5Gen: A Framework for Prototyping Applications Using Multilinear Maps and Matrix Branching Programs*

Kevin Lewi
Stanford University
klewi@cs.stanford.edu

Alex J. Malozemoff[†]
Galois
amaloz@galois.com

Daniel Apon
University of Maryland
dapon@cs.umd.edu

Brent Carmer[‡]
Oregon State University
carmerb@oregonstate.edu

Adam Foltzer
Galois
acfoltzer@galois.com

Daniel Wagner
Galois
dmwit@galois.com

David W. Archer
Galois
dwa@galois.com

Dan Boneh
Stanford University
dabo@cs.stanford.edu

Jonathan Katz
University of Maryland
jkatz@cs.umd.edu

Mariana Raykova
Yale University
mariana.raykova@yale.edu

ABSTRACT

Secure multilinear maps (mmaps) have been shown to have remarkable applications in cryptography, such as multi-input functional encryption (MIFE) and program obfuscation. To date, there has been little evaluation of the performance of these applications. In this paper we initiate a systematic study of mmap-based constructions. We build a general framework, called **5Gen**, to experiment with these applications. At the top layer we develop a compiler that takes in a high-level program and produces an optimized matrix branching program needed for the applications we consider. Next, we optimize and experiment with several MIFE and obfuscation constructions and evaluate their performance. The **5Gen** framework is modular and can easily accommodate new mmap constructions as well as new MIFE and obfuscation constructions, as well as being an open-source tool that can be used by other research groups to experiment with a variety of mmap-based constructions.

1. INTRODUCTION

A multilinear map (mmap) [12] is an extremely powerful tool for constructing advanced cryptographic systems

*The full version of this paper, with up-to-date performance numbers, is available at <https://eprint.iacr.org/2016/619>

[†]Portion of work done while at University of Maryland.

[‡]Portion of work done while at Yale University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS'16, October 24 - 28, 2016, Vienna, Austria

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4139-4/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2976749.2978314>

including program obfuscation [24], n -party non-interactive key exchange [12], multi-input functional encryption [30, 10], optimal broadcast encryption [13], witness encryption [26], and many others. The recent emergence of candidate mmaps [23, 19, 20, 28] bring these proposals closer to reality, although several of the current candidates have been shown to be too weak for some of these applications, as discussed in §5.

Despite the remarkable power of mmaps, few published works study the efficiency of the resulting applications, primarily due to the rapid pace of development in the field and the high resource requirements needed for carrying out experiments. In this paper we develop a generic framework called **5Gen**¹ (available at <https://github.com/5GenCrypto>) that lets us experiment with powerful applications of current and future mmaps. We focus on two applications in particular: multi-input functional encryption (MIFE) and program obfuscation, both of which can be instantiated with some of the existing mmap candidates (see §5). Our framework is built as a multi-layer software stack where different layers can be implemented with any of the current candidates or replaced altogether as new constructions emerge.

The top layer of our framework is a system to compile a high-level program written in the Cryptol language [21] into a matrix branching program (MBP), as needed for the most efficient MIFE and obfuscation constructions. We introduce several novel optimizations for obtaining efficient MBPs and show that our optimizations reduce both the dimension and the total number of matrices needed.

The next layer implements several variants of MIFE and obfuscation using a provided MBP. This lets us experiment with several constructions and to compare their performance.

¹The name **5Gen** comes from the fact that multilinear maps can be considered the “fifth generation” of cryptography, where the prior four are: symmetric key, public key, bilinear maps, and fully homomorphic encryption.

The lowest layer is the multilinear map library, `libmmap`. We demonstrate our framework by experimenting with two leading candidate mmmaps: GGHLite [23, 32, 2] and CLT [19, 20]. Our experiments show that for the same level of security, the CLT mmap performs considerably better than GGHLite in all the applications we tried, as explained in §8. (Although the GGHL15 multilinear map [28] was not included in our implementation or experiments, we hope that future work might integrate this multilinear map into our framework.)

Our framework makes it possible to quickly plug in new mmmaps as new proposals emerge, and easily measure their performance in applications like MIFE and obfuscation.

MIFE experiments. Recall that functional encryption [11] is an encryption scheme where the decryption key sk_f is associated with a function f . If c is the encryption of message m then decrypting c with the key sk_f gives the decryptor the value $f(m)$ and nothing else. An n -input MIFE scheme is the same, except that the function f now takes n inputs. Given independently created ciphertexts $c_1, \dots, c_n$, with each c_i an encryption of a message m_i and associated with “slot” i , decrypting these ciphertexts using sk_f reveals $f(m_1, \dots, m_n)$ and nothing else.

