FINITE FIELDS AND THEIR APPLICATIONS 1, 31–56 (1995)

Tests for Permutation Functions

KEJU MA AND JOACHIM VON ZUR GATHEN*

Department of Computer Science, University of Toronto, Toronto, Ontario M5S 1A4, Canada
E-mail: keju,gathen@cs.toronto.edu

Communicated by Rudolf Lidl

Received October 12, 1993

Let $\mathbb{F}_q$ be a finite field with q elements, $f \in \mathbb{F}_q(x)$ a rational function over $\mathbb{F}_q$, and $\mathbb{D} \subseteq \mathbb{F}_q$ the domain of definition of f. Consider three notions of "permutation functions": f is a permutation on $\mathbb{F}_q$, or on $\mathbb{D}$, or f is injective on $\mathbb{D}$. For each of these, a random polynomial-time test is presented. For the image size of an arbitrary rational function, a fully polynomial-time randomized approximation scheme is given.

1. INTRODUCTION

Let q be a power of a prime, $\mathbb{F}_q$ a finite field with q elements, $f = g/h \in \mathbb{F}_q(x)$ a rational function with $g, h \in \mathbb{F}_q[x]$ and $\gcd(g, h) = 1$. Then f induces a partial mapping $\mathbb{F}_q \to \mathbb{F}_q$ via $a \mapsto f(a)$ for all $a \in \mathbb{F}_q$ with $h(a) \neq 0$. If f is total (i.e., the denominator h has no roots in $\mathbb{F}_q$) and the mapping is bijective, then f is called a *permutation function* over $\mathbb{F}_q$. In the special case $h = 1$, so that $f = g \in \mathbb{F}_q[x]$, it is called a *permutation polynomial* over $\mathbb{F}_q$.

Permutation functions over finite fields have been studied since the past century. Besides monomials, the two best-known classes of permutation functions are *Dickson polynomials* and *Rédei functions*, introduced by Dickson [5] and Rédei [26], respectively. Lidl and Wells [22], Fried and Lidl [6], and Matthews and Lidl [24] considered generalized Dickson polynomials and Rédei functions. In recent years, considerable attention

* Current address: Fachbereich Mathematik-Informatik, Universität-GH Paderborn, Germany. E-mail: gathen@uni-paderborn.de.

1071-5797/95 $6.00

has been given to potential applications of permutation functions in public-key cryptography. We refer the reader to Lidl and Niederreiter [21, Chap. 7], Lidl and Mullen [19, 20], and Mullen [25] for surveys of recent work and related literature.

The interest in permutation functions comes from the fact that they form a connection between two central aspects of the theory of finite fields: combinatorics and algebra. The definition of these objects is rather combinatorial, and applications such as in cryptography can also be considered combinatorial. A computational problem is how to represent permutations efficiently; a general representation requires exponential size (in $\log q$), and it is not clear how to represent interesting classes much more concisely. Algebra presents a partial solution: take those functions that can be represented concisely as polynomials or rational functions. This introduces the degree as an interesting measure. Two fundamental questions for manipulating these objects are to test when an arbitrary rational function is a permutation and when the composition of two is again a permutation. These questions are addressed in this paper.

Some efficient algorithms are known for testing whether a given polynomial is a permutation. If the input $f \in \mathbb{F}_q[x]$ has degree $n < q$, the probabilistic test of von zur Gathen [9] uses $O^{\sim}(n \log q)$ operations in $\mathbb{F}_q$, i.e., softly linear time in the input size $n \log q$. We use the "soft O" notation to ignore logarithmic factors:

$$s = O^{\sim}(t) \Leftrightarrow s = O(t \log^k t) \text{ for some constant } k.$$

A variant of that test applies to "almost permutation polynomials," whose value set contains at least $q - \rho$ elements of $\mathbb{F}_q$, and uses $O^{\sim}(n\rho \log q)$ operations in $\mathbb{F}_q$. Subsequently, Shparlinski [29] obtained a deterministic test for permutation polynomials using $O^{\sim}(n^3 q^{1/2})$ operations in $\mathbb{F}_q$ (if $q > 4n^6$).

In this paper we present efficient probabilistic tests for permutation functions of three flavors. They are all based on the subresultant approach introduced in von zur Gathen [9]. We use the following conventions throughout the paper:

$$\begin{aligned}
&f = g/h \in \mathbb{F}_q(x) \text{ with } g, h \in \mathbb{F}_q[x] \text{ and } \gcd(g, h) = 1,\\
&n = \deg f = \max\{\deg g, \deg h\},\\
&\mathbb{D} = \{a \in \mathbb{F}_q : h(a) \neq 0\} \text{ is the domain of definition of } f,\\
&\sigma = q - \#\mathbb{D} \text{ is the } \textit{singularity} \text{ of } f,\\
&\mathbb{V} = \{f(a) \in \mathbb{F}_q : a \in \mathbb{D}\} \text{ is the value set (or image) of } f,\\
&\nu = \#\mathbb{V} \text{ is the image size of } f.
\end{aligned} \tag{1.1}$$

Section 2 presents a fast probabilistic test for permutation functions in the usual sense [4], namely, functions that induce a bijection from $\mathbb{F}_q$ into itself. The running time of the test is $O^{\sim}(n \log q)$ operations in $\mathbb{F}_q$.

We say that f is a *bijective function* if and only if f is injective on $\mathbb{D}$ (and hence induces a bijection from $\mathbb{D}$ into $\mathbb{V}$). Thus permutation functions are bijective functions of singularity zero. In Section 3, we show that the class of these bijective functions represents exactly all bijections between subsets of $\mathbb{F}_q$ and study how this new class of functions behaves under the formal composition of rational functions. It turns out that the 12 classes definable by this operation actually form a hierarchy of exactly 5 classes, ordered by inclusion (Theorem 3.10). A probabilistic test for bijective functions presented in Section 4 uses $O^{\sim}(n\sigma^2 \log q)$ operations in $\mathbb{F}_q$. We also note that both a bijective function of small degree and its inverse are easy to evaluate. This property is required for applications such as in cryptography.

The class of all one-to-one partial mappings on a finite set has been extensively studied in the context of algebraic semigroup theory, and it forms an inverse semigroup under the set-theoretic composition of partial mappings [3, 13]. Analogous to Cayley's famous theorem saying that any finite group can be represented by permutations on a finite set, the Wagner–Preston Representation Theorem states that any finite inverse semigroup can be represented by one-to-one partial mappings on a finite set. Our results on bijective functions over finite fields may therefore be viewed as a further development in this direction, in the sense that such partial mappings can all be represented by rational functions.

In Section 5 we examine a special subclass of bijective functions. We say that f is a *partial permutation function* if and only if f induces a bijection from $\mathbb{D}$ into itself (and hence $\mathbb{D} = \mathbb{V}$). Thus the class of partial permutation functions represents exactly all permutations on subsets of $\mathbb{F}_q$.

We show that every bijection between subsets of $\mathbb{F}_q$ can be represented by a bijective function that is the composition of a permutation polynomial and a partial permutation function in either order. In Theorem 5.5, we describe the hierarchy of composition classes involving partial permutation functions. Our results are not as complete as Theorem 3.10. It comes as a surprise that some of these classes coincide if and only if $q \leq 5$. A well-known phenomenon is that some general properties of permutation polynomials fail when the field is small compared to the degree; in our case, however, the properties apply to arbitrarily large degrees. In Section 6, we design a probabilistic test for partial permutation functions using only $O^{\sim}(n \log q)$ operations in $\mathbb{F}_q$.

A more general problem than the recognition of permutation functions is to count value sets of arbitrary rational functions over $\mathbb{F}_q$. If q is large, however, the naive counting method for the value set $\mathbb{V}$ of $f \in \mathbb{F}_q(x)$, which evaluates f at all $a \in \mathbb{D}$ and checks how many of these values are distinct, quickly becomes prohibitively expensive even when $\deg f$ is small. Like many other enumeration problems, it seems hard to design

efficient deterministic counting algorithms for $\mathbb{V}$. We are not aware of any deterministic algorithm that computes the size of $\mathbb{V}$ exactly and runs in time sublinear in q.

Section 7 develops a probabilistic approximation algorithm for the image size ν of $f \in \mathbb{F}_q(x)$ of degree $n < q$. We give a fully polynomial-time (ε, δ)-approximation scheme for ν, using $O^{\sim}(n^2\varepsilon^{-2} \log q \log \delta^{-1})$ operations in $\mathbb{F}_q$.

We mention some recent progress in developing (ε, δ)-approximation algorithms for algebraic counting problems. Karpinski and Luby [17] designed an (ε, δ)-approximation algorithm for counting the number of zeros of a multivariate polynomial $f(x_1, \ldots, x_n) \in \mathbb{F}_2[x_1, \ldots, x_n]$. The algorithm was later extended by Karpinski and Lhotzky [16] to an arbitrary multilinear polynomial over $\mathbb{F}_q$. Grigoryev and Karpinski [12] designed an (ε, δ)-approximation algorithm for counting the number of zeros of an arbitrary polynomial $f(x_1, \ldots, x_n)$ over $\mathbb{F}_q$. However, unlike the first algorithm, the latter two (ε, δ)-approximation algorithms for arbitrary $\mathbb{F}_q$ are not fully polynomial-time in the sense that their running times are superlinear in q rather than polynomial in $\log q$.

Viewing the image of a rational function as the projection of its graph, we see that the question of counting the size of a projection of a curve generalizes our problem. This question is addressed in von zur Gathen *et al.* [11].

