into $\mathbf{P}_{(\geq 1)}$) and obtain $\{x[3], 0, 0\}$ which has symbolic rank 1. Hence, we attach the symbolic rank pair $(2, 1)$ to $h' = 100$. We do the same for all 7 non-zero $h' \in \{0, 1\}^3$. The result is a table whose entries are pairs of ranks of the form $(\mathit{maxR}, \mathit{minR}) \in \mathbb{Z}_4 \cdot \mathbb{Z}_4$ (where $\mathit{maxR} \geq \mathit{minR}$), such that entry $(\mathit{maxR}, \mathit{minR})$ stores the set of HSMs that are associated with this pair of ranks.

$$\begin{aligned} (3, 2) &: \{010, 011, 111, 110\} \\ (2, 2) &: \{001\} \\ (2, 1) &: \{100, 101\} \end{aligned}$$

This table is called the *rank table* of \mathbf{F} (with respect to the degree $d = 1$ as we only considered monomials of degree at least 1). The set of HSMs in an entry $(\mathit{maxR}, \mathit{minR})$ of the rank table is called a *rank group* (e.g., the rank group with index $(2, 1)$ is $\{100, 101\}$). Similarly, we compute the rank table of \mathbf{G} with respect to $d = 1$.

$$\begin{aligned} (3, 2) &: \{100, 001, 110, 011\} \\ (2, 2) &: \{010\} \\ (2, 1) &: \{101, 111\} \end{aligned}$$

Although the rank tables are different, the size of each rank group $(\mathit{maxR}, \mathit{minR})$ of \mathbf{F}, \mathbf{G} is identical. We define the *rank histogram* of \mathbf{F} (with respect to d) as a mapping from each $(\mathit{maxR}, \mathit{minR})$ value to the corresponding rank group size (e.g., the histogram entry for \mathbf{F} with index $(2, 1)$ has value $|\{100, 101\}| = 2$). As we show in Lemma 9, that rank histograms of affine equivalent functions (such as \mathbf{F}, \mathbf{G}) are identical.

To explain this, we look at the HSM $h = 101$ in the rank group of $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$ with index $(2, 1)$ and note that it partitions the space into halves $\{000, 101, 010, 111\}$ and $\{001, 011, 100, 110\}$ (according to whether $x[1] + x[3] = 0$ or $x[1] + x[3] = 1$). After applying A_1 , these half-spaces are mapped into $\{100, 101, 110, 111\}$ and $\{000, 001, 010, 011\}$. This is exactly the partition defined by $h' = 100$, which is in the rank group of \mathbf{F} with the same index $(2, 1)$. In terms of canonical affine transformations, $C_{|h',c}$ and $A_1 \circ C_{|h,0}$ have the same half-space range of $\{100, 101, 110, 111\}$ (for $c = 1$ in this case) and we define a mapping $h \mapsto_{A_1} h'$ to capture this. In Lemma 3 we show that this mapping is a bijection as A_1 is invertible.

Exploiting Matchings The central property of the mapping \mapsto_{A_1} is proved in Lemma 6 which asserts that it preserves affine equivalence. Namely, if $\mathbf{F} \equiv \mathbf{G}$ and $h \mapsto_{A_1} h'$, then $\mathbf{F} \circ C_{|h',c} \equiv \mathbf{G} \circ C_{|h,0}$ (for some $c \in \{0, 1\}$). By flipping the half-space ranges and applying the same argument, we also obtain $\mathbf{F} \circ C_{|h',c+1} \equiv \mathbf{G} \circ C_{|h,1}$. Combined with Theorem 1 (which states that symbolic rank is an invariant of affine equivalent functions) we obtain that for any $d \geq 1$, $r_0 \triangleq SR((\mathbf{F} \circ C_{|h',c})_{(\geq d)}) = SR((\mathbf{G} \circ C_{|h,0})_{(\geq d)})$ and $r_1 \triangleq SR((\mathbf{F} \circ C_{|h',c+1})_{(\geq d)}) = SR((\mathbf{G} \circ C_{|h,1})_{(\geq d)})$. Since the ordered rank pairs for h', h in \mathbf{F}, \mathbf{G} are equal to $(maxR, minR)$ (for $maxR = \max\{r_0, r_1\}, minR = \min\{r_0, r_1\}$), they belong to rank groups with the same index $(maxR, minR)$ in the rank tables of \mathbf{F}, \mathbf{G} , respectively. The fact that \mapsto_{A_1} is a bijection leads to Lemma 9 (which asserts that the rank histograms of \mathbf{F}, \mathbf{G} are identical).

The main goal of our affine equivalence algorithm is to recover *matchings* $h \mapsto_{A_1} h'$ for several pairs h, h' . This is useful, as in Lemma 5 we show that each such matching gives n linear equations on the unknown matrix L of $A_1(x) = L(x) + a$. Furthermore, the constant c associated with $h \mapsto_{A_1} h'$ (which determines whether $\mathbf{F} \circ C_{|h',0} \equiv \mathbf{G} \circ C_{|h,0}$ or $\mathbf{F} \circ C_{|h',1} \equiv \mathbf{G} \circ C_{|h,0}$) gives a linear equation on a (once again, by Lemma 5). In total, we need to find about n matchings $h \mapsto_{A_1} h'$ along with their associated constants to completely recover A_1 .

Going back to the example, the rank group with index $(2, 2)$ for \mathbf{G} is $\{010\}$, while the rank group with the same index for \mathbf{F} is $\{001\}$. Therefore, after computing the rank tables we know that

$$010 \mapsto_{A_1} 001. \quad (1)$$

Remark 2. We matched $010 \mapsto_{A_1} 001$ in the rank group $(maxR, minR) = (2, 2)$ and since $maxR = minR$ we cannot derive the constant c associated with this matching (hence we cannot derive a linear equation on a). Such constants can only be derived for matchings $h \mapsto_{(A_1)} h'$ in rank groups where $maxR > minR$, as in such cases we know whether $\mathbf{F} \circ C_{|h',0} \equiv \mathbf{G} \circ C_{|h,0}$ or $\mathbf{F} \circ C_{|h',1} \equiv \mathbf{G} \circ C_{|h,0}$ according to the equality $SR((\mathbf{F} \circ C_{|h',c})_{(\geq d)}) = SR((\mathbf{G} \circ C_{|h,0})_{(\geq d)})$. More precisely, if $maxR > minR$, then either $SR((\mathbf{F} \circ C_{|h',0})_{(\geq d)}) = SR((\mathbf{G} \circ C_{|h,0})_{(\geq d)})$ or $SR((\mathbf{F} \circ C_{|h',1})_{(\geq d)}) = SR((\mathbf{G} \circ C_{|h,0})_{(\geq d)})$, but not both. In this sense, it is more useful to recover matchings for HSMs in rank groups $(maxR, minR)$ such that $maxR > minR$.

By applying similar arguments to the rank group $(2, 1)$, we know that either $101 \mapsto_{A_1} 100$ or $101 \mapsto_{A_1} 101$ (and similarly $111 \mapsto_{A_1} 100$ or $111 \mapsto_{A_1} 101$). Since we have very few possibilities, we can guess which matchings hold, derive A_1 and test our guess. Unfortunately, for larger n we expect the rank groups to be much bigger and it would be inefficient to exhaustively match HSMs for \mathbf{F}, \mathbf{G} only based on their ranks. Thus, to narrow down the number of possibilities and eventually uniquely match sufficiently many pairs h, h' such that $h \mapsto_{A_1} h'$, we need to attach more data to each HSM for \mathbf{F}, \mathbf{G} .

HSM Rank Histograms The main observation that allows attaching more data to each HSM is given in Lemma 4 which shows that the mapping \mapsto_{A_1} is additive.

Consider the two rank groups with index $(3, 2)$ for \mathbf{F}, \mathbf{G} . They are of size 4 and their HSMs cannot be uniquely matched. We first focus on \mathbf{G} and examine the rank group $(3, 2)$ which is $\{100, 001, 110, 011\}$. We take $h_1 = 011$ and compute its *HSM rank histogram* with respect to the rank group $(2, 1)$ (which is $\{101, 111\}$). This is done by computing the $(maxR, minR)$ rank pairs for the set defined by adding all elements of rank group $(2, 1)$ to 011, namely $\{011 + 101, 011 + 111\} = \{110, 100\}$. Looking for 110 and 100 in the rank table of \mathbf{G} , both HSMs have ranks $(3, 2)$. Thus, the HSM rank histogram of $h_1 = 011$ with respect to rank group $(2, 1)$ has a single non-zero entry $(3, 2)$ with the value of 2. We write this HSM rank histogram in short as $[(3, 2) : 2]$.

We now consider the match of $h_1 = 011$ under A_1 which is $h'_1 = 110$ (namely, $h_1 \mapsto_{(A_1)} h'_1$). Similarly to h_1 , we compute the HSM rank histogram of h'_1 with respect to the rank group $(2, 1)$ for \mathbf{F} (which is $\{100, 101\}$) and obtain the same HSM rank histogram $[(3, 2) : 2]$. This is a particular case of Lemma 10, which shows that matching HSMs for \mathbf{F}, \mathbf{G} have identical HSM rank histograms (with respect to a fixed rank group). Lemma 10 is derived using Lemma 4 which asserts that the mapping \mapsto_{A_1} is additive: if $h_1 \mapsto_{(A_1)} h'_1$ and $h_2 \mapsto_{(A_1)} h'_2$, then $(h_1 + h_2) \mapsto_{(A_1)} (h'_1 + h'_2)$.

Fixing $h_1 = 011$ for \mathbf{G} and its match $h'_1 = 110$ under A_1 , let h_2^i, h'_2^i for $i \in \{1, 2\}$ vary over the 2 elements of the rank groups with index $(2, 1)$ in \mathbf{G}, \mathbf{F} , respectively. Then, as $h_1 \mapsto_{(A_1)} h'_1$ and $h_2^i \mapsto_{(A_1)} h'_2^i$ for $i \in \{1, 2\}$, we get $(h_1 + h_2^i) \mapsto_{(A_1)} (h'_1 + h'_2^i)$. By the aforementioned Theorem 1 and Lemma 6 (equating ranks for matching HSMs) we conclude that indeed the HSM rank histograms of 011 and 110 with respect to rank group $(2, 1)$ are identical (which is a special case of Lemma 10).

HSM Rank Histogram Multi-Sets Since we do not know in advance that $h_1 \mapsto_{(A_1)} h'_1$, we have to compute the HSM rank histograms (with respect to rank group $(2, 1)$) for all HSMs in rank group $(3, 2)$. The outcome is the *HSM rank histogram multi-set* of rank group $(3, 2)$ with respect to rank group $(2, 1)$. It is computed by considering all the HSMs in the rank group $(3, 2)$, namely $\{100, 001, 110, 011\}$ for \mathbf{G} and $\{010, 011, 111, 110\}$ for \mathbf{F} . Lemma 11 (whose proof is based on Lemma 10) asserts that these HSM rank histogram multi-sets are identical as \mathbf{F}, \mathbf{G} are affine equivalent.

We hope that these multi-sets contain unique HSM rank histograms (with multiplicity 1), which would allow us to derive more matching between HSMs. Unfortunately, the resultant multi-set (for both \mathbf{F} and \mathbf{G}) is $\{[(3, 2) : 2], [(3, 2) : 2], [(3, 2) : 2], [(3, 2) : 2]\}$. It contains 4 identical elements and does not give us any new information about A_1 . If the multiplicity of the element $[(3, 2) : 2]$ (calculated above for h_1, h'_1) in this multi-set would have been 1, we could have derived the relation $h_1 \mapsto_{(A_1)} h'_1$.

Remark 3. Generally, when n is very small (as in our case), the direct application of the algorithm is more likely to fail to completely recover A_1 . As we show later in this paper, for $n \geq 8$ the fraction of instances for which this occurs is very small (and tends to 0 as n grows). In some cases a failure to retrieve A_1 occurs since the affine mappings A_1, A_2 are not uniquely defined. In particular, if there are several solutions for A_1 , then we cannot hope to obtain unique matchings that completely define A_1 , but we can recover all possible solutions to the affine equivalence problem by enumerating several possibilities for the matchings.