One important application of 2-input MIFE is *order-revealing encryption* (ORE) [30, 10]. Here the function $f(x, y)$ outputs 1 if $x < y$ and 0 otherwise. Thus, the key sk_f applied to ciphertexts c_1 and c_2 reveals the relative order of the corresponding plaintexts. ORE is useful for responding to range queries on an encrypted database. For large domains, the only known constructions for secure ORE are based on mmmaps. We conduct experiments on ORE using real-world security parameters where mmmaps give the best known secure construction.

We also experiment with 3-input MIFE. Here, we choose a DNF formula f that operates on triples of inputs, which is useful in the context of privacy-preserving fraud detection where a partially trusted gateway needs to flag suspicious transactions without learning anything else about the transactions (see the full version for details). Again, the best known construction for such a scheme uses mmmaps.

We use our framework to evaluate the implementation of these schemes using existing mmmaps for which they are currently believed to be secure. Clearly these systems are too inefficient to be used in practice. Nevertheless, our experiments provide a data point for the current cost of using them. Moreover, our framework makes it possible to easily plug in better or more secure mmmaps as they become available.

Obfuscation experiments. Roughly speaking, an obfuscator takes as input a program and outputs a functionally equivalent program such that the only way to learn information about the program is to run it. We experiment with several obfuscators built on the obfuscator described by Barak et al. [7], including those inspired by Sahai and Zhandry [39] and Ananth et al. [3]. These improvements allow for obfuscation of a point function with increased security at less than half the total obfuscation size reported by Apon et al. [4]. We also implemented the Zimmerman [40] obfuscator, but we ultimately found that it was too inefficient for the settings that we consider in our experiments.

1.1 Our Contributions

Summarizing, we make the following contributions:

- An optimizing compiler from programs written in the Cryptol language to MBPs, which are used in many mmap applications including MIFE and obfuscation. Our compiler uses optimizations such as dimension reduction, matrix pre-multiplication, and condensing the input representation, and solves a constraint-satisfaction problem needed to obtain the most efficient MBP. See §4 for details.
- A library providing a clean API to various underlying mmap implementations. This allows researchers to experiment with different mmmaps, as well as to easily plug future mmmaps into our framework. See §5 for more details.
- A general MIFE construction based on the scheme of Boneh et al. [10] using real-world security parameters. We contribute optimized implementations of two-input MIFE (in particular, order-revealing encryption) and three-input MIFE (in particular, a functionality needed for privacy-preserving fraud detection), as well as performance results that characterize our constructions. See §6 for details and §8 for evaluation results.
- Obfuscation constructions [7, 39, 3, 40] using real-world security parameters. We experiment with obfuscating point functions and evaluate their performance. See §7 for details and §8 for evaluation results.

1.2 Related Work

Several groups have previously implemented mmmaps [2, 19, 20] to experiment with their performance. However, they did not go so far as experimenting with cryptographic applications of mmmaps beyond direct applications such as multi-party non-interactive key exchange. The goal of our work is to explore the performance of more advanced applications such as MIFE and obfuscation. An earlier work implementing obfuscation [4] experimented with obfuscating point functions, and was only able to successfully obfuscate a 14-bit point function.

Our work builds on the vast amount of previous work showing applications of mmmaps—most notably, MIFE [30, 10] and obfuscation [15, 7, 38, 3, 40, 5, 29, 6, 34, 35, 27, 37, 22].

2. PRELIMINARIES

In this section, we introduce notation and also define the various cryptographic constructions that we use in the rest of this work.

Notation. For an integer $n > 0$, we use $[n]$ to denote the set of integers $\{1, \dots, n\}$. We use λ to represent the security parameter, where “ λ -bit security” means that security should hold up to 2^λ clock cycles. We assume that all of our procedures run *efficiently*, or more formally, in polynomial-time with respect to the size of the input to the procedure, and polynomial in the security parameter λ .

Multilinear maps. Boneh and Silverberg [12] first proposed the concept of *multilinear maps* (mmmaps), but it was only in 2013 that Garg, Gentry, and Halevi [23] introduced the first plausible construction of an mmap. Since then, mmmaps have been shown to be powerful tools in solving numerous problems in cryptography.