2. A Test for Permutation Functions

For $f = g/h$, $\mathbb{D}$, $\mathbb{V}$ as in (1.1), we say that f is a permutation function if and only if $\mathbb{D} = \mathbb{V} = \mathbb{F}_q$ and denote by $\mathsf{PF} \subseteq \mathbb{F}_q(x)$ the class of all permutation functions over $\mathbb{F}_q$. Then

$$\begin{aligned} f \in \mathsf{PF} &\Leftrightarrow \forall a \in \mathbb{F}_q\ \exists b \in \mathbb{F}_q\ g(b)/h(b) = a \\ &\Leftrightarrow \forall a \in \mathbb{F}_q\ \exists b \in \mathbb{F}_q\ g(b) - ah(b) = 0 \\ &\Leftrightarrow \forall a \in \mathbb{F}_q\ \exists b \in \mathbb{F}_q\ x - b \mid g - ah \\ &\Leftrightarrow \forall a \in \mathbb{F}_q\ \gcd(x^q - x, g - ah) \neq 1 \\ &\Leftrightarrow \forall a \in \mathbb{F}_q\ \operatorname{res}_x(x^q - x, g - ah) = 0 \\ &\Leftrightarrow y^q - y \mid \operatorname{res}_x(x^q - x, g - yh). \end{aligned}$$

Here y is a new indeterminate, $r = \operatorname{res}_x(x^q - x, g - yh) \in \mathbb{F}_q[y]$ is the resultant of the two polynomials $x^q - x$ and $g - yh$ in $\mathbb{F}_q(y)[x]$, and the divisibility condition is in $\mathbb{F}_q[y]$. From the definition of r as the determinant of the $(q + n) \times (q + n)$-Sylvester matrix for $x^q - x$ and $g - yh$ in $\mathbb{F}_q(y)[x]$, it follows that $\deg r \leq q$. Furthermore, the condition $\gcd(g, h) = 1$ implies $r \neq 0$, since

$$\begin{aligned}\operatorname{res}_x(x^q - x,\ g - yh) = 0 &\Rightarrow \exists b \in \mathbb{F}_q\ g(b) - yh(b) = 0\\ &\Rightarrow \exists b \in \mathbb{F}_q\ g(b) = 0 \text{ and } h(b) = 0\\ &\Rightarrow \exists b \in \mathbb{F}_q\ x - b \mid g \text{ and } x - b \mid h\\ &\Rightarrow \gcd(g, h) \neq 1.\end{aligned}$$

THEOREM 2.1. *For $f = g/h$ as in (1.1), $f \in \mathsf{PF}$ if and only if*

$$\operatorname{res}_x(x^q - x,\ g - yh) = c(y^q - y), \text{ for some } c \in \mathbb{F}_q^\times = \mathbb{F}_q\backslash\{0\}.$$

In the special case $h = 1$, so that $f = g \in \mathbb{F}_q[x]$, the constant can be predetermined as $c = -1$ [9].

It would be very costly to compute $r = \operatorname{res}_x(x^q - x, g - yh)$ directly via the determinant of the $(q + n) \times (q + n)$-Sylvester matrix. We can, however, substitute a randomly chosen $a \in \mathbb{F}_{q^m}\backslash\mathbb{F}_q$ for y, from a proper field extension $\mathbb{F}_{q^m} \supset \mathbb{F}_q$, use the Euclidean algorithm to calculate $r(a) \in \mathbb{F}_{q^m}^\times$, and then check whether $r(a)/(a^q - a)$ is a nonzero constant in $\mathbb{F}_q$; with high probability, this will happen if and only if f is a permutation function.

We denote by s rem t the remainder of s on division by t ($t \neq 0$), for s, t in a Euclidean domain such as $F[x]$ (where F is a field) or $\mathbb{Z}$, and recall the definition of the *Euclidean representation* of (s, t) from Knuth [18], Strassen [30], and von zur Gathen [9]. The following fact is derived from von zur Gathen [9, Sect. 5].

Fact 2.2. Let $\mathbb{F}_{q^m} \supset \mathbb{F}_q$ with $m \geq 2$ be a field extension and $a \in \mathbb{F}_{q^m}$. For $f = g/h$ as in (1.1) and $u = x^q - x \operatorname{rem}(g - ah) \in \mathbb{F}_{q^m}[x]$, let $(q_2, \ldots, q_l, a_l)$ be the Euclidean representation of $(g - ah, u)$, $d_i = \deg q_i$, and $\gamma_i \in \mathbb{F}_{q^m}$ be the leading coefficient of q_i, for $2 \leq i \leq l$. Let $\alpha \in \mathbb{F}_{q^m}$ be the leading coefficient of $g - ah$, $n_0 = q$, $n_1 = \deg(g - ah) \leq n$, $n_i = n_1 - \sum_{2 \leq j \leq i} d_j$ for $2 \leq i \leq l$, and $s = \sum_{0 \leq i < l} n_i n_{i+1} \operatorname{rem} 2$. Then, for $r = \operatorname{res}_x(x^q - x, g - yh) \in \mathbb{F}_q[y]\backslash\{0\}$, we have

$$r(a) = \begin{cases} 0 & \text{if } n_l \geq 1,\\ (-1)^s\, \alpha^{n_0 + n_1} \displaystyle\prod_{2 \leq i \leq l} \gamma_i^{-(n_{i-1} + n_i)} & \text{if } n_l = 0. \end{cases} \tag{2.1}$$

Furthermore, if $a \in \mathbb{F}_{q^m}\backslash\mathbb{F}_q$, then $r(a) \in \mathbb{F}_{q^m}^\times$.

Probabilistic Test for Permutation Functions

Input: $f = g/h$ as in (1.1), a monic quadratic irreducible polynomial $\varphi \in \mathbb{F}_q[x]$, and $\varepsilon > 0$.

Output: Yes or No.

1. Set $a = (x \bmod \varphi) \in \mathbb{F}_{q^2} \backslash \mathbb{F}_q$ and calculate
   $r(a) \in \mathbb{F}_{q^2}^{\times}$ by Fact 2.2,
   $c = r(a)/(a^q - a) \in \mathbb{F}_{q^2}^{\times}$ by repeated squaring.
   If $c \notin \mathbb{F}_q^{\times}$, then return No and stop.
2. Repeat $t = \lfloor \log_q \varepsilon^{-1} \rfloor$ times:
   Randomly and uniformly choose $a \in \mathbb{F}_{q^2} = \mathbb{F}_q[x]/(\varphi)$ and calculate $\beta = r(a) - c(a^q - a)$.
   If $\beta \neq 0$, then return No and stop.
3. Return Yes.

To estimate the cost, let $\mathsf{M}: \mathbb{N} \to \mathbb{R}$ be a "universal" cost of multiplication, so that the multiplication st takes $O(\mathsf{M}(n))$ arithmetic operations or bit operations, respectively, for two polynomials $s, t \in F[x]$ of degree at most n, or two n-bit integers $s, t \in \mathbb{Z}$. Similarly, the division with remainder $s \operatorname{rem} t$ (if $t \neq 0$) and the gcd computation $\gcd(s, t)$ can be performed in $O(\mathsf{M}(n))$ and $O(\mathsf{M}(n) \log n)$ operations, respectively [1, Sect. 7.5]. We can choose $\mathsf{M}(n) = n \log n \log\log n$ with "fast arithmetic" [2, 27] and $\mathsf{M}(n) = n^2$ with "classical arithmetic."

THEOREM 2.3. *Let $q \geq n$, $\varepsilon > 0$, and $t = \lceil \log_q \varepsilon^{-1} \rceil$. The test for permutation functions uses $O(2t)$ random choices and $O(t\, \mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$. If $f \in \mathsf{PF}$, the output is* Yes; *if $f \notin \mathsf{PF}$, the output is* No *with probability at least $1 - \varepsilon$. In particular, for $\varepsilon = (1/q)^{O(1)}$, the test can be performed with $O(1)$ random choices and $O^{\sim}(n \log q)$ operations in $\mathbb{F}_q$.*

Proof. If $f \in \mathsf{PF}$, then the test returns Yes, since $r - c(y^q - y) = 0$ by Theorem 2.1. If $f \notin \mathsf{PF}$ and $c \notin \mathbb{F}_q^{\times}$, the test returns No correctly in Step 1.

If $f \notin \mathsf{PF}$ and $c \in \mathbb{F}_q^{\times}$, then by Theorem 2.1, $s = r - c(y^q - y) \in \mathbb{F}_q[y]$ is a nonzero polynomial of degree $\leq q$, and any $a \in \mathbb{F}_{q^2}$ with $\beta = s(a) \neq 0$ is a "witness to the nonbijectivity" of f. The probability of returning a wrong Yes in Step 3 thus equals that of selecting t "liars" $a \in \mathbb{F}_{q^2}$ with $\beta = 0$ in a row, which is at most $(q/q^2)^t = (1/q)^t \leq \varepsilon$.

For the timing analysis, we first note that any field operation in $\mathbb{F}_{q^2}$ can be simulated with $O(1)$ operations in $\mathbb{F}_q$. Therefore, in Step 1, computing $u = x^q - x \operatorname{rem}(g - ah) \in \mathbb{F}_{q^2}[x]$ takes $O(\mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$ via repeated squaring, and calculating $r(a) \in \mathbb{F}_{q^2}^{\times}$ takes $O(\mathsf{M}(n) \log n)$ operations in $\mathbb{F}_q$ via the Euclidean representation for $(g - ah, u)$. The total cost of this step is $O(\mathsf{M}(n) \log(qn))$ operations in $\mathbb{F}_q$, or $O(\mathsf{M}(n) \log q)$ operations if $q \geq n$.

Step 2 uses $O(2t)$ random choices in $\mathbb{F}_q$ to generate t random elements in $\mathbb{F}_{q^2}$ and takes $O(t\, \mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$, or $O^{\sim}(n \log \varepsilon^{-1})$ operations if $\varepsilon \leq q^{-1}$. ∎

Remark 2.4. It is easy to find a quadratic irreducible polynomial $\varphi \in \mathbb{F}_q[x]$ for our test. We distinguish two cases: char $\mathbb{F}_q \neq 2$ and char $\mathbb{F}_q = 2$.

If char $\mathbb{F}_q \neq 2$, then $\varphi = x^2 - \alpha \in \mathbb{F}_q[x]$ is irreducible over $\mathbb{F}_q$ if and only if α is a nonsquare. Therefore, with an expected number of two random choices and $O(\log q)$ operations in $\mathbb{F}_q$, we can find such an α with $\alpha^{(q-1)/2} \neq 1$ and choose $\varphi = x^2 - \alpha$ as the desired irreducible polynomial.

If char $\mathbb{F}_q = 2$, then $q = 2^l$ for some $l \in \mathbb{N}$, and we let $T: \mathbb{F}_q \to \mathbb{F}_2$ be the absolute trace function. It is well known that $x^2 + x + \alpha \in \mathbb{F}_q[x]$ is irreducible over $\mathbb{F}_q$ if and only if $T(\alpha) \neq 0$ [21, Corollary 3.79].

Now suppose that $\omega = x^l + a_1 x^{l-1} + \cdots + a_{l-1}x + a_l \in \mathbb{F}_2[x]$ is irreducible over $\mathbb{F}_2$, and $\mathbb{F}_q = \mathbb{F}_2[x]/(\omega)$ is represented by the basis $\{1, \xi, \ldots, \xi^{l-1}\}$ with $\xi = (x \bmod \omega) \in \mathbb{F}_q$. Clearly, $\beta_i = T(\xi^i) \neq 0$ for some $0 \leq i < l$, $\beta_0 = l \bmod 2$, and $\beta_1 = a_1$, since $\omega = (x - \xi)(x - \xi^2) \cdots (x - \xi^{2^{l-1}})$.

Furthermore, for $2 \leq i < l$, Newton's formula yields that

$$\begin{aligned}\beta_i &= \xi^i + (\xi^i)^2 + \cdots + (\xi^i)^{2^{l-1}} \\ &= (\xi)^i + (\xi^2)^i + \cdots + (\xi^{2^{l-1}})^i \\ &= a_1\beta_{i-1} + a_2\beta_{i-2} + \cdots + a_{i-1}\beta_1 + a_i \cdot (i \bmod 2).\end{aligned}$$

Therefore, if $j = \min\{0 \leq i < l : \beta_0 \neq 0 \text{ or } a_i \cdot (i \bmod 2) \neq 0\}$, then $\beta_j \neq 0$. With virtually no computation, we can find this index j and choose $\varphi = x^2 + x + \xi^j$ as the desired irreducible polynomial.