In conclusion, we attached to each HSM in rank group $(3, 2)$ for \mathbf{F}, \mathbf{G} its HSM rank histogram with respect to rank group $(2, 1)$ and in general such data may allow us to derive additional matchings $h \mapsto_{(A_1)} h'$. Once we obtain about n matchings, we can recover A_1 by solving a system of linear equations.

4 A Basic Property of Affine Equivalent Functions

Before proving the main result of this section, we state two useful lemmas (the first is proved in Appendix A).

Lemma 1. *Let $\mathbf{P} = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^m$, let $A_1 : \{0, 1\}^{n'} \rightarrow \{0, 1\}^n$, $A_2 : \{0, 1\}^m \rightarrow \{0, 1\}^{m'}$ be affine functions, and d be a positive integer. Then,*

1. $(\mathbf{P}_{(\geq d)} \circ A_1)_{(\geq d)} = (\mathbf{P} \circ A_1)_{(\geq d)}$
2. $(A_2 \circ (\mathbf{P}_{(\geq d)}))_{(\geq d)} = (A_2 \circ \mathbf{P})_{(\geq d)}$ and if A_2 is a linear function, $A_2 \circ (\mathbf{P}_{(\geq d)}) = (A_2 \circ \mathbf{P})_{(\geq d)}$.

Essentially, the lemma states that removing all monomials of degree less than d from \mathbf{P} can be done before or after composing it with an affine function and the outcomes are identical.

Note that a potentially simplified first part of the lemma which equates $\mathbf{P}_{(\geq d)} \circ A_1$ and $(\mathbf{P} \circ A_1)_{(\geq d)}$ is generally incorrect, as the first expression may contain monomials of degree less than d . For example, if $d = 2$ and we compose the affine transformation defined by $x[1] = y[1] + y[2]$ and $x[2] = y[2]$ with the polynomial $x[1]x[2]$, then we get the polynomial $y[1]y[2] + y[2]^2$ which has a monomial of degree 1.

Lemma 2. *Let $\mathbf{P} = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^m$, and let $A_1 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an invertible affine function. Then, $\deg(\mathbf{P}) = \deg(\mathbf{P} \circ A_1)$.*

Proof. We show that for $i \in \{1, 2, \dots, m\}$, $\deg(P^{(i)}) = \deg(P^{(i)} \circ A_1)$. Observe that $\deg(P^{(i)}) \geq \deg(P^{(i)} \circ A_1)$ as composition with an affine function cannot increase the algebraic degree of a polynomial. By the same argument and by the invertibility of A_1 , we also obtain $\deg(P^{(i)} \circ A_1) \geq \deg(P^{(i)} \circ A_1 \circ (A_1)^{-1}) = \deg(P^{(i)})$. ■

Theorem 1. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions, represented by \mathbf{P}, \mathbf{Q} , respectively. Then, for every positive integer d , $SR(\mathbf{P}_{(\geq d)}) = SR(\mathbf{Q}_{(\geq d)})$.*

Proof. At a high level, the fact that \mathbf{P} and \mathbf{Q} have the same symbolic rank follows since rank is preserved by composition with invertible affine transformations. Moreover, this rank equality is preserved after truncating low degree monomials since they cannot affect the high degree monomials when composing with invertible affine transformations. The formal proof is below.

Write $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$, implying that $\mathbf{Q} = A_2 \circ \mathbf{P} \circ A_1$. Denote $\mathbf{P}' = \mathbf{P} \circ A_1$ and observe that

$$SR(\mathbf{P}'_{(\geq d)}) = SR((A_2 \circ (\mathbf{P}'_{(\geq d)}))_{(\geq d)}) = SR((A_2 \circ \mathbf{P}')_{(\geq d)}) = SR(\mathbf{Q}_{(\geq d)}),$$

where the first equality holds since rank is preserved by invertible linear transformations⁵ and the second equality is due to the second part of Lemma 1.

It remains to show that $SR(\mathbf{P}'_{(\geq d)}) = SR(\mathbf{P}_{(\geq d)})$, or $SR((\mathbf{P} \circ A_1)_{(\geq d)}) = SR(\mathbf{P}_{(\geq d)})$. We first show that $SR(\mathbf{P}_{(\geq d)}) \geq SR((\mathbf{P} \circ A_1)_{(\geq d)})$.

If $\mathbf{P}_{(\geq d)}$ has full rank of m then the claim is trivial. Otherwise, let $v \in \{0, 1\}^m$ be a non-zero vector in the kernel of $\mathbf{P}_{(\geq d)}$, namely $v(\mathbf{P}_{(\geq d)}) = 0$. Then,

$$v((\mathbf{P} \circ A_1)_{(\geq d)}) = (v((\mathbf{P} \circ A_1)_{(\geq d)}))_{(\geq d)} = (v(\mathbf{P}_{(\geq d)}) \circ A_1)_{(\geq d)} = 0,$$

where the first equality follows from the second part of Lemma 1 and the second equality follows from the first part of this lemma. This implies that v is also in the kernel of $(\mathbf{P} \circ A_1)_{(\geq d)}$, as required.

To prove that $SR(\mathbf{P}_{(\geq d)}) \leq SR((\mathbf{P} \circ A_1)_{(\geq d)})$, observe that if v is a non-zero vector in the kernel of $(\mathbf{P} \circ A_1)_{(\geq d)}$, then by the equality above we have $0 = v((\mathbf{P} \circ A_1)_{(\geq d)}) = (v(\mathbf{P}_{(\geq d)}) \circ A_1)_{(\geq d)}$. This implies that $\deg(v(\mathbf{P}_{(\geq d)}) \circ A_1) < d$ and since A_1 is invertible, by Lemma 2, $\deg(v(\mathbf{P}_{(\geq d)})) = \deg(v(\mathbf{P}_{(\geq d)}) \circ A_1) < d$. This gives $v(\mathbf{P}_{(\geq d)}) = 0$, as the polynomial does not contain monomials of degree less than d . Hence, v is in the kernel of $\mathbf{P}_{(\geq d)}$ which completes the proof. ■

5 The Half-Space Mask Bijection and its Properties

Definition 1. Let $A : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an invertible affine transformation. Define a mapping between HSMs using A as follows: $h \in \{0, 1\}^n$ is mapped to h' if there exists $c \in \{0, 1\}$ such that the affine ranges of $A \circ C_{|h,0}$ and $C_{|h',c}$ are equal. We write $h \mapsto_{(A)} h'$ and say that h and h' match (under A). The bit c is called the associated constant of $h \mapsto_{(A)} h'$.

Lemma 3. Let $A : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an invertible affine transformation. The mapping $\mapsto_{(A)}$ is a bijection and its inverse is given by $\mapsto_{(A^{-1})}$.

Proof. The proof follows from the invertibility of A . Given that $h \mapsto_{(A)} h'$, there exists $c \in \{0, 1\}$ such that the affine ranges of $A \circ C_{|h,0}$ and $C_{|h',c}$ are equal. According to Fact 1, this implies that there exists an invertible affine

⁵ The affine transformation A_2 also adds a constant, but it does contribute to the rank as $d > 0$.

transformation $A' : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^{n-1}$ such that $A \circ C_{|h,0} = C_{|h',c} \circ A'$. Consequently, $C_{|h,0} \circ (A')^{-1} = A^{-1} \circ C_{|h',c}$ and the affine ranges of $C_{|h,0}$ and $A^{-1} \circ C_{|h',c}$ are equal (again, according to Fact 1). This implies that the affine ranges of $A^{-1} \circ C_{|h',0}$ and $C_{|h,c}$ are equal (flipping the HSC of both sides if $c = 1$), namely $h' \mapsto_{(A^{-1})} h$. ■

A property of $\mapsto_{(A)}$ which will be very useful is that it is additive. This is established by the lemma below (proved in Appendix A).

Lemma 4. *Let $A : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an invertible affine transformation. Let $h_1, h'_1, h_2, h'_2 \in \{0, 1\}^n$ be HSMs where $h_1 \neq h_2$ and $h_1 \mapsto_{(A)} h'_1$, $h_2 \mapsto_{(A)} h'_2$ with associated constants c_1, c_2 , respectively. Then $(h_1 + h_2) \mapsto_{(A)} (h'_1 + h'_2)$ with the associated constant $c_1 + c_2$.*

The following lemma (proved in Appendix A) shows that the bijection reveals information about the presumably unknown transformation A .

Lemma 5. *Let $A : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an invertible affine transformation such that $A(x) = L(x) + a$. Let $h, h' \in \{0, 1\}^n$ be HSMs such that $h \mapsto_{(A)} h'$ with associated constant c . Then, A satisfies the following constraints.*

1. For each $i \in \{1, 2, \dots, n\}$, the i 'th column of L , denoted by $L[i]$, satisfies the equation $h'(L[i]) = h[i]$, where $h[i]$ is the i 'th bit of h .
2. The vector a satisfies the equation $h'(a) = c$.

The following lemma asserts that affine equivalence is preserved under composition with matching HSMs.

Lemma 6. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions such that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Let $h, h' \in \{0, 1\}^n$ be HSMs such that $h \mapsto_{(A_1)} h'$ with associated constant c . Then, $\mathbf{F} \circ C_{|h',c} \equiv \mathbf{G} \circ C_{|h,0}$ and $\mathbf{F} \circ C_{|h',c+1} \equiv \mathbf{G} \circ C_{|h,1}$.*

Proof. Since $h \mapsto_{(A_1)} h'$ with associated constant c , the affine ranges of $A_1 \circ C_{|h,0}$ and $C_{|h',c}$ are equal. According to Fact 1, there exists an invertible affine transformation $A'_1 : \{0, 1\}^{n-1} \rightarrow \{0, 1\}^{n-1}$ such that $A_1 \circ C_{|h,0} = C_{|h',c} \circ A'_1$.

We obtain, $A_2 \circ \mathbf{F} \circ C_{|h',c} \circ A'_1 = A_2 \circ \mathbf{F} \circ A_1 \circ C_{|h,0} = \mathbf{G} \circ C_{|h,0}$, implying that $\mathbf{F} \circ C_{|h',c}$ and $\mathbf{G} \circ C_{|h,0}$ are affine equivalent.

The claim that $\mathbf{F} \circ C_{|h',c+1}$ and $\mathbf{G} \circ C_{|h,1}$ are affine equivalent follows by considering the complimentary half-space and observing that the affine ranges of $A_1 \circ C_{|h,1}$ and $C_{|h',c+1}$ are equal. The remainder of the proof is similar. ■

Definition 2. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function represented by \mathbf{P} , let d be a positive integer and let $h \in \{0, 1\}^n$ be a HSM. Let $r_0 = SR((\mathbf{P} \circ C_{|h,0})_{(\geq d)})$, $r_1 = SR((\mathbf{P} \circ C_{|h,1})_{(\geq d)})$, $\max R = \max\{r_0, r_1\}$ and $\min R = \min\{r_0, r_1\}$.*

1. The HSM rank of h (with respect \mathbf{F} , d) is the ordered pair of integers $(\max R, \min R)$, denoted as $R_{\mathbf{F},d,h}$,

2. The attached constant of h is the value $c \in \{0, 1\}$ such that $\max R = SR((\mathbf{P} \circ C_{|h,c})(\geq d))$ (if $\max R = \min R$, the attached constant is undefined).

The lemma below states that the HSM ranks of matching HSMs are equal for affine equivalent functions.

Lemma 7. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions and let d be a positive integer. Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Let $h, h' \in \{0, 1\}^n$ be HSMs such that $h \mapsto_{(A_1)} h'$. Then $R_{\mathbf{G},d,h} = R_{\mathbf{F},d,h'}$.*

Proof. Let $c \in \{0, 1\}$ be the associated constant of $h \mapsto_{(A_1)} h'$. According to Lemma 6, $\mathbf{F} \circ C_{|h',c} \equiv \mathbf{G} \circ C_{|h,0}$ and $\mathbf{F} \circ C_{|h',c+1} \equiv \mathbf{G} \circ C_{|h,1}$.

Assume that \mathbf{F}, \mathbf{G} are represented by \mathbf{P}, \mathbf{Q} , respectively. Denote $r'_0 = SR((\mathbf{F} \circ C_{|h',c})(\geq d))$, $r'_1 = SR((\mathbf{F} \circ C_{|h',c+1})(\geq d))$, $r_0 = SR((\mathbf{G} \circ C_{|h,0})(\geq d))$, $r_1 = SR((\mathbf{G} \circ C_{|h,1})(\geq d))$.