Figure 3.1: Framework architecture. We use `cryfsm` to compile a Cryptol program (here denoted by `prog.cry`) to an MBP, which can either be used as input into our MIFE implementation or our obfuscation implementation. Both these implementations use `libmmap` as a building block, which supports both the CLT (`libclt`) and GGH-Lite (`libgghlite`) mmapps.

A *multilinear map* [12, 23] (or *graded encoding scheme*) is a primitive for producing randomized encodings of plaintexts that may be publicly added, multiplied, and zero-tested but otherwise “do not reveal any information.” Encodings are associated with a “level” that restricts the types of operations that may be performed on that encoding. More formally, a *degree- κ multilinear map* is a tuple of algorithms (`Setup`, `Encode`, `Add`, `Mult`, `ZeroTest`) where:

- `Setup` takes as input the security parameter, and outputs a private parameter `sp` and a public parameter `pp` that, in particular, specifies a ring R .
- `Encode` takes `sp`, an element $x \in R$, and a level $S \subseteq [\kappa]$, and outputs a *level- S encoding* of x denoted $\llbracket x \rrbracket_S$.
- `Add` takes `pp` and two encodings $\llbracket x \rrbracket_S, \llbracket y \rrbracket_S$ at the *same level* S , and outputs an encoding $\llbracket x + y \rrbracket_S$.
- `Mult` takes `pp` and two encodings $\llbracket x \rrbracket_{S_1}, \llbracket y \rrbracket_{S_2}$ for *disjoint* levels S_1, S_2 , and outputs an encoding $\llbracket x \cdot y \rrbracket_{S_1 \cup S_2}$.
- `ZeroTest` takes `pp` and an encoding $\llbracket x \rrbracket_U$ for $U = [\kappa]$. It outputs 1 if and only if $x = 0$.

Informally, an mmap is secure if the only information that an attacker can figure out from the encodings of random elements is exactly the information that can be obtained from running `Add`, `Mult`, and `ZeroTest`, and no more. (We omit any formal definitions since we do not directly rely on them in this work but instead inherit them from prior work.)

Matrix branching programs. A matrix branching program (MBP) on length- n , base- d inputs is a collection of variable-dimension matrices $\mathbf{B}_{i,j}$ for $i \in [n]$ and $j \in \{0, \dots, d-1\}$, along with a “final matrix” $\mathbf{P}$. We require that, for each $i \in [2, n]$ and $j \in \{0, \dots, d-1\}$, the number of columns of $\mathbf{B}_{i-1,j}$ is equal to the number of rows of $\mathbf{B}_{i,j}$, so that the product of these matrices is well-defined. The evaluation of an MBP on input $x \in \{0, \dots, d-1\}^n$ is defined as

$$\text{MBP}(x) = \begin{cases} 1, & \text{if } \prod_{i=1}^n \mathbf{B}_{i,x_i} = \mathbf{P}, \\ 0, & \text{otherwise.} \end{cases}$$

We note that numerous generalizations and extra properties [10, 39] of MBPs have been explored in the literature—however, we will only need to use our simplified definition of MBPs for the remainder of this work.

3. FRAMEWORK ARCHITECTURE

Our framework incorporates several software components that together enable the construction of applications using

mmapps and MBPs. In particular, we use our framework to develop implementations of MIFE and program obfuscation. See Figure 3.1 for the framework architecture.

The top layer of our framework, `cryfsm`, takes as input a program written in Cryptol [21], a high-level language designed to express manipulations over bitstreams in a concise syntax, and compiles the program into an MBP. This process, and the various optimizations we introduce, are described in more detail in §4.

The bottom layer of our framework, `libmmap`, provides an API for using various mmapps, which in our case includes the CLT (through the `libclt` library) and GGHlite (through the `libgghlite` library) mmapps. The `libmmap` library, which we describe in §5, is also designed to allow for a straightforward integration of future mmap implementations.

We combine the above components to realize various applications of mmapps and MBPs: in particular, MIFE and program obfuscation. We demonstrate the applicability of our MIFE implementation (cf. §6) through two examples: order-revealing encryption (ORE) and, in the full version, three-input DNF (3DNF) encryption. We implement program obfuscation based on two main approaches: the techniques described by Sahai and Zhandry [39], and also the scheme by Zimmerman [40], which operates over arithmetic circuits, but only applies to the CLT mmap.