Remark 2.5. The test for permutation functions can be implemented using any field extension $\mathbb{F}_{q^m} \supset \mathbb{F}_q$ with $m \geq 2$. The new implementation relies on the following fact: if $f \notin \mathsf{PF}$, then we can determine a $c \in \mathbb{F}_{q^m}^{\times}$ by $c = r(a)/(a^q - a)$ for some arbitrary $a \in \mathbb{F}_{q^m} \backslash \mathbb{F}_q$, so that either $c \notin \mathbb{F}_q^{\times}$ is a witness to the nonbijectivity of f or $c \in \mathbb{F}_q^{\times}$. In the latter case, $s = r - c(y^q - y) \in \mathbb{F}_q[y]$ is a nonzero polynomial of degree $\leq q$, and any element in the set $\{a \in \mathbb{F}_{q^m} : s(a) \neq 0\}$ constitutes a witness. Therefore, the probability that a random element in $\mathbb{F}_{q^m}$ is a witness is at least $(q^m - q)/q^m = 1 - q^{1-m} \geq 1/2$. In particular, for any $0 < \varepsilon \leq q^{-1}$ and $m = 1 + \lceil \log_q \varepsilon^{-1} \rceil$, the probability is at least $1 - \varepsilon$.

To construct $\mathbb{F}_{q^m}$, we need an irreducible polynomial $\varphi \in \mathbb{F}_q[x]$ of degree m, which can be found with an expected number of $O^{\sim}(m^2 + m \log q)$ operations in $\mathbb{F}_q$ via the probabilistic algorithm of Shoup [28]. Furthermore, any field operation in $\mathbb{F}_{q^m}$ can be implemented with $O(\mathsf{M}(m) \log m)$ operations in $\mathbb{F}_q$.

For $0 < \varepsilon \leq q^{-1}$ and $2 \leq m \leq 1 + \lceil \log_q \varepsilon^{-1} \rceil$, apart from the cost of constructing $\mathbb{F}_{q^m}$, the test uses $O^{\sim}(n \log \varepsilon^{-1})$ operations in $\mathbb{F}_q$. In that case, a returned No answer is always correct, and a Yes answer is correct with probability at least $1 - \varepsilon$.

Remark 2.6. Other tests for permutation functions are available: a deterministic one extends the test for permutation polynomials in Shparlinski [29], and a probabilistic one generalizes the "naive" test for permutation polynomials in von zur Gathen [10]. A full description of these two tests (plus others) can be found in Ma and von zur Gathen [23].

3. Construction and Composition of Bijective Functions

For $f = g/h$, σ, $\mathbb{D}$, $\mathbb{V}$ as in (1.1), we say that f is a *bijective function* if and only if f is injective on $\mathbb{D}$. Thus f *represents* the bijection $a \mapsto f(a)$ between $\mathbb{D}$ and $\mathbb{V}$. We call σ its *singularity* and denote by $\mathsf{BF} \supseteq \mathsf{PF}$ the class of all bijective functions over $\mathbb{F}_q$.

Clearly, if $f \in \mathsf{PF}$, then $1/f \in \mathsf{BF}$ has singularity 1. In general, we have the following.

Proposition 3.1. *If $f \in \mathsf{BF}$ has singularity k, then $1/f \in \mathsf{BF}$ if and only if $k \leq 1$, and if this is the case, then its singularity is $1 - k$.*

A substantial generalization of this fact is given in Theorem 3.7.

Proposition 3.2. *Every bijection between two subsets of $\mathbb{F}_q$ can be represented by a bijective function of degree at most q.*

Proof. Let $\tau: \mathbb{A} \to \mathbb{B}$ be a bijection between two sets $\mathbb{A}, \mathbb{B} \subseteq \mathbb{F}_q$. Then τ can be represented by a bijective function $f = g/h$ as in (1.1), constructed as follows:

- Use interpolation, say, by the Lagrange formula [21, Theorem 1.71], to find $\tilde{g} \in \mathbb{F}_q[x]$ of degree $<q$, such that $\tilde{g}(a) = \tau(a)$ for all $a \in \mathbb{A}$ and $\tilde{g}(a) = 1$ for all $a \in \mathbb{F}_q \backslash \mathbb{A}$.
- If $\mathbb{A} = \emptyset$, then set $\tilde{h} = x^q - x$; otherwise, use interpolation to obtain a nonzero $\tilde{h} \in \mathbb{F}_q[x]$ of degree $<q$, such that $\tilde{h}(a) = 1$ for all $a \in \mathbb{A}$ and $\tilde{h}(a) = 0$ for all $a \in \mathbb{F}_q \backslash \mathbb{A}$.
- Set $w = \gcd(\tilde{g}, \tilde{h})$, $g = \tilde{g}/w$, and $h = \tilde{h}/w$.

Since $\gcd(x^q - x, w) = 1$, we have $\mathbb{D}(g/h) = \mathbb{D}(\tilde{g}/\tilde{h}) = \mathbb{A}$, and $g(a)/h(a) = \tilde{g}(a)/\tilde{h}(a) = \tau(a)$ for all $a \in \mathbb{A}$. ∎

While a polynomial $f \in \mathbb{F}_q[x]$ of degree less than q describes a total mapping $\mathbb{F}_q \to \mathbb{F}_q$ uniquely, a partial mapping $\mathbb{F}_q \to \mathbb{F}_q$ might be represented by different rational functions of degree at most q. Consequently, for $0 \leq \sigma \leq q$, the number of bijections between subsets of $\mathbb{F}_q$ of size $q - \sigma$ yields a lower bound on the number of bijective functions of singularity σ.

Proposition 3.3. *The class BF of bijective functions represents exactly all bijections between subsets of $\mathbb{F}_q$. Over an arbitrary finite field $\mathbb{F}_q$*

and for $0 \le \sigma \le q$, *the number of bijective functions of singularity* σ *and degree at most* q *is at least*

$$\binom{q}{\sigma}^2 (q - \sigma)!.$$

In the rest of this section, we discuss the composition of our various types of functions. It turns out that the 12 classes defined a priori by this operation collapse into 5 distinct classes, ordered by inclusion. Note that by "composition" we mean the formal composition of rational functions, by substituting one function for the indeterminate and then simplifying. The subtle differences of this "syntactic" operation and the "semantic" composition of the set-theoretic partial mappings play a major role in the sequel.

We first observe the following anomalies when composing bijective functions.

EXAMPLE 3.4. If $q > 2$ and $a \in \mathbb{F}_q \backslash \{0, 1\}$, then

$$l = \frac{x + 1}{(x + 1)^{q-1} - a} \in \mathsf{PF}, \qquad r = -\frac{1}{x} \in \mathsf{BF},$$

and r has singularity 1. The composition

$$l \circ r = \frac{x^{q-2}(x - 1)}{(x - 1)^{q-1} - ax^{q-1}}$$

is total but not bijective, since $(l \circ r)(0) = (l \circ r)(1) = 0$.

Over $\mathbb{F}_2$, $l = (x + 1)/(x^2 + x + 1) \in \mathsf{PF}$, $r = 1/x \in \mathsf{BF}$, and r has singularity 1. The composition $l \circ r = x(x + 1)/(x^2 + x + 1)$ is total but not bijective, since $(l \circ r)(0) = (l \circ r)(1) = 0$.

EXAMPLE 3.5. If $q > 2$ and $a \in \mathbb{F}_q \backslash \{0, 1\}$, then both $l = x/(x - a)^{q-1}$ and $r = x/(x - 1)$ are bijective and have singularity 1. The composition

$$l \circ r = \frac{x(x - 1)^{q-2}}{[(1 - a)x + a]^{q-1}}$$

is neither total nor bijective, since it has one singular point $a/(a - 1)$, and $(l \circ r)(0) = (l \circ r)(1) = 0$.

These examples demonstrate that the composition $l \circ r$ of $l, r \in \mathsf{BF}$ is not necessarily bijective. The explanation for such anomalies is that a singular point of r might become *desingularized* during the composition

and subsequently map to an image of l and hence violate the one-to-one nature of bijective functions. However, it will become clear in our proof of Theorem 3.7 that $l \circ r$ either is bijective or has only one "multiple" image to which at most $\sigma + 1$ points from its domain are mapped, where σ is the singularity of r.

LEMMA 3.6 (Zippel [31]). *Let F be an arbitrary field, and $u = u_m x^m + u_{m-1}x^{m-1} + \cdots + u_1 x + u_0$ and $v = v_n x^n + v_{n-1}x^{n-1} + \cdots + v_1 x + v_0 \in F[x]$ relatively prime with $u_m v_n \neq 0$. Let $U, V \in F[x, y]$ be the bivariate homogenizations of u, v of degrees m and n, respectively:*

$$U(x, y) = y^m u\left(\frac{x}{y}\right) = u_m x^m + u_{m-1}x^{m-1}y + \cdots + u_1 x y^{m-1} + u_0 y^m,$$

$$V(x, y) = y^n v\left(\frac{x}{y}\right) = v_n x^n + v_{n-1}x^{n-1}y + \cdots + v_1 x y^{n-1} + v_0 y^n.$$

Then for any relatively prime $g, h \in F[x]$, the polynomials $U(g, h)$, $V(g, h) \in F[x]$ are also relatively prime.

Proof. We give a more direct proof than Zippel's, using the existence of $s, t \in F[x]$ with $us + vt = 1$, $\deg s < n$ and $\deg t < m$. Substituting g/h for x yields

$$u(g/h)s(g/h) + v(g/h)t(g/h) = 1,$$

$$U(g, h)h^{n-1}s(g/h) + V(g, h)h^{m-1}t(g/h) = h^{m+n-1}.$$

Write $w = \gcd(U(g, h), V(g, h))$. The above shows that $w \mid h^{m+n-1}$. Since $\gcd(g, h) = 1$, we have $\gcd(U(g, h), h) = 1$. It follows that $\gcd(w, h) = 1$ and therefore $w = 1$. ∎

THEOREM 3.7. *Let $l = u/v$, $r = g/h \in \mathsf{BF}$ with $u, v, g, h \in \mathbb{F}_q[x]$, $\gcd(u, v) = \gcd(g, h) = 1$, $\deg u = m$, $\deg v = n$, and leading coefficients u_m and v_n, respectively. Let $\sigma = \deg \gcd(x^q - x, h)$, and*

$$\alpha = \begin{cases} u_m/v_m & \text{if } m = n, \\ 0 & \text{if } m \neq n. \end{cases}$$

Then $l \circ r \in \mathsf{BF}$ if and only if one of the following conditions holds:

(i) $m > n$;

(ii) $\sigma = 0$;

(iii) $\sigma = 1$ *and* $\alpha \notin \mathbb{V}(l)$;

(iv) $\sigma = 1$ *and* $\alpha = l(a)$ *for some* $a \in \mathbb{F}_q \backslash \mathbb{V}(r)$.