By the above affine equivalences and Theorem 1, we have $r_0 = r'_0$ and $r_1 = r'_1$. Hence $\max(r_0, r_1) = \max(r'_0, r'_1)$ and $\min(r_0, r_1) = \min(r'_0, r'_1)$ and the lemma follows. ■

6 Rank Tables, Rank Histograms and their Properties

Definition 3. *Given a function $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and a positive integer d , define the following mappings.*

1. The rank table of \mathbf{F} with respect to d is a mapping $\mathcal{T}_{\mathbf{F},d}$, whose keys (indexes) are integer pairs $(\max R, \min R) \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$ such that $\max R \geq \min R$. It is defined as

$$\mathcal{T}_{\mathbf{F},d}(\max R, \min R) = \{h \in \{0, 1\}^n \mid R_{\mathbf{F},d,h} = (\max R, \min R)\}.$$

Moreover, along with each such HSM h , the table stores its attached constant $c \in \{0, 1\}$ (if defined).

An entry in the rank table $\mathcal{T}_{\mathbf{F},d}(\max R, \min R)$ (containing all HSMs with this rank) is called a rank group.

2. The rank histogram of \mathbf{F} with respect to d is a mapping $\mathcal{H}_{\mathbf{F},d} : \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1} \rightarrow \mathbb{Z}$ such that $\mathcal{H}_{\mathbf{F},d}(\max R, \min R) = |\mathcal{T}_{\mathbf{F},d}(\max R, \min R)|$.

To simplify our notation, in the following we refer to a HSM rank $(\max R, \min R) \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$ such that $\max R \geq \min R$ using a single symbol \mathbf{r} .

The lemma below states that if $\mathbf{F} \equiv \mathbf{G}$, then each HSM with rank \mathbf{r} for \mathbf{G} is matched in the rank group with the same HSM rank \mathbf{r} for \mathbf{F} .

Lemma 8. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions and let d be a positive integer. Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$ and let $\mathbf{r} \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$. Then, for each $h \in \mathcal{T}_{\mathbf{G},d}(\mathbf{r})$, there exists $h' \in \mathcal{T}_{\mathbf{F},d}(\mathbf{r})$ such that $h \mapsto_{(A_1)} h'$.*

Proof. By Lemma 7, given $h \in \mathcal{T}_{\mathbf{G},d}(\mathbf{r})$, its match h' under A_1 satisfies $R_{\mathbf{F},d,h'} = R_{\mathbf{G},d,h} = \mathbf{r}$ hence $h' \in \mathcal{T}_{\mathbf{F},d}(\mathbf{r})$ as claimed. \blacksquare

Lemma 9. *Let $\mathbf{F} : \{0,1\}^n \rightarrow \{0,1\}^m$, $\mathbf{G} : \{0,1\}^n \rightarrow \{0,1\}^m$ be two affine equivalent functions and let d be a positive integer. Then the rank histograms of \mathbf{F} and \mathbf{G} with respect to d are equal, namely $\mathcal{H}_{\mathbf{F},d} = \mathcal{H}_{\mathbf{G},d}$.*

Proof. Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Given a histogram entry with index \mathbf{r} , for each $h \in \mathcal{T}_{\mathbf{G},d}(\mathbf{r})$ let h' be its match $h \mapsto_{(A_1)} h'$. Then, by Lemma 8, $h' \in \mathcal{T}_{\mathbf{F},d}(\mathbf{r})$. Since $h \mapsto_{(A_1)} h'$ is a bijection, this shows that $|\mathcal{T}_{\mathbf{G},d}(\mathbf{r})| \leq |\mathcal{T}_{\mathbf{F},d}(\mathbf{r})| = \mathcal{H}_{\mathbf{F},d}(\mathbf{r})$. On the other hand, as $\mathcal{H}_{\mathbf{G},d}(\mathbf{r}) \leq \mathcal{H}_{\mathbf{F},d}(\mathbf{r})$ holds for all histogram entries \mathbf{r} and the sum of entries in both histograms is $2^n - 1$, this implies that $\mathcal{H}_{\mathbf{F},d} = \mathcal{H}_{\mathbf{G},d}$. \blacksquare

Definition 4. *Given a function $\mathbf{F} : \{0,1\}^n \rightarrow \{0,1\}^m$, a positive integer d , a HSM $h_1 \in \{0,1\}^n$ and $\mathbf{r} \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$, we define the HSM rank histogram of h_1 with respect to (or relative to) the rank group \mathbf{r} and denote it by $\mathcal{HG}_{\mathbf{F},d,h_1,\mathbf{r}}$. As the standard histogram, it is a mapping $\mathcal{HG}_{\mathbf{F},d,h_1,\mathbf{r}} : \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1} \rightarrow \mathbb{Z}$, where*

$$\mathcal{HG}_{\mathbf{F},d,h_1,\mathbf{r}}(\mathbf{r}') = |\{h_1 + h_2 \mid h_2 \in \{0,1\}^n \wedge h_1 \neq h_2 \wedge R_{\mathbf{F},d,h_2} = \mathbf{r} \wedge R_{\mathbf{F},d,h_1+h_2} = \mathbf{r}'\}|.$$

Note that unlike the (standard) rank histogram, the HSM rank histogram is defined for a specific HSM with respect to a rank group. We further remark that the HSM rank histogram of h_1 can also be defined with respect to its own the rank group (this is assured by the condition $h_1 \neq h_2$).

The following lemma equates HSM rank histograms for matching HSMs in affine equivalent functions.

Lemma 10. *Let $\mathbf{F} : \{0,1\}^n \rightarrow \{0,1\}^m$, $\mathbf{G} : \{0,1\}^n \rightarrow \{0,1\}^m$ be two affine equivalent functions and let d be a positive integer. Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Let $h_1, h'_1 \in \{0,1\}^n$ be such that $h_1 \mapsto_{(A_1)} h'_1$. Then, for every $\mathbf{r} \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$, $\mathcal{HG}_{\mathbf{G},d,h_1,\mathbf{r}} = \mathcal{HG}_{\mathbf{F},d,h'_1,\mathbf{r}}$.*

Proof. The proof follows from the fact that the mapping $\mapsto_{(A_1)}$ preserves HSM ranks for affine equivalent functions (Lemma 7), and by exploiting its additive property (Lemma 4).

Fix a HSM rank histogram entry $\mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$. Define the following two sets:

$$D_1 = \{h_1 + h_2 \mid h_2 \in \{0,1\}^n \wedge h_1 \neq h_2 \wedge R_{\mathbf{G},d,h_2} = \mathbf{r} \wedge R_{\mathbf{G},d,h_1+h_2} = \mathbf{r}'\}$$

and

$$D_2 = \{h'_1 + h'_2 \mid h'_2 \in \{0,1\}^n \wedge h'_1 \neq h'_2 \wedge R_{\mathbf{F},d,h'_2} = \mathbf{r} \wedge R_{\mathbf{F},d,h'_1+h'_2} = \mathbf{r}'\}.$$

To prove the lemma, we need to show that $|D_1| = |D_2|$. Let $h_1 + h_2 \in D_1$ and denote by $\hat{h} \in \{0, 1\}^n$ the vector such that $h_1 + h_2 \mapsto_{(A_1)} \hat{h}$. We show that $\hat{h} \in D_2$.

Since $h_1 + h_2 \mapsto_{(A_1)} \hat{h}$, by Lemma 7, $R_{\mathbf{F}, d, \hat{h}} = R_{\mathbf{G}, d, h_1 + h_2} = \mathbf{r}'$. Next, write $\hat{h} = h'_1 + (h'_1 + \hat{h})$. Since $h_1 \mapsto_{(A_1)} h'_1$ and $h_1 + h_2 \mapsto_{(A_1)} \hat{h}$, by Lemma 4, $h_2 \mapsto_{(A_1)} h'_1 + \hat{h}$, and by Lemma 7, $R_{\mathbf{F}, d, h'_1 + \hat{h}} = R_{\mathbf{G}, d, h_2} = \mathbf{r}$, giving $\hat{h} \in D_2$. Since $\mapsto_{(A_1)}$ is a bijection this implies that $|D_2| \geq |D_1|$.

As $|D_2| \geq |D_1|$ holds for all HSM histogram entries \mathbf{r}' and the sum of HSM histogram entries in both $\mathcal{HG}_{\mathbf{G}, d, h_1, \mathbf{r}}$ and $\mathcal{HG}_{\mathbf{F}, d, h'_1, \mathbf{r}}$ is equal to size of the rank group⁶ \mathbf{r} (which is equal to $\mathcal{H}_{\mathbf{F}, d}(\mathbf{r}) = \mathcal{H}_{\mathbf{G}, d}(\mathbf{r})$), the equality $|D_1| = |D_2|$ holds. \blacksquare

Definition 5. Given a function $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, a positive integer d , HSMs ranks $\mathbf{r}, \mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$, we define the HSM rank histogram multi-set of rank group \mathbf{r} with respect to rank group \mathbf{r}' as

$$\mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'} = \{\mathcal{HG}_{\mathbf{F}, d, h, \mathbf{r}'} \mid R_{\mathbf{F}, d, h} = \mathbf{r}\}.$$

The HSM rank histogram multi-set collects all the HSM histograms for HSMs in rank group \mathbf{r} with respect to the rank group \mathbf{r}' . Note that it is possible to have $\mathbf{r} = \mathbf{r}'$.

The following lemma equates HSM rank histogram multi-set in affine equivalent functions.

Lemma 11. Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions and let d be a positive integer. Then, for every $\mathbf{r}, \mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$, $\mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'} = \mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'}$.

Proof. Fix $\mathbf{r}, \mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$. We define a mapping between the elements (HSM histograms) of the multi-sets $\mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'}$ and $\mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$. Naturally, the mapping is based on the bijection $\mapsto_{(A_1)}$.

Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Let $h \in \{0, 1\}^n$ be such that $R_{\mathbf{G}, d, h} = \mathbf{r}$ which implies that $\mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}'} \in \mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'}$. Let $h' \in \{0, 1\}^n$ be the HSM such that $h \mapsto_{(A_1)} h'$. By Lemma 10, $\mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}'} = \mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'}$. Furthermore, by Lemma 7 we have $R_{\mathbf{F}, d, h'} = R_{\mathbf{G}, d, h} = \mathbf{r}$, hence $\mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}'} = \mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'} \in \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$. Since $\mapsto_{(A_1)}$ is a bijection, we obtain $\mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'} \subseteq \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$ as multi-sets.

On the other hand, the number of elements (HSM histograms) in both multi-sets is equal to the size of the rank group \mathbf{r} (which is equal to $\mathcal{H}_{\mathbf{F}, d}(\mathbf{r}) = \mathcal{H}_{\mathbf{G}, d}(\mathbf{r})$), hence $\mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'} = \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$. \blacksquare

Definition 6. Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, let d be a positive integer and let $\mathbf{r}, \mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$. A HSM $h \in \{0, 1\}^n$ such that $R_{\mathbf{F}, d, h} = \mathbf{r}$ is called unique (with respect to $\mathbf{F}, d, \mathbf{r}'$) if $\mathcal{HG}_{\mathbf{F}, d, h, \mathbf{r}'} \in \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$ has multiplicity 1 in this multi-set.

⁶ Unless the HSM rank of h_1 is \mathbf{r} , in which case the sum of HSM histogram entries is $\mathcal{H}_{\mathbf{F}, d}(\mathbf{r}) - 1$.

The following theorem establishes the importance of unique HSMs in recovering matchings between HSMs for affine equivalent functions.