4. FROM PROGRAMS TO MBPS

One of our key contributions in this work is a compiler, `cryfsm`, that takes as input a program written in Cryptol [21], a domain-specific language for specifying algorithms over generic streams of bits, and produces an MBP for the given input program. `cryfsm` does this by translating a Cryptol specification into a *layered state machine*, which can then be transformed into an optimized corresponding MBP.

Our toolchain proceeds as follows. The user writes a Cryptol function of type $[n] \rightarrow \text{Bit}$ for some n (that is, the function takes n input bits and produces one output bit). This function is interpreted as deciding membership in a language. The toolchain symbolically evaluates this function to produce a new version of the function suitable for input to an SMT solver, as explained in detail below. Queries to the SMT solver take the form of deciding the prefix equivalence relation between two initial bitstrings, which is sufficient to build the minimal layered state machine, which we then convert to an MBP.

Our solver-based approach results in a substantial dimension reduction of the corresponding output MBPs that we tested. In contrast, the traditional approach would be to heuristically optimize the state machine design in an attempt to achieve a best-effort optimization. The dimension reduction we achieve recovers the most efficient known MBPs for several previously studied bit-string functions, including MBPs for point functions that are smaller than the MBPs constructed from boolean formulas using existing techniques (e.g., [39]). In the remainder of this section, we describe the key steps in this toolchain, along with several optimizations to the MBPs that we use throughout the remainder of this work.

Specifying functions in Cryptol. Cryptol is an existing language widely used in the intelligence community for describing cryptographic algorithms. A well-formed Cryptol program looks like an algorithm specification, and is exe-

cutable. The Cryptol tool suite supports such execution, along with capabilities to state, verify, and formally prove properties of Cryptol specifications, and capabilities to both prove equivalence of implementation in other languages to Cryptol specifications and automatically generate such implementations. In our work, a user specifies an MBP in Cryptol, and then we use `cryfsm` to transform the high-level specification into a minimal layered state machine, and further transform it into an efficient MBP.

Minimal layered state machines. There is a standard translation from traditional finite state machines to MBPs: create a sequence of matrix pairs (or matrix triples for three-symbol alphabets, etc.) that describe the adjacency relation between states. If state i transitions to state j on input symbol number b , then the b th MBP matrix will have a 1 in the i th row and j th column and 0 elsewhere. For many languages of interest, this is inefficient: for an automaton with $|S|$ states, each matrix must be of size $|S|^2$, even though many states may be unreachable.

In the applications of mmaps that we study in this work, we consider functions on inputs of a fixed length. Hence, for a positive integer n , we can take advantage of this property by restricting ourselves to *layered state machines* of depth n , which are simply (deterministic) finite state machines that only accept length- n inputs. Here, the i th “layer” of transitions in the machine is only used when reading the i th digit of the input. As a result, layered state machines are acyclic.

To generate minimal layered state machines, our compiler must introduce machinery to track which states are reachable at each layer, which allows us to reduce the overall MBP matrix dimensions. To do this, `cryfsm` computes the quotient automaton of the layered state machine using an SMT solver to decide the state equivalence relation. The quotient automaton is then used as the new minimal layered state machine for the specified function. Then, from a layered state machine of depth n , we construct the corresponding MBP on base- d inputs of length n in a manner essentially equivalent to the techniques of Ananth et al. [3] for constructing layered branching programs. Intuitively, for each $i \in [n]$ and $j \in [d]$, the i th matrix associated with the j th digit is simply the adjacency matrix corresponding to the transitions belonging to the i th layer of the machine, associated with reading the digit j . Then, the “final matrix” (that defines the output of the MBP being 1) is simply the adjacency matrix linking the initial state to the final state of the layered state machine.

Optimizations for MBP creation. Boneh et al. [10] describe a simple five-state finite state machine appropriate for ORE applications, and describe the translation to MBPs that produces 5×5 matrices at each depth. The MBP we build and use for our ORE application differs from this one via three transformations that can be generalized to other programs: change of base, matrix premultiplication, and dimension reduction. Of these, matrix premultiplication and dimension reduction are a direct consequence of the technique used by `cryfsm` for constructing MBPs and therefore automatically apply to all programs, whereas choosing an input base remains a manual process because it must be guided by outside knowledge about the performance characteristics of the mmap used to encode the MBPs. While the change of base and matrix pre-multiplication optimizations are described by Boneh et al., we introduce dimension reduction as a new optimization that is useful for ORE yet

generalizable to other applications.