Proof. Using the notation from Lemma 3.6, we write

$$u = u_m x^m + u_{m-1}x^{m-1} + \cdots + u_1 x + u_0,$$

$$v = v_n x^n + v_{n-1}x^{n-1} + \cdots + v_1 x + v_0,$$

$$l \circ r = h^{n-m} \frac{u_m g^m + u_{m-1} g^{m-1} h + \cdots + u_1 g h^{m-1} + u_0 h^m}{v_n g^n + v_{n-1} g^{n-1} h + \cdots + v_1 g h^{n-1} + v_0 h^n}$$

$$= h^{n-m} \frac{U(g, h)}{V(g, h)}.$$

It follows from Lemma 3.6 that

$$\gcd(U(g, h), h) = \gcd(V(g, h), h) = \gcd(U(g, h), V(g, h)) = 1.$$

Let $\mathbb{A} = \{a \in \mathbb{F}_q : v(a) = 0\}$ and $\mathbb{B} = \{a \in \mathbb{F}_q : h(a) = 0\}$ be the sets of roots of v and h in $\mathbb{F}_q$, respectively, and let $r^{-1}(\mathbb{A}) = \{a \in \mathbb{F}_q \backslash \mathbb{B} : r(a) \in \mathbb{A}\}$ be the set of preimages of $\mathbb{A}$ under r. Then, the domain of $l \circ r$ is

$$\mathbb{D}(l \circ r) = \begin{cases} \mathbb{F}_q \backslash (r^{-1}(\mathbb{A}) \cup \mathbb{B}) & \text{if } m > n, \\ \mathbb{F}_q \backslash r^{-1}(\mathbb{A}) & \text{if } m \leq n. \end{cases} \tag{3.1}$$

Since l and r are injective on their domains, respectively, we have

$$\forall a, b \in \mathbb{F}_q \backslash (r^{-1}(\mathbb{A}) \cup \mathbb{B}) \qquad a \neq b \Rightarrow l(r(a)) \neq l(r(b)),$$

and this shows that $l \circ r \in \mathsf{BF}$ if $m > n$.

If $m \leq n$, then $(l \circ r)(a) = \alpha$ for all $a \in \mathbb{B}$. Thus, for $m \leq n$ and $\sigma = \#\mathbb{B}$,

$l \circ r \in \mathsf{BF}$

$\Leftrightarrow \sigma = 0$, or $\sigma = 1$ and $\forall a \in \mathbb{F}_q \backslash (r^{-1}(\mathbb{A}) \cup \mathbb{B})\ l(r(a)) \neq \alpha$

$\Leftrightarrow \sigma = 0$, or $\sigma = 1$ and $\alpha \notin \mathbb{V}(l)$, or $\sigma = 1$ and $\alpha = l(a)$ for some $a \in \mathbb{F}_q \backslash \mathbb{V}(r)$

$\Leftrightarrow \sigma = 0$, or $\sigma = 1$ and $\gcd(x^q - x, u - \alpha v) = 1$, or $\sigma = 1$, $\gcd(x^q - x, u - \alpha v) = x - a$ and $\gcd(x^q - x, g - ah) = 1$ for some $a \in \mathbb{F}_q$. ∎

Remark 3.8. Theorem 3.7 is asymmetric, since for $l \in \mathsf{PF}$ and $r \in \mathsf{BF}$ as in Example 3.4, $l \circ r \notin \mathsf{BF}$, but $r \circ l \in \mathsf{BF}$.

Let l, r, σ, α be as in Theorem 3.7 and $d = \max\{\deg l, \deg r\} \leq q$. The proof of the theorem shows that we can test deterministically whether

$l \circ r \in \mathsf{BF}$ using $O^{\sim}(d \log q)$ operations in $\mathbb{F}_q$. If $l \circ r \notin \mathsf{BF}$, then it has only one multiple image α to which $l \circ r$ maps all σ distinct roots of h in $\mathbb{F}_q$ and possibly another point $b \in \mathbb{F}_q$, determined by

$$\gcd(x^q - x, u - \alpha v) = x - a \quad \text{and} \quad \gcd(x^q - x, g - ah) = x - b,$$

for some $a \in \mathbb{F}_q$.

DEFINITION 3.9. Let $\mathsf{PP} \subseteq \mathsf{PF} \subseteq \mathsf{BF}$ be the classes of permutation polynomials, permutation functions, and bijective functions over $\mathbb{F}_q$, respectively. For $\mathsf{A}, \mathsf{B} \in \{\mathsf{PP}, \mathsf{PF}, \mathsf{BF}\}$, write $\mathsf{A} \circ \mathsf{B} = \{l \circ r : l \in \mathsf{A} \text{ and } r \in \mathsf{B}\}$.

THEOREM 3.10. *Over an arbitrary finite field* $\mathbb{F}_q$, *we have* (*see Fig.* 3.1)

$$\begin{aligned}\mathsf{PP} = \mathsf{PP} \circ \mathsf{PP} \subsetneq \mathsf{PF} = \mathsf{PP} \circ \mathsf{PF} = \mathsf{PF} \circ \mathsf{PP} = \mathsf{PF} \circ \mathsf{PF} \subsetneq \mathsf{BF} \\ = \mathsf{PP} \circ \mathsf{BF} = \mathsf{BF} \circ \mathsf{PP} = \mathsf{BF} \circ \mathsf{PF} \subsetneq \mathsf{PF} \circ \mathsf{BF} \subsetneq \mathsf{BF} \circ \mathsf{BF}.\end{aligned}$$

Proof. All equations follow from Theorem 3.7 (ii) or (i), with $n = 0$. In the other cases, the inclusions "$\subseteq$" follow from $\mathsf{PP} \subseteq \mathsf{PF} \subseteq \mathsf{BF}$. Clearly, $\mathsf{PP} \neq \mathsf{PF} \neq \mathsf{BF}$, and Example 3.4 shows that $\mathsf{BF} \neq \mathsf{PF} \circ \mathsf{BF}$. It remains to show that $\mathsf{PF} \circ \mathsf{BF} \neq \mathsf{BF} \circ \mathsf{BF}$.

We claim that $\mathsf{PF} \circ \mathsf{BF}$ contains only functions that are either total or bijective. To show this, let $l = u/v \in \mathsf{PF}$, $r = g/h \in \mathsf{BF}$, m and n be as in

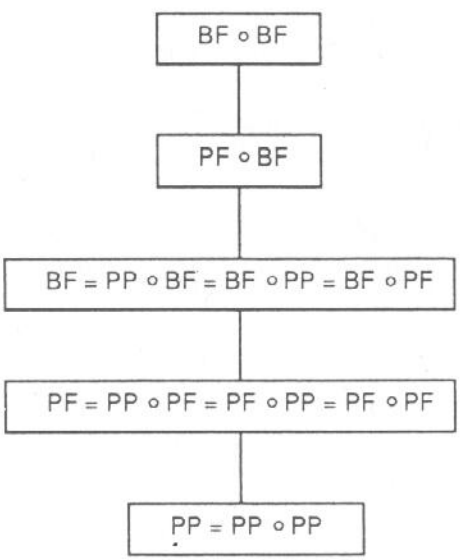

FIG. 3.1. The classes of Theorem 3.10.

Theorem 3.7. It follows from Theorem 3.7 (i) that $l \circ r \in \mathsf{BF}$ if $m > n$, and (3.1) implies that $l \circ r$ is total if $m \leq n$, since $l \in \mathsf{PF}$ and thus $\mathbb{A} = \emptyset$.

Furthermore, we claim that over $\mathbb{F}_2$, every $l \circ r \in \mathsf{PF} \circ \mathsf{BF}$ as above with $m \leq n$ is total with a nonconstant denominator. To prove the claim, we note that since v has no zeros in $\mathbb{F}_2$, it has the form

$$v = x^{n_{2k}} + x^{n_{2k-1}} + \cdots + x^{n_2} + x^{n_1} + 1$$

with $n = n_{2k} > n_{2k-1} > \cdots > n_2 > n_1 > 0$ for some $k \geq 1$.

Let $U, V \in \mathbb{F}_2[x, y]$ be the bivariate homogenizations of u, v of degrees m, n, respectively. In the proof of Theorem 3.7, we showed that $l \circ r = h^{n-m}U(g, h)/V(g, h)$ with $\gcd(V(g, h), h) = \gcd(U(g, h), V(g, h)) = 1$. Let $\hat{n} = \max\{\deg g, \deg h\}$ be the degree of r. It is easily seen that $\deg V(g, h) = n\hat{n} > 0$. This proves our claim.

Now we present $l, r \in \mathsf{BF}$ such that $l \circ r \notin \mathsf{PF} \circ \mathsf{BF}$. We distinguish two cases: $q > 2$ and $q = 2$. Over $\mathbb{F}_q$ with $q > 2$ and for $l, r \in \mathsf{BF}$ as in Example 3.5, $l \circ r$ is neither total nor bijective and hence not in $\mathsf{PF} \circ \mathsf{BF}$.

Over $\mathbb{F}_2$ and for $l = 1/x$, $r = 1/[x(x + 1)] \in \mathsf{BF}$, $l \circ r = x(x + 1)$ is neither bijective nor total with a nonconstant denominator and hence not in $\mathsf{PF} \circ \mathsf{BF}$. ■

4. A Test for Bijective Functions

For f, ν as in (1.1) and $0 \leq \rho \leq q$, we say that f is ρ-large if and only if $\nu \geq q - \rho$. Let $r = \operatorname{res}_x(x^q - x, g - yh)$ and $k = \gcd(y^q - y, r) \in \mathbb{F}_q[y]$. Then

f is ρ-large

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \ \exists b \in \mathbb{F}_q \quad f(b)/g(b) = a$

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \ \exists b \in \mathbb{F}_q \quad f(b) - ag(b) = 0$

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \ \exists b \in \mathbb{F}_q \quad x - b \mid f - ag$

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \ \gcd(x^q - x, f - ag) \neq 1$

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \ \operatorname{res}_x(x^q - x, f - ag) \neq 1$

$\Leftrightarrow \exists \mathbb{W} \subseteq \mathbb{F}_q \ \#\mathbb{W} \geq q - \rho$ and $\forall a \in \mathbb{W} \quad y - a \mid r$

$\Leftrightarrow \deg(\gcd(y^q - y, r)) \geq q - \rho$.

$\Leftrightarrow \deg((y^q - y)/k)) \leq \rho$.