Theorem 2. *Let $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be two affine equivalent functions and let d be a positive integer. Then for every $\mathbf{r}, \mathbf{r}' \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$, if $h \in \{0, 1\}^n$ (such that $R_{\mathbf{G}, d, h} = \mathbf{r}$) is unique with respect to $\mathbf{G}, d, \mathbf{r}'$, then the following statements hold:*

1. *There exists $h' \in \{0, 1\}^n$ such that $R_{\mathbf{F}, d, h'} = \mathbf{r}$ and h' is unique with respect to $\mathbf{F}, d, \mathbf{r}'$.*
2. *$\mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}'} = \mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'}$.*
3. *Assume that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. Then, $h \mapsto_{(A_1)} h'$. Moreover, if the attached constants of h, h' are defined and equal to c, c' , respectively, then the associated constant of $h \mapsto_{(A_1)} h'$ is $c + c'$.*

Proof. By Lemma 11, we have equality of the multi-sets $\mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'} = \mathcal{HM}_{\mathbf{G}, d, \mathbf{r}, \mathbf{r}'}$ which immediately implies the first two statements. Denote by $h'' \in \{0, 1\}^n$ the HSM such that $h \mapsto_{(A_1)} h''$. To complete the proof of the third statement we show that $h'' = h'$.

By Lemma 7, we have $R_{\mathbf{F}, d, h''} = R_{\mathbf{G}, d, h} = \mathbf{r}$. Hence $\mathcal{HG}_{\mathbf{F}, d, h'', \mathbf{r}} \in \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$ (and also $\mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'} \in \mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$ from the first statement). Since h' is unique with respect to $\mathbf{F}, d, \mathbf{r}'$, then $\mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'}$ has multiplicity 1 in $\mathcal{HM}_{\mathbf{F}, d, \mathbf{r}, \mathbf{r}'}$. Thus if we show that $\mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'} = \mathcal{HG}_{\mathbf{F}, d, h'', \mathbf{r}}$, then $h'' = h'$ must hold.

According to Lemma 10, $\mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}} = \mathcal{HG}_{\mathbf{F}, d, h'', \mathbf{r}}$ and by the second statement we obtain $\mathcal{HG}_{\mathbf{F}, d, h', \mathbf{r}'} = \mathcal{HG}_{\mathbf{G}, d, h, \mathbf{r}} = \mathcal{HG}_{\mathbf{F}, d, h'', \mathbf{r}}$ as required.

Finally, we examine the attached constants c, c' of h, h' , respectively (assuming they are defined). If $c = c'$, then the affine ranges of $A_1 \circ C_{|h, 0}$ and $C_{|h', 0}$ are equal implying that the associated constant of $h \mapsto_{(A_1)} h'$ is $0 = c + c'$. Otherwise $c = c' + 1$ and the affine ranges of $A_1 \circ C_{|h, 0}$ and $C_{|h', 1}$ are equal implying that the associated constant of $h \mapsto_{(A_1)} h'$ is $1 = c + c'$. ■

7 Analysis of the Distribution of Rank Histogram Entries for Random Permutations

In this section we analyze the distribution of entries of the rank histogram $\mathcal{H}_{\mathbf{F}, d}$ for a random permutation $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$. The analysis is performed for $d = n - 2$, which is the value that we use in our algorithm as explained in detail next.

Assume that \mathbf{F} is represented by $\mathbf{P} = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^n$. For a given $h \in \{0, 1\}^n$, we consider $SR((\mathbf{P} \circ C_{|h, 0})_{(\geq n-2)})$ and $SR((\mathbf{P} \circ C_{|h, 1})_{(\geq n-2)})$. For $c \in \{0, 1\}$, every one of the n polynomials of $(\mathbf{P} \circ C_{|h, c})_{(\geq n-2)}$ has $n - 1$ variables (the number of variables in \mathbf{P} is reduced by 1 after composition with $C_{|h, c}$). Hence, the number of possible non-zero monomial coefficients in each such polynomial is $\binom{n-1}{n-1} + \binom{n-1}{n-2} = 1 + n - 1 = n$. Therefore, $(\mathbf{P} \circ C_{|h, c})_{(\geq n-2)}$ can be represented by an $n \times n$ Boolean matrix and we are interested in its rank.

Choosing $d = n - 1$ would leave at most one non-zero monomial which almost always would be present in $(\mathbf{P} \circ C_{|h,c})_{(\geq n-1)}$. Hence, essentially all HSMs would fall into a single rank group and the affine equivalence algorithm would not be able to distinguish and match them. On the other hand, choosing $d \leq n - 3$ would leave $\Omega(n^2)$ non-zero monomials and $(\mathbf{P} \circ C_{|h,c})_{(\geq r)}$ would almost always have full rank, leading once again to a single rank group. We conclude that $d = n - 1$ is indeed the optimal choice.

Our analysis is based on the following heuristic assumption.

Assumption 1 *For a random permutation $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ represented by \mathbf{P} , for every $h \in \{0, 1\}^n$ and $c \in \{0, 1\}$, the entries of the $n \times n$ Boolean matrix $(\mathbf{P} \circ C_{|h,c})_{(\geq n-2)}$ are uniform independent random variables.*

The $n \times n$ Boolean matrix $(\mathbf{P} \circ C_{|h,c})_{(\geq n-2)}$ is indeed uniform for a random function \mathbf{F} (rather than a random permutation), given any $h \in \{0, 1\}^n$ and $c \in \{0, 1\}$. However, even for a random function the Boolean matrices obtained for different h, c values are correlated. Nevertheless, these correlations (and the fact that \mathbf{F} is a permutation) do not seem to have a noticeable influence on our algorithm in practice (as we demonstrate in Section 8.5 and Appendix B).

The rank of random matrices is a well-studied problem. For large n and a non-negative integer $r \leq n$, we denote the probability that a random Boolean $n \times n$ matrix has rank r by β_r . We can lower bound β_r by considering the event where we first select r linearly independent rows to form a subspace of size 2^r (which occurs with constant probability) and then select the remaining $n - r$ rows within this subspace (which occurs with probability $2^{-(r-n)^2}$). This gives a lower bound of $\Omega(2^{-(r-n)^2})$ on β_r . The exact formula is given by the theorem below, taken and adapted from [12].

Theorem 3 ([12], page 126, adapted). *For $n \rightarrow \infty$, the probability that a random Boolean $n \times n$ matrix has rank r is*

$$\beta_r = 2^{-(r-n)^2} \cdot \alpha \cdot \prod_{i=1}^{n-r} (1 - 1/2^i)^{-2}, \quad (2)$$

where $\alpha = \prod_{i=1}^{\infty} (1 - 1/2^i) \approx 0.2888$.

Since $\alpha \leq \alpha \cdot \prod_{i=1}^{n-r} (1 - 1/2^i)^{-2} < 1/\alpha$, the initial probability estimation of $\approx 2^{-(r-n)^2}$ is correct up to a small constant. We also note that (2) is a good estimation even for relatively small values of n (e.g., $n \geq 8$).

Let

$$p_{\max R, \min R} = \begin{cases} 2\beta_{\max R} \beta_{\min R} & \text{if } \min R < \max R \\ \beta_{\max R} \beta_{\min R} & \text{otherwise } (\min R = \max R), \end{cases} \quad (3)$$

where

$$\beta_{\max R} \beta_{\min R} = \alpha^2 \cdot 2^{-(\max R - n)^2 - (\min R - n)^2} \cdot \prod_{i=1}^{n - \max R} (1 - 1/2^i)^{-2} \cdot \prod_{i=1}^{n - \min R} (1 - 1/2^i)^{-2}.$$

Then, based on Assumption 1 and Theorem 3, for every $(maxR, minR) \in \mathbb{Z}_{m+1} \times \mathbb{Z}_{m+1}$ such that $maxR \geq minR$, given $h \in \{0, 1\}^n$, we have $Pr[R_{\mathbf{F}, n-2, h} = (maxR, minR)] \approx p_{maxR, minR}$. Hence, according to Assumption 1, the entries of $\mathcal{H}_{\mathbf{F}, n-2}$ are distributed multinomially, with parameter 2^n (the number of HSMs⁷) and probabilities given by $p_{maxR, minR}$. In particular, each individual histogram entry $\mathcal{H}_{\mathbf{F}, n-2}(maxR, minR)$ is distributed binomially with parameter 2^n and probability $p_{maxR, minR}$.

Experimental results that support this conclusion are given in Appendix B).

Asymptotic Analysis of Specific Histogram Entries For large n , the binomial variable $\mathcal{H}_{\mathbf{F}, n-2}(maxR, minR)$ is with high probability very close to its expectation, which is about

$$2^n \cdot p_{maxR, minR}.$$

If we ignore constant multiplicative factors, we can approximate this expectation by

$$2^n \cdot 2^{-(maxR-n)^2 - (minR-n)^2}, \quad (4)$$

as $p_{maxR, minR} \approx 2^{-(maxR-n)^2 - (minR-n)^2}$.

We now approximate (up to constant multiplicative factors) the expected values of two specific histogram entries which will be useful for our algorithm. Denote $\gamma_n = \lfloor (n/2)^{1/2} \rfloor$, and let $\mathbf{r}_1 = (n+1 - \gamma_n, n - \gamma_n)$ and $\mathbf{r}_2 = (n, n - \gamma_n)$. Define the random variables $S_1 = \mathcal{H}_{\mathbf{F}, d}(\mathbf{r}_1)$ and $S_2 = \mathcal{H}_{\mathbf{F}, d}(\mathbf{r}_2)$. Below, we estimate their expected values according to (4).

Write $\gamma_n = \lfloor (n/2)^{1/2} \rfloor = (n/2)^{1/2} - k$, where $0 \leq k < 1$. Hence, with very high probability we have $S_2 = \mathcal{H}_{\mathbf{F}, d}(\mathbf{r}_2) = \mathcal{H}_{\mathbf{F}, d}(n, n - \gamma_n) \approx 2^n \cdot p_{n, n - \gamma_n} \approx 2^n \cdot 2^{-(\gamma_n)^2} = 2^n \cdot 2^{-((n/2)^{1/2} - k)^2} = 2^n \cdot 2^{-n/2 + 2k(n/2)^{1/2} - k^2} = 2^{n/2 + O(n^{1/2})}$. Therefore, S_2 is close to $2^{n/2}$.

Similarly $S_1 = \mathcal{H}_{\mathbf{F}, d}(\mathbf{r}_1) = \mathcal{H}_{\mathbf{F}, d}(n+1 - \gamma_n, n - \gamma_n) \approx 2^n \cdot p_{n+1 - \gamma_n, n - \gamma_n} \approx 2^n \cdot 2^{-(\gamma_n - 1)^2 - (\gamma_n)^2} = 2^n \cdot 2^{-2(\gamma_n)^2 + 2\gamma_n - 1} = 2^n \cdot 2^{-n + (4k+2)(n/2)^{1/2} - 2k^2 - 2k - 1} = 2^{\Theta(n^{1/2})}$. Hence S_1 is sub-exponential in n .

8 Details of the New Affine Equivalence algorithm

In this section we describe and analyze our new affine equivalence algorithm. We start with a description of the auxiliary algorithms it uses.

8.1 The Rank Table and Histogram Algorithm

For $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ represented by $\mathbf{P} = \{P^{(i)}(x[1], \dots, x[n])\}_{i=1}^n$, the following algorithm computes the rank histogram $\mathcal{H}_{\mathbf{F}, d}$ and rank table $\mathcal{T}_{\mathbf{F}, d}$ for $d = n - 2$. The algorithm is given as input $\mathbf{P}_{\geq (n-2)}$.

⁷ More accurately, the parameter is $2^n - 1$ as HSMs are non-zero.

1. For each non-zero HSM $h \in \{0, 1\}^n$:
 - (a) Compute $R_{\mathbf{F}, n-2, h} = (maxR, minR)$ as follows. Compute $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,0}(\geq n-2)) = (\mathbf{P} \circ C_{|h,0}(\geq n-2))$ and calculate its symbolic rank r_0 using Gaussian elimination. Similarly, compute $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,1}(\geq n-2)) = (\mathbf{P} \circ C_{|h,1}(\geq n-2))$ and its symbolic rank r_1 . Let $maxR = \max\{r_0, r_1\}$ and $minR = \min\{r_0, r_1\}$.
 - (b) Insert h into $\mathcal{T}_{\mathbf{F}, n-2}(maxR, minR)$, along with the value of the attached constant $c \in \{0, 1\}$ such that $maxR = SR((\mathbf{F} \circ C_{|h,c}(\geq n-2)))$ (if $maxR > minR$). In addition, increment entry $\mathcal{H}_{\mathbf{F}, n-2}(maxR, minR)$.