For each optimization, we use the integer d to represent the “input base”, the integer n to represent the length (number of digits) of each input, the integer N to represent the input domain size (so, we have that $d^n \geq N$), the integer m to represent the length of the MBP, and the integer M to represent the total number of elements across all the matrices of the MBP.

At a high level, the optimizations are as follows.

- **Condensing the input representation** corresponds to processing multiple bits of the input, by increasing d , to reduce the length of the MBP, at the expense of increasing the number M of total elements.
- **Matrix premultiplication** also aims to reduce the parameter m , but without increasing the parameter M .
- **Dimension reduction** aims to directly reduce the number M of total elements, but may not be fully compatible with matrix premultiplication, depending on the function.

To help with the understanding of the intuition behind these optimizations, we use the simple comparison state machine as a running example—however, we stress that these optimizations are in no way specific to the comparison function, and can be applied more generally to any function expressed as a layered state machine.

Condensing the input representation. The most immediate optimization that we apply is to condense the representation of inputs fed to our state machines. MBPs are traditionally defined as operating on bitstrings, so it is natural to begin with state machines that use bits as their alphabet, but using larger alphabets can cut down on the number of state transitions needed (at the potential cost of increasing the state space).

As an example, for evaluating the comparison state machine, this optimization translates to representing the input strings in a larger base $d > 2$, and to adjust the comparison state machine to evaluate using base- d representations. The resulting state machine consists of $d + 3$ total states.

A naive representation of an input domain of size N with a state machine that processes the inputs bit-by-bit (in other words, $d = 2$) would induce an MBP length of $m = 2 \cdot \lceil \log_2(N) \rceil$ and $M = 50 \cdot m$ total elements (in two 5×5 matrices). However, by using the corresponding comparison state machine that recognizes the language when the inputs are in base- d , we can then set $m = 2 \cdot \lceil \log_d(N) / \log_2(d) \rceil$ and $M = 2 \cdot (d + 3)^2 \cdot m$.

Concretely, setting $N = 10^{12}$, without condensing the input representation, we require $m = 80$ and $M = 2000$ for the resulting MBP. However, if we represent the input in base-4, we can then obtain $m = 20$ and $M = 1960$, a strict improvement in parameters.

Matrix premultiplication. Boneh et al. [10] informally describe a simple optimization to the comparison state machine, which we explain in more detail here. The natural state machine for evaluating the comparison function on two n -bit inputs x and y reads the bits of x and y in the order $x_1 y_1 x_2 y_2 \dots x_n y_n$.

However, Boneh et al. show that a slight reordering of the processing of these input bits can result in reduced MBP length without compromising in correctness. When the in-

puts are instead read in the following order:

$$x_1 y_1 y_2 x_2 x_3 y_3 \cdots y_n x_n, \quad (1)$$

then, rather than producing one matrix for each input bit position during encryption, the two matrices corresponding to y_1 and y_2 can be pre-multiplied, and the result is a *single* matrix representing two digit positions. Naturally, this premultiplication can be performed for each pair of adjacent bit positions belonging to the same input string (such as for $x_2 x_3$, $y_3 y_4$, and so on), and hence the number of matrices produced is slightly over half of the number of matrices in the naive ordering of input bits.

As a result, for evaluating the comparison state machine, where n is the length of the base- d representation of an input, applying this optimization implies $m = n + 1$, a reduction from the naive input ordering, which would result in $m = 2n$, and a reduction from $M = 2 \cdot (d + 3)^2 \cdot m$ to $M = (d + 3)^2 \cdot m$. When applying this optimization in conjunction with representing the input in base $d = 4$, for example, setting $N = 10^{12}$ only requires $m = 21$ and $M = 1029$, a huge reduction in cost that was emphasized by Boneh et al., and another strict improvement in parameters.

A new optimization: dimension reduction. We now describe a more sophisticated optimization that can be applied to general MBPs which also results in a reduced ciphertext size. As an example, we describe this optimization, called dimension reduction, as it applies to the comparison function state machine (without applying the reordering of input bits from matrix premultiplication), but we emphasize that the technique does not inherently use the structure of this state machine in any crucial way, and can naturally be extended to general MBPs.

Our new optimization stems from the observation that, for each bit position in the automaton evaluation, the transitions in the automaton do not involve all of the states in the automaton. This is the same observation that motivates the use of layered state machines over finite state machines.