THEOREM 4.1. *For $f = g/h$ as in* (1.1) *and k as above, f is ρ-large if and only if* $\deg((y^q - y)/k) \le \rho$. *In particular, $f \in$ BF if and only if f is σ-large, where $\sigma = \deg \gcd(x^q - x, h)$ is the singularity of f.*

Our test for bijective functions is based on an extension of the test for ρ-large polynomials in von zur Gathen [9] to rational functions. We refer the reader to that paper for an account of the subresultant theory, arithmetic circuits, and Kaltofen's Algorithm Rational Numerator and Denominator.

Test for Bijective Functions

Input: $f = g/h$ as in (1.1), a monic quadratic irreducible polynomial $\varphi \in \mathbb{F}_q[x]$, and $\varepsilon > 0$.
Output: Yes or No.

1. Compute $\sigma = \deg \gcd(x^q - x, h)$.
2. Set $a = (x \bmod \varphi) \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$ and work in two stages:

   (i) Calculate $a_q = a \in \mathbb{F}_{q^2}^\times$ by repeated squaring, and $r(a) = \operatorname{res}_x(x^q - x, g - ah) \in \mathbb{F}_{q^2}^\times$ by Fact 2.2.

   By the subresultant theory, this computation yields an arithmetic circuit A for calculating $(y^q - y)/r \in \mathbb{F}_q(y)$ with no division by zero on input $y \leftarrow a$.

   (ii) Call Kaltofen's Algorithm Rational Numerator and Denominator with input A and a, and degree bound σ both for numerator and denominator.

   The output is an arithmetic circuit B computing two polynomials $c_1, c_2 \in \mathbb{F}_q[y]$ of degree at most σ. Remove divisions from the circuit B to get a division-free circuit C.
3. Repeat $t = \lceil \log_{q/2} \varepsilon^{-1} \rceil$ times:
   Randomly and uniformly choose $a \in \mathbb{F}_{q^2} = \mathbb{F}_q[x]/(\varphi)$, calculate $c_1(a)$ and $c_2(a)$ by executing C with input $y \leftarrow a$, and compute $\beta = (a^q - a)c_2(a) - r(a)c_1(a)$.
   If $\beta \neq 0$, then return No and stop.
4. Return Yes.

THEOREM 4.2. *Let $q \ge n$, $\varepsilon > 0$, $t = \lceil \log_{q/2} \varepsilon^{-1} \rceil$, and $\sigma = \deg \gcd(x^q - x, h)$. The test uses $O(2t)$ random choices and $O(t\, \mathrm{M}^2(\sigma)\mathrm{M}(n) \log q)$ operations in $\mathbb{F}_q$. If $f \in$ BF, the output is* YES; *if $f \notin$ BF, the output is* No *with probability at least $1 - \varepsilon$. In particular, for $\varepsilon = (1/q)^{O(1)}$, the test can be performed with $O(1)$ random choices and $O^\sim(n\sigma^2 \log q)$ operations in $\mathbb{F}_q$.*

Proof. For $r = \operatorname{res}_x(x^q - x, g - yh)$ and $k = \gcd(y^q - y, r) \in \mathbb{F}_q[y]$, let $w_1 = (y^q - y)/k$ and $w_2 = r/k \in \mathbb{F}_q[y]$. Then, $\gcd(w_1, w_2) = 1$, $\deg w_2 \le \deg w_1$ since $\deg r \le q$, and $(y^q - y)/r = w_1/w_2$.

If $f \in \mathsf{BF}$, then $\deg w_1 \le \sigma$ by Theorem 4.1, and Kaltofen's algorithm produces $c_1 = w_1$ and $c_2 = w_2$ in Step 2. In that case, $(y^q - y)c_2 - rc_1 = 0$, and the test returns Yes.

If $f \notin \mathsf{BF}$, then $\deg w_1 > \sigma$ by Theorem 4.1, and Kaltofen's algorithm produces two essentially unrelated polynomials $c_1, c_2 \in \mathbb{F}_q[y]$ of degree at most σ. This shows that $w_1/w_2 \neq c_1/c_2$, and hence $s = (y^q - y)c_2 - rc_1 \in \mathbb{F}_q[y]$ is a nonzero polynomial of degree $\le q + \sigma \le 2q$. Any $a \in \mathbb{F}_{q^2}$ with $\beta = s(a) \neq 0$ constitutes a witness to the nonbijectivity, and the probability of returning a wrong Yes in Step 4 equals that of selecting t liars $a \in \mathbb{F}_{q^2}$ with $\beta = 0$ in a row, which is at most $(2q/q^2)^t = (2/q)^t \le \varepsilon$.

By repeated squaring and a gcd calculation, we can compute σ in Step 1 with $O(\mathsf{M}(n)\log(nq))$ or $O(\alpha)$ operations in $\mathbb{F}_q$, for $\sigma \le n \le q$ and $\alpha = \mathsf{M}(n)\log q$.

Similarly, the circuit size of A in Step 2(i) is $O(\alpha)$. Kaltofen's algorithm uses $O(\mathsf{M}(\sigma)(\alpha + \log \sigma))$ or $O(\mathsf{M}(\sigma)\alpha)$ operations in $\mathbb{F}_q$ to compute the circuit B of size $O(\mathsf{M}(\sigma)\alpha)$ and $O(\mathsf{M}^2(\sigma)\alpha)$ operations to construct a division-free circuit C of size $O(\mathsf{M}^2(\sigma)\alpha)$. The total cost of Steps 1 and 2 is $O(\mathsf{M}^2(\sigma)\alpha)$ operations in $\mathbb{F}_q$.

Step 3 uses $O(2t)$ random choices in $\mathbb{F}_q$ to generate t random elements in $\mathbb{F}_{q^2}$. Each trial takes $O(\mathsf{M}^2(\sigma)\alpha)$ operations in $\mathbb{F}_q$ to calculate $c_1(a)$ and $c_2(a)$ and $O(\alpha)$ operations in $\mathbb{F}_q$ to calculate β. Therefore, the test uses a total number of $O(t\,\mathsf{M}^2(\sigma)\alpha)$ operations in $\mathbb{F}_q$, or $\tilde{O}(n\sigma^2 \log \varepsilon^{-1})$ operations if $\varepsilon \le q^{-1}$. ■

PROPOSITION 4.3. *Let $f \in \mathsf{BF}$ as in (1.1) have degree n. Then f can be evaluated with $O(n)$ operations in $\mathbb{F}_q$ and f^{-1} with $O(\mathsf{M}(n)\log(qn))$ operations.*

Proof. Let $a, b \in \mathbb{F}_q$ with $f(a) = b$. Then $g(a) - bh(a) = 0$, and

$$\gcd(x^q - x,\ g - bh) = x - a.$$

Thus is it easy to calculate a from b. ■

5. PARTIAL PERMUTATION FUNCTIONS

For $f = g/h$, σ, $\mathbb{D}$, $\mathbb{V}$ as in (1.1) we say that f is a partial permutation function if and only if $\mathbb{D} = \mathbb{V}$. We call σ its singularity and denote by $\mathsf{PPF} \subseteq \mathsf{BF}$ the class of all partial permutation functions over $\mathbb{F}_q$.

As a special case of Proposition 3.3, we have the following.

PROPOSITION 5.1. *The class PPF of partial permutation functions represents exactly all permutations on subsets of $\mathbb{F}_q$. Over an arbitrary finite field $\mathbb{F}_q$ and for $0 \le \sigma \le q$, the number of partial permutation functions of*

singularity σ and degree at most q is at least

$$\binom{q}{\sigma}(q-\sigma)!.$$

We identify a special class of partial permutation functions. For $f \in \mathsf{PPF}$, we say that f is a *partial identity function* if and only if $f(a) = a$ for all $a \in \mathbb{D}$ and denote by $\mathsf{PIF} \subseteq \mathsf{PPF}$ the class of all partial identity functions over $\mathbb{F}_q$.

EXAMPLE 5.2. Over an arbitrary finite field $\mathbb{F}_q$ and for $a \in \mathbb{F}_q^\times = \mathbb{F}_q \backslash \{0\}$, $x/(x-a)^{q-1} \in \mathsf{PIF}$ with domain $\mathbb{D} = \mathbb{F}_q \backslash \{a\}$.

For $q > 2$ and $q = 2$, respectively, both $1/x^{q-2}$ and $1/x$ are in PIF with $\mathbb{D} = \mathbb{F}_q^\times$. If char $\mathbb{F}_q \neq 2$, then both

$$\frac{2x}{x^{(q-1)/2}+1} \quad \text{and} \quad \frac{-2x}{x^{(q-1)/2}-1}$$

are in PIF with the domains of squares and nonsquares of $\mathbb{F}_q^\times$ plus the zero element, respectively.

The following is a special case of Proposition 5.1.

PROPOSITION 5.3. *The class* PIF *of partial identity functions represents exactly all identity mappings on subsets of $\mathbb{F}_q$. Over an arbitrary finite field $\mathbb{F}_q$ and for $0 \leq \sigma \leq q$, the number of partial identity functions of singularity σ is at least* $\binom{q}{\sigma}$. *The total number of partial identity functions is at least* 2^q.

It is evident that $\mathsf{PP} \subsetneq \mathsf{PF} \subsetneq \mathsf{PPF} \subsetneq \mathsf{BF}$. In contrast to PP, PF, and BF (see Theorem 3.10), PPF is not closed under composition with any of the four classes, as demonstrated by the following example.

EXAMPLE 5.4. Over an arbitrary finite field $\mathbb{F}_q$, $l = x + 1 \in \mathsf{PP}$ and $r = -1/x \in \mathsf{PPF}$. However, $l \circ r = 1 - 1/x \notin \mathsf{PPF}$, since $\mathbb{D}(l \circ r) = \mathbb{F}_q \backslash \{0\}$ and $\mathbb{V}(l \circ r) = \mathbb{F}_q \backslash \{1\}$. Similarly, $r \circ l = -1/(x+1) \notin \mathsf{PPF}$, since $\mathbb{D}(r \circ l) = \mathbb{F}_q \backslash \{-1\}$ and $\mathbb{V}(r \circ l) = \mathbb{F}_q \backslash \{0\}$.