Note that $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,0}(\geq n-2)) = (\mathbf{P} \circ C_{|h,0}(\geq n-2))$ holds according to the first part of Lemma 1.

The time complexity of the algorithm depends on how a polynomial is represented. Here, we represent it using a bit array that specifies the values of its monomial coefficients.

We first analyze the complexity of computing the composition $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,c}(\geq n-2))$ in Step 1.(a), which is performed for each non-zero $h \in \{0, 1\}^n$ and $c \in \{0, 1\}$. Each of the n polynomials of $\mathbf{P}_{(\geq n-2)}$ contains at most $\binom{n}{n} + \binom{n}{n-1} + \binom{n}{n-2} < n^2$ non-zero monomials. As described in Section 2, computing the composition $\mathbf{P}_{(\geq n-2)} \circ C_{|h,c}$ requires substituting one of the n variables with a linear combination of the remaining $n - 1$ variables (while renaming the variables of the monomials).

In total, for each polynomial of $\mathbf{P}_{(\geq n-2)} \circ C_{|h,c}$, we compose its n^2 monomials with a linear combination of size n , which requires $n^2 \cdot n = n^3$ bit operations. However, as we are only interested in monomials of degree at least $n - 2$, the outcome $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,0}(\geq n-2))$ is a polynomial of at most $\binom{n-1}{n-1} + \binom{n-1}{n-2} = 1 + n - 1 = n$ monomials, and the average complexity can be easily reduced to n^2 using low-level optimization techniques.⁸

In conclusion, the average complexity of computing the n polynomials of $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,0}(\geq n-2))$ is $n \cdot n^2 = n^3$ and the total time spent on composition is $n^3 \cdot 2^n$ bit operations (up to multiplicative constant factors). Similarly, Gaussian elimination requires n^3 bit operations, hence the total time complexity of the algorithm is $n^3 \cdot 2^n$ bit operations.

8.2 The Unique HSM Algorithm

The following algorithm computes the HSM rank histogram multi-set $\mathcal{HM}_{\mathbf{F}, n-2, \mathbf{r}, \mathbf{r}'}$ and uses it to compute a set of unique HSMs, denoted by $U_{\mathbf{F}}$. This set contains triplets of the form $(h, c, \mathcal{HG}_{\mathbf{F}, n-2, h, \mathbf{r}'})$, where $h \in \mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r})$ is unique with respect to $\mathbf{F}, n - 2, \mathbf{r}'$ and $c \in \{0, 1\}$ is its attached constant. Note that

⁸ For example, we can exploit the fact that the composition $(\mathbf{P}_{(\geq n-2)} \circ C_{|h,c}(\geq n-2))$ is computed for each $h \in \{0, 1\}^n$ and $c \in \{0, 1\}$, and the effect of flipping a bit in h on the outcome can be precomputed. Consequently, we iterate over $h \in \{0, 1\}^n$ using a Gray code.

for the attached constant to be defined, we must have $\max R > \min R$, where $(\max R, \min R) = \mathbf{r}$.

The algorithm is given as input the rank table $\mathcal{T}_{\mathbf{F}, n-2}$ and rank group indexes \mathbf{r}, \mathbf{r}' .

1. For each $h \in \mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r})$, compute $\mathcal{HG}_{\mathbf{F}, n-2, h, \mathbf{r}'}$ as follows:
 - (a) for each $h' \in \mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r}')$:
 - i. Compute $h + h'$, find its rank $\mathbf{r}'' = R_{\mathbf{F}, n-2, h+h'}$ in $\mathcal{T}_{\mathbf{F}, n-2}$ and increment $\mathcal{HG}_{\mathbf{F}, n-2, h, \mathbf{r}'}(\mathbf{r}'')$.
 - (b) Insert $\mathcal{HG}_{\mathbf{F}, n-2, h, \mathbf{r}'}$ along with h and its attached constant c into the multi-set $\mathcal{HM}_{\mathbf{F}, n-2, \mathbf{r}, \mathbf{r}'}$.
2. For each unique HSM h in $\mathcal{HM}_{\mathbf{F}, n-2, \mathbf{r}, \mathbf{r}'}$, add the triplet $(h, c, \mathcal{HG}_{\mathbf{F}, n-2, h, \mathbf{r}'})$ to $U_{\mathbf{F}}$.

The time complexity of the algorithm is the product of sizes of the rank groups $|\mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r})| \cdot |\mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r}')| = \mathcal{H}_{\mathbf{F}, n-2}(\mathbf{r}) \cdot \mathcal{H}_{\mathbf{F}, n-2}(\mathbf{r}')$.

Since the goal of the affine equivalence algorithm will be to find n linearly independent unique HSMs, it is useful to estimate their number. In Appendix C we lower bound the expected number of unique HSMs in $\mathcal{HM}_{\mathbf{F}, n-2, \mathbf{r}, \mathbf{r}'}$ asymptotically (ignoring constant factors) based on Assumption 1, given that \mathbf{F} is a random permutation. More specifically, we obtain the lower bound of $S - S^2/\sqrt{S'}$, where $S = \mathcal{H}_{\mathbf{F}, n-2}(\mathbf{r})$, and $S' = \mathcal{H}_{\mathbf{F}, n-2}(\mathbf{r}')$.

8.3 The Affine Transformation A_1 Recovery Algorithm

Assume that we have affine equivalent functions \mathbf{F} and \mathbf{G} such that $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$ and $A_1(x) = L(x) + a$.

The following algorithm recovers A_1 using sets of unique HSMs $U_{\mathbf{F}}$ and $U_{\mathbf{G}}$, computed with the previous algorithm of Section 8.2 (where its invocations for \mathbf{F} and \mathbf{G} use the same parameters values of \mathbf{r}, \mathbf{r}'). Since \mathbf{F} and \mathbf{G} are affine equivalent, the HSM rank histograms of the HSMs in these sets have to match according to Theorem 2. Each equal HSM histogram pair reveals the matching $h \mapsto_{A_1} h'$ and its associated constant is revealed by adding the attached constants of h, h' (which are defined in case $\max R > \min R$, where $(\max R, \min R) = \mathbf{r}$), again by Theorem 2.

Each matching $h \mapsto_{A_1} h'$ and its associated constant give linear equations on the columns of L and on a (respectively) according to Lemma 5. Assuming that $U_{\mathbf{F}}$ and $U_{\mathbf{G}}$ contain n linearly independent unique HSMs, A_1 is recovered by linear algebra.

1. Allocate $n+1$ linear equation systems $\{E_i\}_{i=1}^{n+1}$, each of dimension $n \times n$: the first n equation systems are on the columns $L[i]$ of L and the final equation system E_{n+1} is on a .
2. Locate n linearly independent HSMs in $U_{\mathbf{G}}$. For each such HSM h :
 - (a) Recover the triplet $(h, c, \mathcal{HG}_{\mathbf{G}, n-2, h, \mathbf{r}'})$ from $U_{\mathbf{G}}$.

- (b) Search $U_{\mathbf{F}}$ for a triplet $(h', c', \mathcal{H}\mathcal{G}_{\mathbf{F}, n-2, h', r'})$ such that $\mathcal{H}\mathcal{G}_{\mathbf{F}, n-2, h', r'} = \mathcal{H}\mathcal{G}_{\mathbf{G}, n-2, h, r'}$. If no match exists, return “Not Equivalent”.
 - (c) Based on Lemma 5, for $i = 1, 2, \dots, n$ add equation $h'(L[i]) = h[i]$ to E_i .
 - (d) Based on Lemma 5, add equation $h'(a) = c + c'$ to E_{n+1} .
3. Solve each one of $\{E_i\}_{i=1}^{n+1}$, recover A_1 and return its matrix L and vector a .

The complexity of the algorithm is about $n \cdot n^3 = n^4$ bit operations, which is polynomial in n . Since we solve the same linear equation (with coefficients given by the h' vectors) $n + 1$ times with different constants, the complexity can be reduced to n^3 by inverting the matrix which defines the linear equations.

8.4 The New Affine Equivalence Algorithm

We describe the new affine equivalence algorithm below. Let $\mathbf{r}_1 = (n+1-\gamma_n, n-\gamma_n)$ and $\mathbf{r}_2 = (n, n-\gamma_n)$ for $\gamma_n = \lfloor (n/2)^{1/2} \rfloor$, as defined in Section 7.

1. Given $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, $\mathbf{G} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, compute their corresponding ANF representations $\mathbf{P}_{\geq(n-2)}$ and $\mathbf{Q}_{\geq(n-2)}$.
2. Run the algorithm of Section 8.1 to compute the rank table $\mathcal{T}_{\mathbf{F}, n-2}$ and rank histogram $\mathcal{H}_{\mathbf{F}, n-2}$ for \mathbf{F} using $\mathbf{P}_{\geq(n-2)}$, and similarly compute $\mathcal{T}_{\mathbf{G}, n-2}$ and $\mathcal{H}_{\mathbf{G}, n-2}$ for \mathbf{G} . If $\mathcal{H}_{\mathbf{F}, n-2} \neq \mathcal{H}_{\mathbf{G}, n-2}$, return “Not Equivalent”.
3. Run the unique HSM algorithm of Section 8.2 for \mathbf{F} on inputs $\mathcal{T}_{\mathbf{F}, n-2}$ and $\mathbf{r}_1, \mathbf{r}_2$ defined above, and obtain the set $U_{\mathbf{F}}$. Similarly, obtain the set for $U_{\mathbf{G}}$ by running this algorithm on inputs $\mathcal{T}_{\mathbf{G}, n-2}$ and $\mathbf{r}_1, \mathbf{r}_2$. If $|U_{\mathbf{F}}| \neq |U_{\mathbf{G}}|$, return “Not Equivalent”. Otherwise, if $U_{\mathbf{F}}$ does not contain n linearly independent HSMs, return “Fail”.
4. Run the affine transformation recovery algorithm of Section 8.3 on inputs $U_{\mathbf{F}}$ and $U_{\mathbf{G}}$. If it returns “Not Equivalent”, return the same output. Otherwise, it returns a candidate for A_1 .
5. Recover a candidate for $A_2 = L_2(x) + a_2$ by evaluating inputs $v \in \{0, 1\}^n$ to $\mathbf{F} \circ A_1$ and \mathbf{G} : each input v gives n linear equations on L_2 and a_2 . Hence, after a bit more than n evaluations, we expect the linear equation system to have a single solution which gives a candidate for A_2 .
6. Test the candidates A_1, A_2 by equating the evaluations of \mathbf{G} and $A_2 \circ \mathbf{F} \circ A_1$ on all 2^n possible inputs. If $\mathbf{G}(v) \neq A_2 \circ \mathbf{F} \circ A_1(v)$ for some $v \in \{0, 1\}^n$, return “Not Equivalent”. Otherwise, return A_1, A_2 .

The correctness of the algorithm follows from the correctness of the sub-procedures it executes and from the results obtained so far. In particular, Step 2 is correct according to Lemma 9, while the correctness of Step 3 is based on Theorem 2. The correctness of the final step is trivial.

Step 3 is the most complex to analyze in terms of success probability and complexity (which is the product $\mathcal{H}_{\mathbf{F},n-2}(\mathbf{r}_1) \cdot \mathcal{H}_{\mathbf{F},n-2}(\mathbf{r}_2)$). We first focus on the time complexity analysis of the other steps.

The complexities of steps 4, 5 are at most polynomial in n and can be neglected. Step 1 interpolates the 2^n ANF coefficients for each of the n output bits of \mathbf{F}, \mathbf{G} . Each such interpolation can be performed in $n2^n$ bit operations using the Moebius transform [11]. Hence this step requires n^22^n bit operations and 2^n function evaluations in total. The complexity of Step 2 was shown to be n^32^n bit operations, while the complexity of Step 6 is 2^n function evaluations.

In total, the time complexity of the algorithm is at most n^32^n bit operations and 2^n function evaluations, assuming that the complexity of Step 3 does not dominate the algorithm (as we show below).

The memory complexity is 2^n words of n bits, but it can be significantly reduced in some cases as described in Section 8.6.