In particular, for the even-numbered bit positions, the transitions map from a set of d states to a set of only 3 states. Similarly, for the odd-numbered bit positions, the transitions map from a set of (at most) 3 states to a set of d states. As a result, the corresponding matrices for each bit position need only be of dimension $d \times 3$ or $3 \times d$ (depending on the parity), as opposed to the naive interpretation of the Boneh et al. construction which requires matrices of dimension $(d + 3) \times (d + 3)$.

Note, however, that the dimension reduction optimization is not fully compatible with matrix premultiplication, since the effectiveness of dimension reduction can degrade if matrix premultiplication is also applied. In particular, when applying matrix premultiplication to the comparison state machine, we notice that there is less room for improvements with dimension reduction, as the transitions for the position $y_1 y_2$ correlate from a domain of d states to a range of also d states.

In §6.1, we concretely show how to apply a mixture of these optimizations to the comparison automaton, and then use these optimizations to obtain asymptotically shorter ciphertexts for order-revealing encryption.

vtable	function	comments
mmap_pp_vtable	fread/fwrite clear	read/write public parameters clear public params
mmap_sk_vtable	init/clear fread/fwrite	initialize/clear secret key read/write secret key
mmap_enc_vtable	init/clear fread/fwrite	initialize/clear encoding read/write encoding
	set	copy encoding
	add	implements Add
	mul	implements Mult
	is_zero	implements ZeroTest
	encode	implements Encode

Table 5.1: Interfaces exported by the libmmap library.

5. A LIBRARY FOR MULTILINEAR MAPS

In this section we describe our library, **libmmap**, which provides an API for interacting with different mmap backends. In this work we implement GGHLite (**libgghlite**) and CLT (**libclt**) backends², although we believe that it should be relatively straightforward to support future mmap implementations.

The **libmmap** library exports as its main interface a virtual method table **mmap_vtable**, which in turn contains virtual method tables for the public parameters (**mmap_pp_vtable**), the secret key (**mmap_sk_vtable**), and the encoded values (**mmap_enc_vtable**). Table 5.1 lists the available functions within each table. Each underlying mmap library must export functions matching these function interfaces and write a wrapper within **libmmap** to match the virtual method table interface. A user of **libmmap** then defines a pointer **const mmap_vtable *** which points to the virtual method table corresponding to the mmap of the user’s choice (in our case, either **clt_vtable** or **gghlite_vtable**). In the following, we describe the two mmap schemes we support within **libmmap**: **libgghlite** (§5.1) and **libclt** (§5.2).

Figure 5.1 presents estimates for the size of an encoding using GGHLite and CLT for security parameters $\lambda = 80$ and $\lambda = 40$. We describe our parameter choices for arriving at these estimates in the full version. As we can see, the CLT mmap produces smaller encodings than GGHLite as we vary both λ and κ . This appears to be due to the growth of the lattice dimension in GGHLite compared to the number of secret primes required by the CLT scheme, among other factors.

5.1 The GGHLite Multilinear Map

Building off of the original mmap candidate construction of Garg et al. (GGH) [23], Langlois et al. [32] proposed a modification called GGHLite, along with parameter and performance estimates for the resulting encodings of the scheme. More recently, Albrecht et al. [2] proposed further modifications and optimizations on top of GGHLite, along with an implementation of their scheme under an open-source license. In this work, we refer to GGHLite as the construction from the work of Albrecht et al., as opposed to the original work of Langlois et al.

Our GGHLite implementation. We use as our starting point the implementation of GGHLite³ released by Albrecht et al. [2]. We modified this implementation to add

²We also have a “dummy” mmap implementation for testing purposes.

³<https://bitbucket.com/malb/gghlite-flint>

Figure 5.1: Estimates for the size of a single encoding in megabytes (MB) produced for security parameters $\lambda = 80$ and $\lambda = 40$ and varying the multilinearity degree $\kappa \in [2, 30]$ for the GGHLite and CLT mmap.

functionality for handling the reading and writing of encodings, secret parameters, and public parameters to disk. We also extended the implementation to handle more expressive index sets, which are used in MIFE and obfuscation, as follows.

Typically, multilinear maps only support “levels”, where each encoding is created with respect to an integer $i \in [\kappa]$ (for an mmap of degree κ). The GGHLite implementation supports more advanced labelings of encodings, by allowing for a universe U of κ indices to be defined, and each encoding can be created with respect to a singleton subset (containing only one element) of this universe U . Multiplication of two encodings with respect to sets of indices S_1 and S_2 produces an encoding with respect to the multiset union of S_1 and S_2 . The zero-testing parameter is then created to test for encodings which are labeled with respect to U . However, this functionality is still not sufficiently expressive to match the needs of our implementation and our definition of mmaps.