THEOREM 5.5. *Over an arbitrary finite field $\mathbb{F}_q$, we have*

$$\mathsf{PPF} \subsetneq \begin{matrix} \mathsf{PPF} \circ \mathsf{PP} \subseteq \mathsf{PPF} \circ \mathsf{PF} \\ \mathsf{PP} \circ \mathsf{PPF} \end{matrix} \subseteq \mathsf{BF} \not\supseteq \mathsf{PF} \circ \mathsf{PPF},$$

$$\mathsf{PP} \circ \mathsf{PPF} \subsetneq \mathsf{PF} \circ \mathsf{PPF} \subsetneq \mathsf{PPF} \circ \mathsf{PPF} \not\subseteq \mathsf{PF} \circ \mathsf{BF},$$

$$\mathsf{PF} \circ \mathsf{BF} \subsetneq \mathsf{PPF} \circ \mathsf{BF},$$

$$\mathsf{PP} \circ \mathsf{PPF} = \mathsf{PPF} \circ \mathsf{PP} = \mathsf{PPF} \circ \mathsf{PF} = \mathsf{BF} \text{ if } q \leq 5,$$

$$\mathsf{PP} \circ \mathsf{PPF} \subsetneq \mathsf{BF} \text{ if } q > 5.$$

Proof. All inclusions "$\subseteq$" follow from $\mathsf{PP} \subset \mathsf{PF} \subseteq \mathsf{PPF} \subseteq \mathsf{BF}$ and Theorem 3.10. Example 5.4 shows that $\mathsf{PPF} \neq \mathsf{PP} \circ \mathsf{PPF}$ and $\mathsf{PPF} \neq \mathsf{PPF} \circ \mathsf{PP}$.

We recall Examples 3.4 and 3.5 and note that those l, r stated as in BF are actually in PPF. For $l \in \mathsf{PF}$ and $r \in \mathsf{PPF}$ as in Example 3.4, $l \circ r \notin \mathsf{BF}$ and hence $l \circ r \notin \mathsf{PP} \circ \mathsf{PPF}$. For l, $r \in \mathsf{PPF}$ as in the proof of Theorem 3.10, $l \circ r \notin \mathsf{PF} \circ \mathsf{BF}$ and hence $l \circ r \notin \mathsf{PF} \circ \mathsf{PPF}$. It remains to show that $\mathsf{PP} \circ \mathsf{PPF} = \mathsf{PPF} \circ \mathsf{PP} = \mathsf{PPF} \circ \mathsf{PF} = \mathsf{BF}$ if $q \leq 5$, and $\mathsf{PP} \circ \mathsf{PPF} \subsetneq \mathsf{BF}$ if $q > 5$.

Suppose that $0 \leq k \leq q \leq 5$, and that $f \in \mathsf{BF}$ induces a bijection $\mathbb{A} \to \mathbb{B}$ between $\mathbb{A}, \mathbb{B} \subseteq \mathbb{F}_q$ with $\#\mathbb{A} = \#\mathbb{B} = k$. For such $\mathbb{A}$ and $\mathbb{B}$, we claim that some linear polynomial $l = \alpha x + \beta \in \mathbb{F}_q[x]$ with $\alpha \neq 0$ maps $\mathbb{B}$ onto $\mathbb{A}$. To see this, we note that $l(\mathbb{B}) = \mathbb{A}$ if and only if $l(\mathbb{F}_q \backslash \mathbb{B}) = \mathbb{F}_q \backslash \mathbb{A}$. Since $\min(k, q - k) \leq 2$, such l exists.

For $l = \alpha x + \beta$ with $\mathbb{A} = l(\mathbb{B})$, $\mathbb{D}(\alpha f + \beta) = \mathbb{V}(\alpha f + \beta) = \mathbb{A}$, and $\mathbb{D}(f(\alpha x + \beta)) = \mathbb{V}(f(\alpha x + \beta)) = \mathbb{B}$. Thus $\alpha f + \beta, f(\alpha x + \beta) \in \mathsf{PPF}$, and clearly

$$f = \left(\frac{x}{\alpha} - \frac{\beta}{\alpha}\right) \circ (\alpha f + \beta) = f(\alpha x + \beta) \circ \left(\frac{x}{\alpha} - \frac{\beta}{\alpha}\right).$$

It follows that $\mathsf{PP} \circ \mathsf{PPF} = \mathsf{PPF} \circ \mathsf{PP} = \mathsf{PPF} \circ \mathsf{PF} = \mathsf{BF}$ if $q \leq 5$.

For $q > 5$ (i.e., $q \geq 7$), we first show that there are $\mathbb{A}, \mathbb{B} \subseteq \mathbb{F}_q$ with $\#\mathbb{A} = \#\mathbb{B} \geq 3$, such that no linear polynomial $l \in \mathbb{F}_q[x]$ with $\mathbb{B} = l(\mathbb{A})$ exists.

If $q \geq 8$, we choose $\mathbb{A} = \{0, 1, a, b\} \subseteq \mathbb{F}_q$ with $1 + a^{-1} + b^{-1} \neq 0$ and $\#\mathbb{A} = 4$. Since the number of linear polynomials $l \in \mathbb{F}_q[x]$ is $q(q - 1)$, this is an upper bound on the number of image sets of $\mathbb{A}$ under such l. From

$$\#\{\mathbb{B} \subseteq \mathbb{F}_q : \#\mathbb{B} = 4\} = \binom{q}{4} > q(q - 1),$$

it follows that there exists $\mathbb{B}$ with $\#\mathbb{A} = \#\mathbb{B}$ and such that no linear polynomial maps $\mathbb{A}$ onto $\mathbb{B}$. If $q = 7$, we choose $\mathbb{A} = \{0, 1, 4\}$ and $\mathbb{B} = \{0, 1, 3\}$ and verify that no linear polynomial over $\mathbb{F}_7$ maps $\mathbb{A}$ onto $\mathbb{B}$.

We now use such $\mathbb{A}$ and $\mathbb{B}$ to construct a bijective f such that $f \notin \mathsf{PP} \circ \mathsf{PPF}$, for $q \geq 7$. Let $f = g/h \in \mathsf{BF}$ represent a bijection $\mathbb{F}_q \backslash \mathbb{A} \to \mathbb{F}_q \backslash \mathbb{B}$, with $g, h, \tilde{g}, \tilde{h}$ as constructed in the proof of Proposition 3.2. We may assume that h is monic. We claim that x divides h and x^2 does not.

Since $\tilde{h}(0) = 0$ and $\tilde{g}(0) = 1$, x divides $\tilde{h}$ but not $\tilde{g}$, and it is sufficient to show that $x^2 \nmid \tilde{h}$. Noting that $\tilde{h}(c) = 1$ for all $c \in \mathbb{F}_q \backslash \mathbb{A}$, and $\tilde{h}(c) = 0$ for all $c \in \mathbb{A}$, we have by the Lagrange formula that

$$\tilde{h} = \sum_{c \in \mathbb{F}_q \backslash \mathbb{A}} \prod_{d \in \mathbb{F}_q \backslash \{c\}} \frac{x - d}{c - d} = - \sum_{c \in \mathbb{F}_q \backslash \mathbb{A}} \prod_{d \in \mathbb{F}_q \backslash \{c\}} (x - d),$$

since $\prod_{d \in \mathbb{F}_q^\times} d = -1$ by the general form of Wilson's Theorem (see, e.g., [8, Chap. 1.3]).

Using the fact that $\sum_{c \in \mathbb{F}_q} c = 0$ if $q > 2$, we find that the coefficient s of x in $\hat{h}$ is

$$s = -\sum_{c \in \mathbb{F}_q \setminus \mathbb{A}} \prod_{d \in \mathbb{F}_q \setminus \{0,c\}} (-d) = -\sum_{c \in \mathbb{F}_q \setminus \mathbb{A}} \frac{1}{c} = \sum_{c \in \mathbb{A} \setminus \{0\}} \frac{1}{c}.$$

Our choice of $\mathbb{A}$ implies that

$$s = \begin{cases} 1 + a^{-1} + b^{-1} \neq 0 & \text{if } q \geq 8, \\ 1 + 4^{-1} \neq 0 & \text{if } q = 7. \end{cases}$$

Now we claim that no $l \in \mathsf{PP}$ and $u/v \in \mathsf{PPF}$ exist with $f = l \circ (u/v)$. Suppose that we have such $l, u, v \in \mathbb{F}_q[x]$ with v monic and $\gcd(u, v) = 1$. Let $n = \deg l \geq 1$, $l = \sum_{0 \leq i \leq n} l_i x^i$, and $w = \sum_{0 \leq i \leq n} l_i u^i v^{n-i} \in \mathbb{F}_q[x]$. Then $\gcd(v, w) = 1$, since v divides all summands of w except $l_n u^n$, and $\gcd(v, l_n u^n) = 1$. Furthermore,

$$f = \frac{g}{h} = \sum_{0 \leq i \leq n} l_i \left(\frac{u}{v}\right)^i = \frac{w}{v^n},$$

which implies that $h = v^n$ and hence $n = 1$, since x divides h exactly once. Thus l is a linear polynomial over $\mathbb{F}_q$, and the domain and value set of u/v are

$$\mathbb{D}\left(\frac{u}{v}\right) = \mathbb{V}\left(\frac{u}{v}\right) = \mathbb{F}_q \setminus \mathbb{A}.$$

Since $f = l \circ (u/v)$ and $\mathbb{V}(f) = \mathbb{F}_q \setminus \mathbb{B}$, we have $l(\mathbb{A}) = \mathbb{B}$. By our construction, no such l exists. ■

Several properties of permutation polynomials, such as Carlitz's conjecture [7] or the implication that permutation polynomials are exceptional [4, 10], hold only when the field is sufficiently large compared to the degree. The proper inclusion $\mathsf{PP} \circ \mathsf{PPF} \subsetneq \mathsf{BF}$ exhibits a similar phenomenon; a difference is that these classes contain rational functions of arbitrarily large degrees.

We note that PP and PF are syntactically distinct classes of functions over $\mathbb{F}_q$; however, they are semantically equivalent, or have the same expressive power, in the sense that each represents exactly all permutations on $\mathbb{F}_q$.

The following proposition establishes an interesting connection among the classes PP, PPF, and BF. Combining it with Theorem 5.5, we can conclude that the four classes of functions BF, PP $\circ$ PPF, PPF $\circ$ PP, and PPF $\circ$ PF are semantically equivalent, since each class represents exactly all bijections between subsets of $\mathbb{F}_q$.

PROPOSITION 5.6. *Every bijection between two subsets of $\mathbb{F}_q$ can be represented by a bijective function that is the composition of a permutation polynomial and a partial permutation function in either order.*

Proof. Let $\tau : \mathbb{A} \to \mathbb{B}$ be a bijection between two sets $\mathbb{A}, \mathbb{B} \subseteq \mathbb{F}_q$. We want to show that τ can be represented by two bijective functions f and u such that $f = g \circ h$ and $u = v \circ w$ with $g, w \in \mathsf{PP}$ and $h, v \in \mathsf{PIF} \subseteq \mathsf{PPF}$.

Let $g = w$ be a permutation polynomial over $\mathbb{F}_q$ inducing the restricted bijection $a \mapsto \tau(a)$ between $\mathbb{A}$ and $\mathbb{B}$, and let h be a partial identity function with domain $\mathbb{A}$ and v a partial identity function with domain $\mathbb{B}$. By Theorem 3.10, both $f = g \circ h$ and $g = v \circ w$ are bijective functions, $\mathbb{D}(f) = \mathbb{D}(g) = \mathbb{A}$, and $f(a) = g(a) = \tau(a)$ for all $a \in \mathbb{A}$. ■

Combining Theorems 3.10 and 5.5, we obtain the following corollary.