Asymptotic Analysis of the Unique HSM Algorithm As in Section 7, denote $S_1 = \mathcal{H}_{\mathbf{F},d}(\mathbf{r}_1)$ and $S_2 = \mathcal{H}_{\mathbf{F},d}(\mathbf{r}_2)$ and recall that their expected values are $2^{\Theta(n^{1/2})}$ and $2^{n/2+O(n^{1/2})}$, respectively.

The expected asymptotic complexity of the unique HSM algorithm is therefore at most

$$S_1 \cdot S_2 = 2^{n/2+O(n^{1/2})} \ll 2^n.$$

Hence, the complexity of Step 3 is negligible compared to the complexity of the remaining steps of the affine equivalence algorithm described above.

According to the analysis of the unique HSM algorithm given in Appendix C (see (6)), the asymptotic lower bound on the expected number of unique HSMs in $\mathcal{HM}_{\mathbf{F},n-2,\mathbf{r},\mathbf{r}'}$ is

$$S_1 - (S_1)^2 / \sqrt{S_2} > 2^{\Theta(n^{1/2})} - 2^{\Theta(n^{1/2})} / 2^{n/4+O(n^{1/4})} = 2^{\Theta(n^{1/2})} \gg n.$$

Out of these unique HSMs, n are very likely to be linearly independent. This shows that asymptotically the algorithm succeeds with overwhelming probability.

We remark that there are many possible ways to select the rank group indexes \mathbf{r}_1 and \mathbf{r}_2 that give similar results.

8.5 Experimental results

In practice we do not pre-fix the rank groups of \mathbf{F}, \mathbf{G} for which we run the unique HSM algorithm. Instead, we select a *reference rank group* \mathbf{r}' such that $|\mathcal{T}_{\mathbf{F},n-2}(\mathbf{r}')| \approx 2^{n/2}$ (as \mathbf{r}_2 defined above). We then iterate over the rank groups \mathbf{r} from the smallest to the largest, while collecting unique HSMs using repeated executions of the unique HSM algorithm with inputs \mathbf{r}, \mathbf{r}' . We stop once we collect n linearly independent unique HSMs. This practical variant is more flexible and succeeds given that the variant above succeeds.

We implemented the algorithm and tested it for various values of $8 \leq n \leq 28$. In each trial we first selected the permutation \mathbf{F} uniformly at random. We then chose invertible affine mappings A_1, A_2 uniformly at random and defined $\mathbf{G} = A_2 \circ \mathbf{F} \circ A_1$. After calculating these inputs, we executed the algorithm and verified that it correctly recovered A_1, A_2 .

Following this initial verification, our goal was to collect statistics that support the asymptotic complexity analysis of the unique HSM algorithm above. For this purpose, we selected the permutation \mathbf{F} at random and calculated the success rate and complexity of Step 3, which executes the unique HSM algorithm on \mathbf{F} (after running steps 1,2). If Step 3 succeeds to return n linearly independent unique HSMs with a certain complexity for \mathbf{F} , then it would succeed with identical complexity on any linearly equivalent \mathbf{G} , hence analyzing a single permutation is sufficient for the purpose of gathering statistics.

Our results for $n \in \{8, 12, 16, 20, 24, 28\}$ are summarized in Table 1. This table shows that all the trials for the various choices of n were successful. In terms of complexity, for $n = 8$, the unique HSM algorithm had to iterate over $2^{10.5} > 2^8$ HSMs on average in order to find 8 unique linearly independent HSMs. This relatively high complexity is due to the fact that our asymptotic analysis ignores constants whose effect is more pronounced for smaller values of n . Nevertheless, the complexity of the algorithm for $n = 8$ in terms of bit operations remains roughly $8^3 \cdot 2^8$, as the unique HSM algorithm does not perform linear algebra.

For $n \geq 12$, the average complexity of the unique HSM algorithm is below 2^n , and this gap increases as n grows (as predicted by the asymptotic analysis). Note that the complexity drops in two cases (between $n = 16$ and $n = 20$ and between $n = 24$ and $n = 28$) since for larger n we have more non-empty rank groups of various sizes and hence more flexibility in the algorithm (which happens to be quite substantial for $n = 20$ and $n = 28$). Finally, we note that we did not optimize the index r' of the reference rank group and better options that improve the complexities are likely to exist. However, since the unique HSM algorithm does not dominate the overall complexity, such improvements would have negligible effect.

In addition to the experiments on random permutations, we also performed simulations on random functions and obtained similar results.

Table 1. Experimental Results for the Unique HSM Algorithm

n	Number of Trials	Number of Successful Trials	Average Complexity
8	1000	1000	$1486 \approx 2^{10.5}$
12	1000	1000	$3229 \approx 2^{11.7}$
16	1000	1000	$25599 \approx 2^{14.6}$
20	1000	1000	$15154 \approx 2^{13.9}$
24	1000	1000	$126777 \approx 2^{17}$
28	100	100	$40834 \approx 2^{15.3}$

8.6 Additional Variants of the New Affine Equivalence Algorithm

We describe several variants of the affine equivalence algorithm.

Using Rank Group Sums The first additional variant we describe uses the rank tables of \mathbf{F}, \mathbf{G} to directly recover several matchings in the initial stage of the algorithm. It is based on the observation that for each non-empty rank group \mathbf{r} , the HSM obtained by summing of all HSMs in $\mathcal{T}_{\mathbf{G}, n-2}(\mathbf{r})$ has to match (under $\mapsto_{(A_1)}$) the HSM obtained by summing of all HSMs in $\mathcal{T}_{\mathbf{F}, n-2}(\mathbf{r})$ due to the additive property of the HSM bijection. Simple analysis (based on Assumption 1 and backed up by experimental results) shows that the number of non-empty rank groups for a random permutation \mathbf{F} is at least $n/4$ with very high probability. Hence we can initially recover at least $n/4$ matchings using this approach. There are several ways to recover the remaining matchings by exploiting the fact that we have essentially reduced the size of the problem from 2^n to at most $2^{3n/4}$. We can also continue in a similar way, further exploiting additive properties of the bijection: we take a uniquely matched HSM pair h, h' . For \mathbf{G} , we compute the *HSM rank table* for h with respect to some rank group \mathbf{r}' by adding it to all HSMs in this group. We do the same for \mathbf{F} by computing the HSM rank table of h' with respect to \mathbf{r}' . As in the initial observation, the sum of HSMs in each non-empty rank group of these smaller tables for \mathbf{F}, \mathbf{G} match under $\mapsto_{(A_1)}$, revealing additional matchings. We repeat this process for several uniquely matched HSM pairs (computing additional HSM rank tables) until we identify the required n linearly independent matchings.

Reducing the Memory complexity The memory complexity of the algorithm is about 2^n words of n bits. If the functions \mathbf{F}, \mathbf{G} are given as truth tables, then the memory complexity cannot be reduced by much. However, if we are given access to \mathbf{F}, \mathbf{G} via oracles (e.g., they are implemented by block ciphers with a fixed key), then we can significantly reduce the memory complexity with no substantial effect on the time complexity.

First, instead of using the Moebius transform in Step 1 in order to interpolate all the coefficients of \mathbf{F}, \mathbf{G} , we simply interpolate each of the relevant $\approx n^2$ coefficients of degree at least $n - 2$ independently, increasing the complexity of Step 1 by a factor of about n . Next, in Step 2 we do not store the entire rank table, but only the relevant rank groups with indexes \mathbf{r}_1 and \mathbf{r}_2 . As a result, we now have to recompute the ranks of $S_1 \cdot S_2$ HSMs in Step 3, but this requires much lower complexity than 2^n .

Overall, the memory complexity of this low-memory variant is dominated by the size of largest rank group stored in memory S_2 , which is bit more than $2^{n/2}$. Finally, by a different choice of rank groups of indexes \mathbf{r}_1 and \mathbf{r}_2 , it is possible reduce the memory to be sub-exponential in n .

Multiple Solutions to the Affine Equivalence Problem Consider the case where there are two or more solutions of the form $(A_1^{(i)}, A_2^{(i)})$ to an instance

of the affine equivalence problem \mathbf{F}, \mathbf{G} . This may occur (for example) if \mathbf{F} is self-affine equivalent, namely, there exist (A_1, A_2) (that are not both identities) such that $\mathbf{F} = A_2 \circ \mathbf{F} \circ A_1$. We note that this case is extremely unlikely if \mathbf{F} is chosen uniformly at random for $n \geq 8$, but it may occur for specific choices (e.g., the AES Sbox is self-affine equivalent).

In case of multiple solutions, a straightforward application of the affine equivalence algorithm would fail, as a HSM h would most likely match a different $h^{(i)}$ for each solution $A_1^{(i)}$, namely $h \mapsto_{A_1^{(i)}} h^{(i)}$. Consequently, we would not be able to find sufficiently many unique HSMs in Step 4. However, we can tweak the algorithm to deal with this case by working on each match $h \mapsto_{A_1^{(i)}} h^{(i)}$ separately. More specifically, according to Lemma 6 we know that $\mathbf{F} \circ C_{|h^{(i)}, c}^1 \equiv \mathbf{G} \circ C_{|h, 0}$ and we can apply the algorithm recursively on these functions.

Affine Equivalences Among a Set of Functions We consider a generalization of the affine equivalence problem that was described in [3]. Given a set of K functions $\{\mathbf{F}_i\}_{i=1}^K$, our goal is to partition them into groups of affine equivalent functions. The naive approach is to run the affine equivalence algorithm on each pair of functions, which results in complexity of $K^2 \cdot n^{32^n}$.

We can improve this complexity by noticing that up to Step 4 of the affine equivalence algorithm the functions \mathbf{F}, \mathbf{G} are analyzed independently. In particular, we can compute the rank histogram $\mathcal{H}_{\mathbf{F}_i, n-2}$ for each function \mathbf{F}_i independently (as done in Step 2) in time n^{32^n} , and then sort the functions and classify them according to their rank histograms.⁹ This reduces the time complexity to about $K \cdot n^{32^n} + \tilde{O}(K^2)$ (where \tilde{O} hides a small polynomial factor in n), improving upon the time complexity of $K \cdot n^{32^{2n}} + \tilde{O}(K^2)$, obtained in [3].

9 Applications

We describe applications of the affine equivalence algorithm and then focus on additional applications of the new objects and algorithms defined in this paper.

9.1 Applications of the New Affine Equivalence Algorithm

Algorithms for the affine equivalence problem are useful in several contexts such as classification of Sboxes [6, 13], producing equivalent representations of block ciphers [3] and attacking white-box ciphers [14]. In all of these contexts, if the goal is to apply the algorithm a few times to functions with a small domain size n , then the main algorithm of Biryukov et al. [3] is already practical and there is little to be gained by using our algorithm.

On the other hand, our algorithm may provide an advantage if the goal is to solve the affine equivalence problem on functions with a larger domain

⁹ We can also attach more data to each function by computing HSM histogram multi-sets between groups $\mathcal{HM}_{\mathbf{F}_i, n-2, r, r'}$.

sizes (e.g., the domain size of the CAST Sbox [1] is $n = 32$). Furthermore, our algorithm may be beneficial if we need to solve the affine equivalence problem for many functions with domain size $n \geq 8$. For example, if we want to classify a large set of 8-bit Sboxes produced based on some design criteria, we can use the variant that searches for affine equivalences among a set of functions (described in Section 8.6).

An additional application (which is also described in [3]) is cryptanalysis of a generalization of the Even-Mansour scheme. The original scheme [10] builds a block cipher using a public permutation $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a pair of n -bit keys k_1, k_2 by defining the encryption function on a plaintext $p \in \{0, 1\}^n$ as $\mathbf{E}(p) = \mathbf{F}(p+k_1)+k_2$. Breaking the scheme may be considered as a special case of solving the affine equivalence problem where the linear matrices are identities. Thus, in the generalized scheme, arbitrary affine transformations A_1, A_2 are used as the key and the encryption function is defined as $\mathbf{E}(p) = A_2 \circ \mathbf{F} \circ A_1(p)$. Clearly, breaking the generalized Even-Mansour scheme reduces to solving the affine equivalence problem. The currently best known attack on this scheme (given in [3]) requires about $2^{3n/2}$ time and memory. It uses a birthday paradox based approach that generalizes Daemen’s attack on the original Even-Mansour cipher [8]. Hence, we improve the complexity of the best known attack on the generalized Even-Mansour cipher from about $2^{3n/2}$ to 2^n .