Consequently, we upgraded the handling of these encodings to support labelings of an encoding with respect to *any* subset S of indices of the universe U . Then, when two encodings labeled with two different subsets are multiplied, the resulting encoding is labeled with respect to their multi-set union. Finally, as before, the zero-testing parameter allows to check for encodings of 0 labeled at U , only.

Finally, we isolated and rewrote the randomness generation procedures used by GGHLite, since the original implementation relied on the randomness obtained from the GMP library, which is not generated securely. We split this into a separate library, `libaesrand`, which uses AES-NI for efficient randomness generation, and which may be useful in other contexts.

Attacks on GGHLite. Recently, Hu and Jia [31] showed how to perform “zeroizing” attacks on GGHLite, to recover the secret parameters given certain public encodings of 0. However, since neither MIFE nor obfuscation publish any encodings of 0, these applications seem to be unaffected by the zeroizing attacks. More recently, Albrecht, Bai, and Ducas [1] gave a quantum break for GGHLite *without* using any encodings of 0 or the public zero-testing parameter. Subsequently, Cheon, Jeong, and Lee [17] showed how to give a (classical) polynomial-time attack on GGHLite, again without using any encodings of 0. However, their attack re-

quires exponential time if the parameters of GGHLite are sufficiently increased (by a polynomial amount).

In concurrent work, Miles, Sahai, and Zhandry [36] gave a completely different form of attack, known as an “annihilation” attack, on *applications* of GGHLite, specifically, MIFE and program obfuscation. They show that provably secure instantiations of these primitives from mmaps are in fact insecure when the mmap is instantiated with GGHLite. Despite the annihilation attacks, our implementations of these primitives from GGHLite still serve as a useful benchmark for the efficiency of GGHLite and for the efficiency of future GGH-like schemes resistant to annihilation attacks, which will inevitably arise from improvements to the GGH framework.

5.2 The CLT Multilinear Map

Coron, Lepoint, and Tibouchi [19] proposed a candidate multilinear map over the integers, which works over a composite modulus that is assumed to be hard to factor.

Our CLT implementation. Our implementation started with the implementation⁴ of CLT in C++ by Coron et al. [19]. We rewrote it in C and added functionality to save and restore encodings and the public parameters. As in the GGHLite case, we also modified its basic functionality to support indices instead of levels.

Furthermore, in our extension of CLT, we improve the efficiency of the encoding process which allows for us to apply the CLT multilinear map to the large parameter settings that we consider in the remainder of this work. The original CLT implementation applies the Chinese Remainder Theorem in the procedure that produces encodings of plaintext elements. Our implementation employs a certain trade-off that allows for the application of the Chinese Remainder Theorem in a recursive manner, resulting in more multiplications to compute the encoding, but with the efficiency gain that the elements being multiplied are much smaller. Experimentally, this yields a large speedup in the encoding time, more noticeably with larger parameters. In particular, for $\lambda = 80$ and $\kappa = 19$, without this optimization, it takes 134 seconds to produce a CLT encoding, whereas with our optimization, this time drops to 33 seconds.

Attacks on CLT. Similarly to other candidate constructions for multilinear maps, the CLT construction was not based on an existing hardness assumption but rather introduced a new assumption. Subsequently Cheon et al. [16] demonstrated a zeroizing attack against the construction of CLT, which succeeds in recovering the secret parameters of the scheme. This attack was further extended in the work of Coron et al. [18], which demonstrated how it can be generalized and applied against some proposed countermeasures [14, 25] to the attack by Cheon et al. [16]. But again, as with the zeroizing attacks on GGHLite, these results do *not* apply directly to the constructions we consider in this work.