COROLLARY 5.7. *Over a finite field $\mathbb{F}_q$, we have the composition classes shown in Fig.* 5.1.

6. A TEST FOR PARTIAL PERMUTATION FUNCTIONS

Since any $f = g/h \in \mathsf{PPF}$ is a special bijective function that induces a bijection from its domain into itself, one way of testing whether $f \in \mathsf{PPF}$ is to first test whether $f \in \mathsf{BF}$ and then check that none of the roots of h lie in the value set $\mathbb{V}$ of f. So suppose that $f \in \mathsf{BF}$ and that $x_1, \ldots, x_\sigma$ are the roots of h in $\mathbb{F}_q$, and set

$$k = \prod_{1 \le i \le \sigma} (x - x_i) = \gcd(x^q - x, h) \in \mathbb{F}_q[x].$$

Then

$$\prod_{1 \le i \le \sigma} (g - x_i h) = h^\sigma \prod_{1 \le i \le \sigma} (g/h - x_i) = h^\sigma k(g/h) \in \mathbb{F}_q[x].$$

Denote this polynomial by K; it is the bivariate homogenization of k, evaluated at (g, h). Then we have

$$x_i \notin \mathbb{V} \Leftrightarrow \gcd(x^q - x, g - x_i h) = 1,$$

$$f \in \mathsf{PPF} \Leftrightarrow \gcd\left(x^q - x, \prod_{1 \le i \le \sigma} (g - x_i h)\right) = 1 \Leftrightarrow \gcd(x^q - x, K) = 1.$$

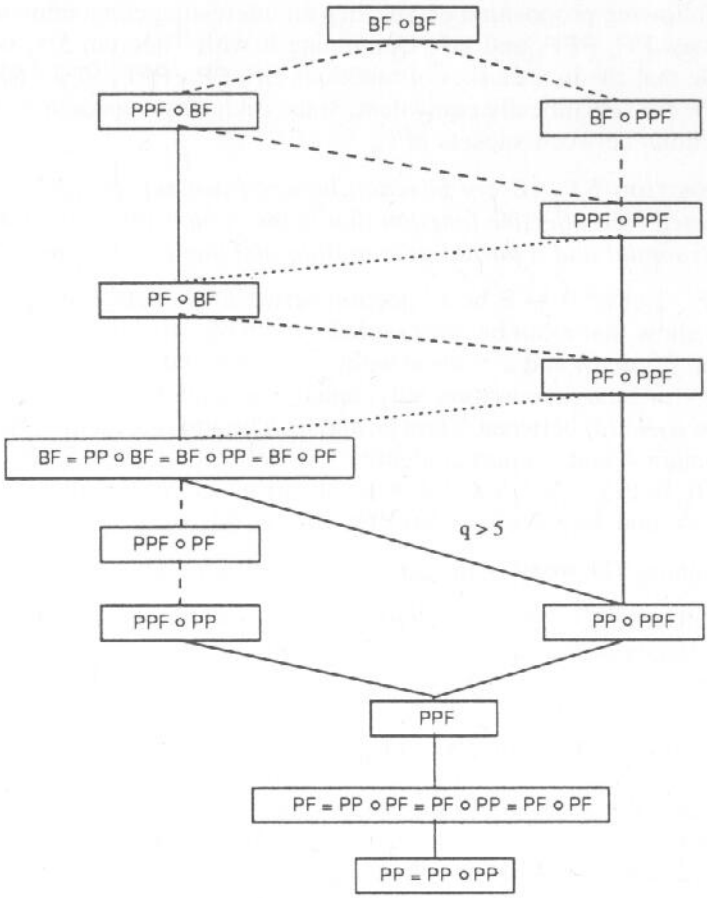

FIG. 5.1. The composition classes of Corollary 5.7. Solid line, lower class is properly contained in upper class; dashed line, containment; dotted line, upper class is not contained in lower class.

The running time of this procedure is $O^{\sim}(n\sigma^2 \log q)$ operations in $\mathbb{F}_q$. In the remainder of this section, we design a simpler and more efficient test, using only $O^{\sim}(n \log q)$ operations in $\mathbb{F}_q$.

For $\mathbb{D} = \mathbb{F}_q \setminus \{x_1, \ldots, x_\sigma\}$, let $d \in \mathbb{F}_q[x]$ be the following polynomial of degree $q - \sigma$:

$$d = \prod_{a \in \mathbb{D}} (x - a) = \frac{x^q - x}{\gcd(x^q - x, h)}. \tag{6.1}$$

Then

$$\begin{aligned} f \in \mathsf{PPF} &\Leftrightarrow \forall a \in \mathbb{D}\ \exists b \in \mathbb{D}\ g(b)/h(b) = a \\ &\Leftrightarrow \forall a \in \mathbb{D}\ \exists b \in \mathbb{D}\ g(b) - ah(b) = 0 \\ &\Leftrightarrow \forall a \in \mathbb{D}\ \exists b \in \mathbb{D}\ x - b \mid g - ah \\ &\Leftrightarrow \forall a \in \mathbb{D}\ \gcd(d, g - ah) \neq 1 \\ &\Leftrightarrow \forall a \in \mathbb{D}\ \mathrm{res}_x(d, g - ah) = 0 \\ &\Leftrightarrow d(y) \mid \mathrm{res}_x(d, g - yh). \end{aligned}$$

Here y is a new indeterminate, $r = \mathrm{res}_x(d, g - yh) \in \mathbb{F}_q[y]$ is the resultant of d and $g \to yh$ in $\mathbb{F}_q(y)[x]$, and the divisibility condition is in $\mathbb{F}_q(y]$. Using the same argument as that in Section 2, we see that $\deg r \leq q - \sigma$ and $r \neq 0$.

THEOREM 6.1. *For $f = g/h$ as in* (1.1), *let $k = \gcd(x^q - x, h)$ and $d = (x^q - x)/k \in \mathbb{F}_q[x]$. Then $f \in$ PPF if and only if*

$$k(y)\mathrm{res}_x(d, g - yh) = c(y^q - y) \qquad \text{for some } c \in \mathbb{F}_q^{\times}.$$

Remark 6.2. An explicit representation of $d \in \mathbb{F}_q[x]$ via polynomial coefficients would require exponential size in $\log q$, and it would be prohibitively expensive to compute $r = \mathrm{res}_x(d, g - yh) \in \mathbb{F}_q[y]$ directly via the determinant of the $(q - \sigma + n) \times (q - \sigma + n)$-Sylvester matrix for d and $g - yh$ in $\mathbb{F}_q(y)[x]$.

For any $a \in \mathbb{F}_{q^m} \supset \mathbb{F}_q$ with $m \geq 2$, however, we can apply Fact 2.2 to compute $r(a) \in \mathbb{F}_{q^m}$ efficiently by the following trick:

- First compute $v = x^q - x \ \mathrm{rem}[k(g - ah)] \in \mathbb{F}_{q^m}[x]$ by repeated squaring and then $u = v/k = d \ \mathrm{rem}(g - ah) \in \mathbb{F}_{q^m}[x]$ by polynomial division. Since g, h, k, v, and u each have degree at most $n + \sigma \leq 2n$, the cost of this step is $O(\mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$.
- Set $n_0 = q - \sigma$, compute the Euclidean representation of $(g - ah, u)$, and then use (2.1) to compute $r(a)$. The cost of this step is $O(\mathsf{M}(n) \log n)$ operations in $\mathbb{F}_q$.

Test for Partial Permutation Functions

Input: $f = g/h$ as in (1.1), a monic quadratic irreducible polynomial $\varphi \in \mathbb{F}_q[x]$, and $\varepsilon > 0$.
Output: Yes or No.

1. Compute $k = \gcd(x^q - x, h) \in \mathbb{F}_q[x]$.

2. Set $a = (x \bmod \varphi) \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$ and calculate
   $r(a) = \mathrm{res}_x((x^q - x)/k, g - ah) \in \mathbb{F}_{q^2}^\times$ by Remark 6.2, $c = k(a)r(a)/(a^q - a) \in \mathbb{F}_{q^2}^\times$ by repeated squaring.
   If $c \notin \mathbb{F}_q^\times$, then return No and stop.
3. Repeat $t = \lceil \log_q \varepsilon^{-1} \rceil$ times:
   Randomly and uniformly choose $a \in \mathbb{F}_{q^2} = \mathbb{F}_q[x]/(\varphi)$, and calculate $\beta = k(a)r(a) - c(a^q - a)$.
   If $\beta \neq 0$, then return No and stop.
4. Return Yes.

The proof of the following theorem is similar to that of Theorem 2.3.

THEOREM 6.3. *Let $q \geq n$, $\varepsilon > 0$, and $t = \lceil \log_q \varepsilon^{-1} \rceil$. The test for partial permutation functions uses $O(2t)$ random choices and $O(t\,\mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$. If $f \in$ PPF, the output is* Yes; *if $f \notin$ PPF, the output is* No *with probability at least $1 - \varepsilon$. In particular, for $\varepsilon = (1/q)^{O(1)}$, the test can be performed with $O(1)$ random choices and $O^\sim(n \log q)$ operations in $\mathbb{F}_q$.*

7. A Probabilistic Approximation Scheme for Counting Value Sets

The basic idea of our probabilistic algorithm for estimating the image size of an arbitrary rational function is the well-known "dart throwing" scheme. Let U be a finite universe of known size #U and $\mathsf{S} \subseteq \mathsf{U}$ a subset of unknown size. If we have an algorithm to decide the membership of a in S for all $a \in \mathsf{U}$, then we can approximate #S as follows:

1. Predetermine an $N \in \mathbb{N}$ as the number of trials and initialize the estimator $\Delta \leftarrow 0$.
2. Repeat N times:
   Randomly and uniformly choose $a \in \mathsf{U}$ (throw a dart) and test whether $a \in \mathsf{S}$ (see where the dart lands). If $a \in \mathsf{S}$, then set
   $$\Delta \leftarrow \Delta + \#\mathsf{U}.$$
3. Output $\Delta \leftarrow \Delta/N$.

DEFINITION 7.1. An (ε, δ)-approximation scheme for #S is a Monte Carlo algorithm with two additional input parameters ε and δ, which outputs an estimate Δ of #S with probability at least $1 - \delta$ and relative error at most ε, i.e.,

$$\mathrm{Prob}[\#\mathsf{S}(1 - \varepsilon) \leq \Delta \leq \#\mathsf{S}(1 + \varepsilon)] \geq 1 - \delta.$$

An (ε, δ)-approximation scheme for #S is said to be *fully polynomial-time* if the running time of the algorithm is polynomial in ε^{-1}, $\log \delta^{-1}$, and the input size of the counting problem.