9.2 Additional Applications

We describe additional applications of the rank table and histogram objects defined in this paper, and the algorithm used to compute them.

Application to Decomposition of the ASASA Construction The ASASA construction is an SP-network that consists of three secret affine layers (A) interleaved with two secret Sbox layers (S). At ASIACRYPT 2014, Biryukov et al. [2] proposed several concrete ASASA block cipher designs as candidates for white-box cryptography, whose security was based on the alleged difficulty of recovering their internal components. These designs were subsequently broken in [15] and [9].

Of particular interest is the integral attack of [9]. While the full details of this attack are out of the scope of this paper, we focus on its heaviest computational step that consists of summing over about 2^n affine subspaces of dimension slightly less than n (where n is the block size of the scheme). This step was performed in [9] in complexity of about $2^{3n/2}$. We can improve the complexity of this step (and the complexity of the full attack) to about 2^n by using a symbolic algorithm which is similar to the one used for computing the rank table.

Application to Distinguishers on Sboxes and Block Ciphers In [4] Biryukov and Perrin considered the problem of reverse-engineering Sboxes and proposed techniques to check whether a given Sbox was selected at random or was designed according to some unknown criteria. These techniques are based on

the linear approximation table (LAT) and difference distribution table (DDT) of the Sbox. Here, we provide another method based on the distribution of entries in the rank histogram of the Sbox. More specifically, an Sbox would be considered suspicious if its rank histogram entry sizes differ significantly from their expected values according to the distribution derived in Section 7 (supported by the experimental results of Appendix B). An advantage of our proposal is that the LAT and DDT require about 2^{2n} time and memory to compute and store, whereas the rank histogram can be computed in time of about 2^n . Hence, our proposal can be used to analyze larger Sboxes. We can also use additional properties of HSM rank histogram multi-sets (such as the number of unique HSMs) as possible distinguishing techniques.

In a related application, the rank table (and additional structures defined in this paper) can be used to experimentally construct distinguishers on block ciphers with a small block size (e.g., 32 bits). This is done by selecting a few keys for the block cipher at random and detecting consistent deviations from random among the resultant permutations. In particular, if there is a linear combination of the output bits that is a low-degree function of some $(n - 1)$ -dimensional input subspace, then we can detect it in time complexity of about 2^n . Since there are 2^{n+1} possible $(n - 1)$ -dimensional affine subspaces and 2^n linear combinations of output bits, we search over a space of 2^{2n+1} possible distinguishers in about 2^n time. This can be viewed as an improvement over known experimental methods [18] that search a much smaller space containing about n^2 potential high-order differential distinguishers in similar complexity (these methods only consider the input and output bits, but not their linear combinations). Finally, the technique can also be used on block ciphers with larger block sizes by considering linear subspaces of the input domain and output range.

10 Conclusions and Open Problems

In this paper we described an improved algorithm for the affine equivalence problem, focusing on randomly chosen permutations. The main open problem is to further improve the algorithm's complexity and applicability. An additional future work item is to find more applications for the rank table and related structures defined in this paper.

References

1. C. M. Adams. Constructing Symmetric Ciphers Using the CAST Design Procedure. *Des. Codes Cryptography*, 12(3):283–316, 1997.
2. A. Biryukov, C. Bouillaguet, and D. Khovratovich. Cryptographic Schemes Based on the ASASA Structure: Black-Box, White-Box, and Public-Key (Extended Abstract). In P. Sarkar and T. Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology*

- and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. *Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*, pages 63–84. Springer, 2014.
3. A. Biryukov, C. D. Cannière, A. Braeken, and B. Preneel. A Toolbox for Cryptanalysis: Linear and Affine Equivalence Algorithms. In E. Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 33–50. Springer, 2003.
 4. A. Biryukov and L. Perrin. On Reverse-Engineering S-Boxes with Hidden Design Criteria or Structure. In R. Gennaro and M. Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 116–140. Springer, 2015.
 5. C. Bouillaguet, P. Fouque, and A. Véber. Graph-Theoretic Algorithms for the "Isomorphism of Polynomials" Problem. In T. Johansson and P. Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 211–227. Springer, 2013.
 6. M. Brinkmann and G. Leander. On the classification of APN functions up to dimension five. *Des. Codes Cryptography*, 49(1-3):273–288, 2008.
 7. A. Canteaut and J. Roué. On the Behaviors of Affine Equivalent Sboxes Regarding Differential and Linear Attacks. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 45–74. Springer, 2015.
 8. J. Daemen. Limitations of the Even-Mansour Construction. In H. Imai, R. L. Rivest, and T. Matsumoto, editors, *Advances in Cryptology - ASIACRYPT '91, International Conference on the Theory and Applications of Cryptology, Fujiyoshida, Japan, November 11-14, 1991, Proceedings*, volume 739 of *Lecture Notes in Computer Science*, pages 495–498. Springer, 1991.
 9. I. Dinur, O. Dunkelman, T. Kranz, and G. Leander. Decomposing the ASASA Block Cipher Construction. *IACR Cryptology ePrint Archive*, 2015:507, 2015.
 10. S. Even and Y. Mansour. A Construction of a Cipher from a Single Pseudorandom Permutation. *J. Cryptology*, 10(3):151–162, 1997.
 11. A. Joux. *Algorithmic Cryptanalysis*. Chapman & Hall, 2009.
 12. V. F. Kolchin. *Random Graphs*. Cambridge Univ. Press, 1999.
 13. G. Leander and A. Poschmann. On the Classification of 4 Bit S-Boxes. In C. Carlet and B. Sunar, editors, *Arithmetic of Finite Fields, First International Workshop, WAIFI 2007, Madrid, Spain, June 21-22, 2007, Proceedings*, volume 4547 of *Lecture Notes in Computer Science*, pages 159–176. Springer, 2007.
 14. W. Michiels, P. Gorissen, and H. D. L. Hollmann. Cryptanalysis of a Generic Class of White-Box Implementations. In R. M. Avanzi, L. Keliher, and F. Sica, editors, *Selected Areas in Cryptography, 15th International Workshop, SAC 2008, Sackville, New Brunswick, Canada, August 14-15, Revised Selected Papers*, volume 5381 of *Lecture Notes in Computer Science*, pages 414–428. Springer, 2008.
 15. B. Minaud, P. Derbez, P. Fouque, and P. Karpman. Key-Recovery Attacks on ASASA. In T. Iwata and J. H. Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of*

Cryptography and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II, volume 9453 of *Lecture Notes in Computer Science*, pages 3–27. Springer, 2015.

16. J. Patarin. Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): Two New Families of Asymmetric Algorithms. In U. M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, volume 1070 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 1996.
17. J. Patarin, L. Goubin, and N. Courtois. Improved Algorithms for Isomorphisms of Polynomials. In K. Nyberg, editor, *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 184–200. Springer, 1998.
18. Q. Wang, Z. Liu, K. Varici, Y. Sasaki, V. Rijmen, and Y. Todo. Cryptanalysis of Reduced-Round SIMON32 and SIMON48. In W. Meier and D. Mukhopadhyay, editors, *Progress in Cryptology - INDOCRYPT 2014 - 15th International Conference on Cryptology in India, New Delhi, India, December 14-17, 2014, Proceedings*, volume 8885 of *Lecture Notes in Computer Science*, pages 143–160. Springer, 2014.

A Missing Proofs

Proof (of Lemma 1). We begin by proving the first part. For $i \in \{1, 2, \dots, m\}$, we show that $(P_{(\geq d)}^{(i)} \circ A_1)_{(\geq d)} = (P^{(i)} \circ A_1)_{(\geq d)}$ by examining these compositions. Denote the coefficient of the monomial M_u in $P^{(i)}$ by α_u , then

$$(P^{(i)} \circ A_1)_{(\geq d)} = \sum_{u \in \{0,1\}^n} \alpha_u \cdot (M_u \circ A_1)_{(\geq d)},$$

and

$$(P_{(\geq d)}^{(i)} \circ A_1)_{(\geq d)} = \sum_{u \in \{0,1\}^n \wedge wt(u) \geq d} \alpha_u \cdot (M_u \circ A_1)_{(\geq d)}.$$

To show that these two expressions are equal, it is sufficient to show that for each $u \in \{0,1\}^n$ such that $wt(u) < d$ (it is not included in the second sum), $\alpha_u \cdot (M_u \circ A_1)_{(\geq d)} = 0$ (it contributes nothing to the first sum). This holds because A_1 is affine and hence $deg(M_u \circ A_1) \leq deg(M_u) < d$. Thus, $(M_u \circ A_1)_{(\geq d)} = 0$ when $wt(u) < d$ which completes the proof of the first part.

To prove the second part of the lemma, let $A_2^{(i)}(x[1], \dots, x[n]) = \sum_{j=1}^n \beta_j x[j] + \alpha$ denote the symbolic form of the i 'th output bit of A_2 . Then, the i 'th polynomial of $(A_2 \circ (P_{(\geq d)}))_{(\geq d)}$ and $(A_2 \circ P)_{(\geq d)}$ are respectively computed by the expressions

$$\begin{aligned} (\alpha + \sum_{j=1}^n \beta_j P_{(\geq d)}^{(j)})_{(\geq d)} &= \sum_{j=1}^n \beta_j P_{(\geq d)}^{(j)}, \\ (\alpha + \sum_{j=1}^n \beta_j P^{(j)})_{(\geq d)} &= (\sum_{j=1}^n \beta_j P^{(j)})_{(\geq d)}. \end{aligned}$$

These expressions are equal as monomials of degree less than d in each $P^{(j)}$ (which are removed in the first expression) do not contribute to the final outcome in the second expression.

Finally, note that if A_2 is linear then $\alpha = 0$ and $A_2 \circ (\mathbf{P}_{(\geq d)}) = (A_2 \circ \mathbf{P})_{(\geq d)}$. ■

Proof (of Lemma 4). Let v be a vector in the affine range of $A \circ C_{|h_1+h_2,0}$. We need to show that it is in the affine range of $C_{|h'_1+h'_2,c}$ for $c = c_1 + c_2$, or equivalently, it satisfies the equation $(h'_1 + h'_2)(v) = c$.

Since v is in the affine range of $A \circ C_{|h_1+h_2,0}$, there exists $x \in \{0, 1\}^{n-1}$ such that $A \circ C_{|h_1+h_2,0}(x) = v$, namely, $C_{|h_1+h_2,0}(x) = A^{-1}(v)$. Hence $A^{-1}(v)$ is in the affine range of $C_{|h_1+h_2,0}$ and satisfies the equation $(h_1 + h_2)(A^{-1}(v)) = 0$.

Assume that v is in the affine range of $A \circ C_{|h_1,0}$, then similarly, $A^{-1}(v)$ is in the affine range of $C_{|h_1,0}$ and satisfies the equation $h_1(A^{-1}(v)) = 0$. By adding the equations above, we obtain $h_2(A^{-1}(v)) = 0$, namely, $A^{-1}(v)$ is in the affine range of $C_{|h_2,0}$. This implies that v is in the affine ranges of both $A \circ C_{|h_2,0}$ and $A \circ C_{|h_1,0}$. Since $h_1 \mapsto_{(A)} h'_1$, $h_2 \mapsto_{(A)} h'_2$ with associated constants c_1, c_2 , then v is in the affine ranges of $C_{|h'_1,c_1}$ and $C_{|h'_2,c_2}$. Hence, v satisfies the equations $h'_1(v) = c_1$ and $h'_2(v) = c_2$, which gives $(h'_1 + h'_2)(v) = c_1 + c_2$ as required.

The proof for the case where v is not in the affine range of $A \circ C_{|h_1,0}$ is completely symmetric. By similar arguments we get that v is not in the affine ranges of $C_{|h'_1,c_1}$ and $C_{|h'_2,c_2}$. Hence, it satisfies $h'_1(v) = c_1 + 1$ and $h'_2(v) = c_2 + 1$, which gives the same equality $(h'_1 + h'_2)(v) = c_1 + c_2$. ■

Proof (of Lemma 5). Since $h \mapsto_{(A)} h'$ with associated constant c then A transforms the $(n-1)$ -dimensional affine range of $C_{|h,0}$ to the $(n-1)$ -dimensional affine range specified by $C_{|h',c}$. The condition that A maps between affine subspaces can be broken down into a condition that L maps between the corresponding linear subspaces and a condition that A maps a vector from one affine subspace to the other. More specifically, first, the matrix L transforms the (linear) range of $C_{|h,0}$ to the linear range of $C_{|h',c}$ (whose canonical representative is $C_{|h',0}$). Since L is invertible, it must also transform each remaining vector (in the affine range of $C_{|h,1}$) to a vector in the affine range of $C_{|h',1}$ (otherwise, L would map two different inputs to the same output).