Fact 7.2. (Karp *et al.* [15]). Let $\beta = \#\mathsf{U}/\#\mathsf{S}$ and $\varepsilon \leq 1$. If $N \geq 4\beta \ln(2/\delta)\varepsilon^{-2}$, then the dart throwing scheme is an (ε, δ)-approximation algorithm for #S.

LEMMA 7.3. *For $f = g/h$, σ, ν, n as in* (1.1), *we have*

$$\frac{q - \sigma}{n} \leq \nu \leq q - \sigma.$$

Proof. We recall the notation of (1.1). It is clear that $\nu \leq \#\mathbb{D} = q - \sigma$.

For $i \in \mathbb{N}$, let $R_i = \{a \in \mathbb{F}_q : \#f^{-1}(\{a\}) = i\}$ be the set of points with exactly i preimages under f, and $r_i = \#R_i$. Clearly, $R_i = \varnothing$ for $i > n$, since for any $a \in \mathbb{F}_q$, $g - ah \in \mathbb{F}_q[x]$ has at most n roots in $\mathbb{F}_q$. Since

$$\bigcup_{0 \leq i \leq n} R_i = \mathbb{F}_q \qquad \text{and} \qquad \bigcup_{1 \leq i \leq n} R_i = \mathbb{V}$$

are partitions of $\mathbb{F}_q$ and $\mathbb{V}$, respectively, it follows that

$$\nu = \sum_{1 \leq i \leq n} ir_i, \quad \sum_{1 \leq i \leq n} ir_i = \#\mathbb{D} = q - \sigma,$$

$$n\nu = n \sum_{1 \leq i \leq n} r_i \geq \sum_{1 \leq i \leq n} ir_i = q - \sigma. \qquad \blacksquare$$

If $n < q$, then $\sigma = \deg \gcd(x^q - x, h) \leq n < q$, and thus $\nu \geq 1$. This yields the upper bound

$$\frac{q}{\nu} \leq n + \frac{\sigma}{\nu} \leq n + \sigma \leq 2n. \tag{7.1}$$

Monte Carlo Approximation Algorithm for ν

Input: $f = g/h$ as in (1.1), $n = \deg f < q$, $0 < \varepsilon$, and $0 < \delta$.
Output: Δ such that $\text{Prob}[\nu(1 - \varepsilon) \leq \Delta \leq \nu(1 + \varepsilon)] \geq 1 - \delta$.

1. Set $N = \lceil 8n \ln(2/\delta)\varepsilon^{-2} \rceil$ and initialize $\Delta \leftarrow 0$.
2. Repeat N times:

   Randomly and uniformly choose an element $a \in \mathbb{F}_q$, and test whether it is an image of f; if so, i.e., if $\gcd(x^q - x, g - ah) \neq 1$, then set

$$\Delta \leftarrow \Delta + q.$$

3. Output $\Delta \leftarrow \Delta/N$.

THEOREM 7.4. *The Monte Carlo approximation algorithm is a fully polynomial-time* (ε, δ)*-approximation scheme for the image size of* $f \in \mathbb{F}_q(x)$ *of degree* $n < q$, *and it uses* $O(n\mathsf{M}(n)\varepsilon^{-2} \log q \log \delta^{-1})$ *operations in* $\mathbb{F}_q$.

Proof. It follows from Fact 7.2 and the upper bound in (7.1) that the algorithm is an (ε, δ)-approximation scheme for ν. Each of the $\lceil 8n \ln(2/\delta)\varepsilon^{-2} \rceil$ iterations takes $O(\mathsf{M}(n) \log q)$ operations in $\mathbb{F}_q$ to calculate $\gcd(x^q - x, g - ah)$.

The total cost of the algorithm is therefore $O(n\mathsf{M}(n)\varepsilon^{-2} \log q \log \delta^{-1})$, or $O^{\sim}(n^2\varepsilon^{-2} \log q \log \delta^{-1})$ operations in $\mathbb{F}_q$. ■

This scheme has been generalized in von zur Gathen *et al.* [11] to estimate the size of the projection of a plane curve.

8. OPEN QUESTIONS

We have developed random polynomial-time tests for permutation functions of three flavors and a fully polynomial-time randomized approximation scheme for counting value sets of arbitrary rational functions over any finite field. It would be very interesting to know whether there exist efficient deterministic algorithms for such decision and counting problems.

Here are some open problems we would like to answer deterministically in time sublinear in q (ideally in time polynomial in $\log q$).

1. Count exactly the number ν of distinct values of an arbitrary rational function $f \in \mathbb{F}_q(x)$. Clearly, this problem is in the complexity class $\#\mathcal{P}$ (see Johnson [14] for terminology): A nondeterministic Turing machine guesses an element $a \in \mathbb{F}_q$ and then checks whether a is an image of f; if so, it accepts and otherwise it rejects. The number of accepting computations thus equals the image size ν of f. Is this problem $\#\mathcal{P}$-complete?

2. If an exact counting algorithm is hard to obtain, approximate ν deterministically in time polynomial in ε^{-1} for the relative error estimate ε.

3. Count the classes PF, PPF, BF with a degree restriction; e.g., determine $\#\{f = g/h \in \mathsf{PF} : \deg g, \deg h \leq d\}$, with $d < q$. How many distinct bijections are represented by these rational functions?

4. Complete the picture of composition classes in Corollary 5.7.

5. Find explicit classes of permutation functions, partial permutation functions, or bijective functions, if possible, easy to compute or with small degree.

ACKNOWLEDGMENTS

This research was supported by the Information Technology Research Centre and the Natural Sciences and Engineering Research Council of Canada. Parts of the second author's work were done during a sabbatical visit to the Institute for Scientific Computation at ETH Zürich, whose hospitality is gratefully acknowledged. We thank two referees for their valuable remarks and suggestions.

REFERENCES

1. A. V. Aho, J. E. Hopcroft, and J. D. Ullman, "The Design and Analysis of Computer Algorithms," Addison–Wesley, Reading, MA, 1974.
2. D. G. Cantor and E. Kaltofen, On fast multiplication of polynomials over arbitrary algebras, *Acta. Inform.* **28** (1991), 693–701.
3. A. H. Clifford and G. B. Preston, The algebraic theory of semigroups, Vol I, *in* "Mathematical Surveys, Number 7," American Mathematical Society, Providence, RI, 1961.
4. S. D. Cohen, The distribution of polynomials over finite fields, *Acta Arith.* **17** (1970), 255–271.
5. L. E. Dickson, The analytic representation of substitutions on a power of a prime number of letters with a discussion of the linear group, *Ann. Math.* **11** (1897), 65–120, 161–183.
6. M. D. Fried and R. Lidl, On Dickson polynomials and Rédei functions, *in* "Proceedings, Salzburg Conference, Contributions to General Algebra," Vol. 5, pp. 139–149, 1986.
7. M. Fried, R. Guralnick, and J. Saxl, Schur covers and Carlitz's conjecture, *Israel J. Math.* **82** (1993), 157–225.
8. L. Gårding and T. Tambour, "Algebra for Computer Science," Springer Verlag, New York, 1988.
9. J. von zur Gathen, Tests for permutation polynomials, *SIAM J. Comput.* **20** (1991), 591–602.
10. J. von zur Gathen, Values of polynomials over finite fields, *Bull. Austral. Math. Soc.* **43** (1991), 141–146.
11. J. von zur Gathen, M. Karpinski, and I. E. Shparlinski, Counting curves and their projections, *in* "Proceedings, 25th ACM Symposium on the Theory of Computing," pp. 805–812, 1993; *Comput. Complex.* **5**, to appear (1995).
12. D. Yu. Grigoryev and M. Karpinski, An approximation algorithm for the number of zeros of arbitrary polynomials over $GF[q]$, *in* "Proceedings, 20th IEEE Symposium on Foundations of Computer Science," pp. 662–669, 1991.
13. J. M. Howie, "An Introduction to Semigroup Theory," Academic Press, London, 1976.
14. D. S. Johnson, A catalog of complexity classes, *in* "Handbook of Theoretical Computer Science" (J. van Leeuwen, Ed.), Vol. A, pp. 68–161, North-Holland, Amsterdam/New York, 1990.
15. R. M. Karp, M. Luby, and N. Madras, Monte-Carlo approximation algorithms for enumeration problems, *J. Algorithms* **10** (1989), 429–448.
16. M. Karpinski and B. Lhotzky, An (ε, δ)-approximation algorithm of the number of zeros of a multilinear polynomial over $GF[q]$, Technical Report TR-91-022, International Computer Science Institute, Berkeley, 1991.
17. M. Karpinski and M. Luby, Approximating the number of solutions of a $GF[2]$ polynomial, *J. Algorithms* **14** (1993), 280–287.

18. D. E. Knuth, The analysis of algorithms, *in* "Proceedings, International Congress of Mathematicians," Vol. 3, pp. 269–274, Nice, 1970.
19. R. Lidl and G. L. Mullen, When does a polynomial over a finite field permute the elements of the field? *Amer. Math. Monthly* **95** (1988), 243–246.
20. R. Lidl and G. L. Mullen, When does a polynomial over a finite field permute the element sof the field?, II, *Amer. Math. Monthly* **100** (1993), 71–74.
21. R. Lidl and H. Niederreiter, "Encyclopedia of Mathematics and Its Applications, Vol. 20, Finite Fields," Addison–Wesley, Reading, MA, 1983.
22. R. Lidl and C. Wells, Chebyshev polynomials in several variables, *J. Reine Angew. Math.* **255** (1972), 104–111.
23. K. Ma and J. von zur Gathen, The computational complexity of recognizing permutation functions, *in* "Proceedings, 26th ACM Symposium on the Theory of Computing," pp. 392–401, Montreal, Quebec, 1994.
24. R. Matthews and R. Lidl. On generalized Rédei functions, *Internat. J. Math. Math. Sci.* **11** (1988), 625–634.
25. G. L. Mullen, Permutation polynomials, *in* "Proceedings, Conference on Finite Fields and Their Applications," Lecture Notes in Pure and Applied Mathematics, Vol. 141, pp. 131–151, Dekker, New York, 1993.
26. L. Rédei, Über eindeutig umkehrbare Polynome in endlichen Körpern, *Acta Sci. Math. Szeged* **11** (1946), 85–92.
27. A. Schönhage and V. Strassen, Schnelle Multiplikation großer Zahlen, *Computing* **7** (1971), 281–292.
28. V. Shoup, Fast construction of irreducible polynomials over finite fields, *J. Symbolic Comput.* **17** (1994), 371–391.
29. I. E. Shparlinski, A deterministic test for permutation polynomials, *Comput. Complexity* **2** (1992), 129–132.
30. V. Strassen, The computational complexity of continued fractions, *SIAM J. Comput.* **12** (1983), 1–27.
31. R. Zippel, Rational function decomposition, *in* "Proceedings, ACM Symposium on Symbolic and Algebraic Computation," pp. 1–6, 1991.