In short, the first condition is that for every $c' \in \{0, 1\}$, L transforms the affine range of $C_{|h,c'}$ to the affine range of $C_{|h',c'}$. The second condition is that A transforms the first unit vector e_1 (or any other vector) into the affine range of $C_{|h',c}$ if and only if e_1 is in the affine range of $C_{|h,0}$.

We prove that the first constraint above holds by analyzing the first condition. Let e_i denote the i 'th unit vector. Then, for some $c' \in \{0, 1\}$, e_i is in the affine range of $C_{|h,c'}$ implying that $h(e_i) = c'$, namely $h[i] = c'$. Moreover, L transforms e_i into the affine range of $C_{|h',c'}$, giving $h'(L(e_i)) = c'$, or $h'(L[i]) = c' = h[i]$ as claimed.

We now analyze the second condition. For some $c' \in \{0, 1\}$, e_1 is in the affine range of $C_{|h,c'}$. Then $h(e_1) = c'$ or $h[1] = c'$. Moreover, it is transformed by A into the affine range of $C_{|h',c+c'}$, giving $h'(L[1] + a) = c + c'$, namely

$h'(a) = h'(L[1]) + c + c'$. From the first constraint we have $h'(L[1]) = h[1]$, and as $h[1] = c'$ we obtain $h'(a) = c + c' + c' = c$ as required. ■

B Experimental Results on Rank Histogram Entries for Random Permutations

In Section 7 we concluded based on Assumption 1 that for a random permutation \mathbf{F} , each rank histogram entry $\mathcal{H}_{\mathbf{F},n-2}(maxR, minR)$ is distributed binomially with parameters 2^n and $p \approx p_{maxR, minR}$ (given in (3)). In this section we experimentally estimate the expectation and standard deviation of $\mathcal{H}_{\mathbf{F},n-2}(maxR, minR)$ for various choices of n and demonstrate that they indeed closely match those of a binomial random variable with parameters 2^n and $p \approx p_{maxR, minR}$.

More specifically, we give experimental evidence that the expectation of $\mathcal{H}_{\mathbf{F},n-2}(maxR, minR)$ is about $2^n \cdot p_{maxR, minR}$, while its standard deviation is close to $\sqrt{2^n \cdot p_{maxR, minR} \cdot (1 - p_{maxR, minR})}$.

The tables below present our results for $n \in \{8, 12, 16, 20, 24, 28\}$, where 1000 trials were performed for each value of $n < 28$ and 100 trials were performed for $n = 28$. In each trial, we selected a permutation $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ uniformly at random and calculated its rank histogram. We then calculated the expected value and standard deviation for each histogram entry $(maxR, minR)$ and compared them to the analytic estimates above. Results are rounded down to nearest integer and given only for histogram entries whose expected values are at least 1.

In terms of expectation, the experimentally calculated values of the histogram entries are very close (generally within 1%) to the analytical calculation. In terms of standard deviation, the experimentally calculated histogram entries seem to deviate slightly more compared to binomial variables, as standard deviations are generally within 30% (and typically less) of the analytical calculation. This small increase in standard deviation has negligible effect on the algorithm's performance (both asymptotically and concretely).

Finally, we note that we also performed (somewhat less extensive) experiments on random functions and obtained similar results.

Table 2. Rank Histogram Entry Statistics for $n = 8$ (Obtained for 1000 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(7,5)	1 (1)	1 (1)
(6,6)	4 (4)	2 (2)
(8,6)	18 (18)	4 (4)
(8,8)	21 (21)	5 (4)
(7,6)	37 (37)	7 (5)
(7,7)	85 (85)	7 (7)
(8,7)	85 (85)	9 (7)

Table 3. Rank Histogram Entry Statistics for $n = 12$ (Obtained for 1000 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(10,9)	5 (5)	2 (2)
(12,9)	12 (12)	3 (3)
(11,9)	24 (24)	5 (4)
(10,10)	67 (67)	8 (8)
(12,10)	302 (303)	16 (16)
(12,12)	342 (341)	20 (17)
(11,10)	607 (607)	30 (22)
(11,11)	1365 (1366)	30 (30)
(12,11)	1366 (1366)	36 (30)

Table 4. Rank Histogram Entry Statistics for $n = 16$ (Obtained for 1000 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(13,13)	1 (1)	1 (1)
(16,12)	1 (1)	1 (1)
(15,12)	3 (3)	1 (1)
(14,13)	88 (88)	10 (9)
(16,13)	198 (198)	14 (14)
(15,13)	396 (396)	21 (19)
(14,14)	1078 (1079)	36 (32)
(16,14)	4864 (4858)	68 (67)
(16,16)	5468 (5466)	81 (70)
(15,14)	9717 (9717)	115 (90)
(15,15)	21857 (21864)	122 (120)
(16,15)	21858 (21864)	145 (120)

Table 5. Rank Histogram Entry Statistics for $n = 20$ (Obtained for 1000 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(18,16)	12 (12)	3 (3)
(17,17)	28 (28)	5 (5)
(20,16)	28 (28)	5 (5)
(19,16)	56 (56)	7 (7)
(18,17)	1409 (1410)	38 (37)
(20,17)	3175 (3173)	58 (56)
(19,17)	6345 (6346)	85 (79)
(18,18)	17275 (17275)	148 (130)
(20,18)	77732 (77739)	275 (268)
(20,20)	87431 (87456)	316 (283)
(19,18)	155460 (155479)	471 (363)
(19,19)	349798 (349827)	484 (482)
(20,19)	349818 (349827)	586 (482)

Table 6. Rank Histogram Entry Statistics for $n = 24$ (Obtained for 1000 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(23,19)	1 (1)	1 (1)
(21,20)	8 (8)	2 (2)
(22,20)	201 (200)	13 (14)
(24,20)	450 (451)	21 (21)
(21,21)	460 (460)	22 (21)
(23,20)	901 (902)	28 (30)
(22,21)	22566 (22563)	155 (150)
(24,21)	50760 (50768)	236 (224)
(23,21)	101542 (101537)	343 (317)
(22,22)	276408 (276407)	575 (521)
(24,22)	1243735 (1243832)	1147 (1073)
(24,24)	1399243 (1399311)	1277 (1132)
(23,22)	2487495 (2487664)	1897 (1455)
(23,23)	5596692 (5597244)	1932 (1931)
(24,23)	5596745 (5597244)	2492 (1931)

C Asymptotic Analysis of the Unique HSM Algorithm

In this appendix we revisit the unique HSM algorithm of Section 8.2. We lower bound the expected number of unique HSMs in $\mathcal{HM}_{\mathbf{F},n-2,\mathbf{r},\mathbf{r}'}$ asymptotically (ignoring constant factors) based on Assumption 1, given that \mathbf{F} is a random permutation. Denote $S = \mathcal{H}_{\mathbf{F},n-2}(\mathbf{r})$, and $S' = \mathcal{H}_{\mathbf{F},n-2}(\mathbf{r}')$.

Table 7. Rank Histogram Entry Statistics for $n = 28$ (Obtained for 100 Trials)

Histogram Entry	Computed Expected Value (Analytic Expected Value)	Computed Standard Deviation (Analytic Standard Deviation)
(26,23)	6 (6)	2 (2)
(28,23)	14 (15)	4 (3)
(27,23)	30 (30)	5 (5)
(25,24)	133 (130)	12 (11)
(26,24)	3205 (3209)	60 (56)
(28,24)	7209 (7220)	91 (84)
(25,25)	7368 (7367)	83 (85)
(27,24)	14430 (14440)	118 (120)
(26,25)	361131 (361021)	722 (600)
(28,25)	812111 (812298)	892 (899)
(27,25)	1624526 (1624596)	1338 (1270)
(26,26)	4421944 (4422514)	2018 (2085)
(28,26)	19899724 (19901313)	4403 (4292)
(28,28)	22387702 (22388977)	5826 (4530)
(27,26)	39798878 (39802626)	6892 (5822)
(27,27)	89548460 (89555909)	9332 (7725)
(28,27)	89548573 (89555909)	9209 (7725)

Let p_{col} denote the HSM histogram collision probability $p_{col} = Pr[\mathcal{HG}_{\mathbf{F},n-2,h_1,\mathbf{r}'} = \mathcal{HG}_{\mathbf{F},n-2,h_2,\mathbf{r}'}]$ for $h_1, h_2 \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r})$. By the union bound, the probability that $h \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r})$ is unique in $\mathcal{HM}_{\mathbf{F},n-2,\mathbf{r},\mathbf{r}'}$ is at least $1 - p_{col} \cdot S$ and the expected number of unique HSMs is lower bounded by

$$S - p_{col} \cdot S^2. \quad (5)$$

We now upper bound p_{col} . Essentially, we have two HSM rank histograms, each with a sum of entries equal to S' , and our goal is to upper bound the probability that all of their entries are equal. According to Assumption 1, the HSMs in $\mathcal{T}_{\mathbf{F},n-2}(\mathbf{r})$ and $\mathcal{T}_{\mathbf{F},n-2}(\mathbf{r}')$ are uniformly distributed. Similarly to the analysis of Section 7, this implies that for each $h \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r})$, the entries of $\mathcal{HG}_{\mathbf{F},n-2,h,\mathbf{r}'}$ are distributed multinomially with parameter S' and probabilities given by $p_{maxR,minR}$ (specified in (3)).¹⁰

Since each possible outcome of a multinomial distribution is distributed binomially, to upper bound p_{col} we can focus on the single entry $(maxR, minR) = (n, n)$ and upper bound the collision probability of two independent binomial random variables, each with parameters S' and $p_{n,n}$. By (2), $p_{n,n} \approx \alpha^2 \approx 0.083$ is constant.

¹⁰ Note that $\mathcal{HG}_{\mathbf{F},n-2,h,\mathbf{r}'}$ for $h \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r})$ are essentially independent assuming $S \ll 2^n$ and $S' \ll 2^n$, since the fraction of collisions in the multi-set $\{h + h' \mid h \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r}) \wedge h' \in \mathcal{T}_{\mathbf{F},n-2}(\mathbf{r}')\}$ is negligible.

Therefore we have two identically distributed binary random variables $X \sim B(k, p)$, $Y \sim B(k, p)$ with parameters $k = S'$ and $p = p_{n,n}$ and we would like to approximate the probability that they are equal, namely $Pr[X - Y = 0]$. For this purpose, we use the normal approximation of a binary variable (given that S' is sufficiently large). A binomial random variable $B(k, p)$ can be approximated by $\mathcal{N}(kp, kp(1 - p))$, hence $X \approx \mathcal{N}(kp, kp(1 - p))$ and $Y \approx \mathcal{N}(kp, kp(1 - p))$. By properties of the normal distribution $X - Y \approx \mathcal{N}(0, 2kp(1 - p)) = \mathcal{N}(0, \sigma^2)$ (for $\sigma^2 = 2kp(1 - p)$). We obtain

$$p_{col} \leq Pr[X - Y = 0] = Pr[X - Y \leq 0] - Pr[X - Y \leq -1] \approx \int_{x=-1}^0 \frac{1}{\sqrt{2\pi\sigma^2}} e^{-x^2/2\sigma^2} dx \leq \frac{1}{\sqrt{2\pi\sigma^2}} \approx 1/\sigma = \frac{1}{\sqrt{2kp(1-p)}} = \frac{1}{\sqrt{2S' \cdot p_{n,n}(1-p_{n,n})}} \approx \frac{1}{\sqrt{S'}}$$

where the second inequality is obtained by upper bounding the probability density function at $x = 0$.

Finally, plugging p_{col} into (5), we obtain the asymptotic lower bound of

$$S - S^2/\sqrt{S'}. \tag{6}$$

Note that this gives a positive value only for $S' > S^2$.