6. MULTI-INPUT FUNCTIONAL ENCRYPTION

The notion of multi-input functional encryption (MIFE), introduced by Goldwasser et al. [30], extends the concept of functional encryption [11] so that a decryption key is asso-

⁴<https://github.com/tlepoint/multimap>

ciated with a *multi-input* function which is evaluated over multiple ciphertexts. More formally, a secret-key, fixed-key MIFE scheme for a function f , on m inputs and with output in a range $\mathcal{R}$, is a tuple of algorithms (keygen , encrypt , eval) such that:

- $\text{keygen}(1^\lambda) \rightarrow (\text{pp}, \text{sk})$. The algorithm takes as input the security parameter and generates the public parameters pp and a secret key sk .
- $\text{encrypt}(\text{sk}, i, x) \rightarrow \text{ct}$. The algorithm takes as input the secret key sk , an input position index i , and an input x , and outputs a ciphertext ct .
- $\text{eval}(\text{pp}, \text{ct}_1, \dots, \text{ct}_m) \rightarrow z$. The algorithm takes as input a secret key sk and m ciphertexts $\text{ct}_1, \dots, \text{ct}_m$, and produces an output $z \in \mathcal{R}$.

Correctness requires that for any inputs $x_1, \dots, x_m$, for $(\text{pp}, \text{sk}) \leftarrow \text{keygen}(1^\lambda)$, letting $\text{ct}_i = \text{encrypt}(\text{sk}, i, x_i)$ for each $i \in [m]$, we have that

$$\text{eval}(\text{pp}, \text{ct}_1, \dots, \text{ct}_m) = f(x_1, \dots, x_m).$$

Informally, an MIFE scheme is *secure* if the information revealed by a collection of ciphertexts is exactly the information that can be obtained by running eval , and no more. We omit formal security definitions since we do not directly rely on them in this work.

Goldwasser et al. [30] gave a general MIFE construction that uses indistinguishability obfuscation in a black-box manner. Boneh et al. [10] proposed a secret-key MIFE construction that is based directly on mmaps (instead of obfuscation) in order to obtain better efficiency. A particular instantiation of this MIFE construction, where the function used in the decryption key is the comparison function, results in a construction for *order revealing encryption* (ORE), which allows comparisons over ciphertexts while hiding all other information about the encrypted messages. More specifically, an ORE scheme is a MIFE scheme with the function $f(x, y)$ that outputs the ordering between x and y . This ORE scheme achieves the optimal security definition for a scheme that allows the comparison functionality over encrypted data, improving over the security level provided by order-preserving encryption schemes.

MIFE implementation. We implemented the Boneh et al. [10] MIFE construction on top of the `libmmap` library, and provide interfaces for `keygen`, `encrypt`, and `eval`, which perform the respective operations supported by MIFE. We parallelize the computation performed during `keygen`, but for `encrypt`, we choose to sequentially construct the encodings belonging to the ciphertext, and instead defer the parallelism to the underlying mmap implementation for producing encodings, in the interest of reducing memory usage at the cost of potentially increased running times. We note that, since CLT enjoys much more parallelism than GGHLite when constructing encodings, this optimization causes the `encrypt` time for GGHLite to be less efficient. Finally, for `eval`, we multiply encodings in parallel for CLT, since the multiplication of CLT encodings natively does not support parallelism. However, for GGHLite, we choose to multiply encodings sequentially, and instead rely on the parallelism afforded by GGHLite encoding multiplication.

The ciphertexts produced by a call to `encrypt` on an ℓ -length input are split into ℓ components (one for each input slot), which can be easily separated and combined with different components from other ciphertexts in a later call to

`eval`. Hence, with a collection of full ciphertexts, an evaluator can specify which components from each ciphertext should be passed as input to `eval`, in order to evaluate the function on components originating from different sources.

6.1 Optimizing Comparisons

In this section, we describe a case study of applying the optimizations detailed in §4 to the comparison function. We establish two distinct “variants” of the comparison function which result in shorter ciphertext sizes. Both variants are built from a combination of condensing the input representation into a larger base $d > 2$, followed by dimension reduction, and optionally applying matrix pre-multiplication.

- **DC-variant.** The *degree-compressed* optimization is to first apply matrix pre-multiplication to re-order the reading of the input bits as in Equation (1). Then, the dimensions of the resulting matrices from the layered state machine are slightly reduced.
- **MC-variant.** The *matrix-compressed* optimization is to directly apply dimension reduction in the normal interleaved ordering of the bits (as $x_1y_1x_2y_2 \dots x_ny_n$). Here, the dimensions of the matrices can be reduced to depend only linearly in the base representation d , as opposed to quadratically.

We now discuss each optimization in more detail.

The degree-compressed variant (DC-variant). By optimizing the (layered) comparison state machine, we obtain that not all matrices need to be of dimension $(d+3) \times (d+3)$. For example, the first matrix need only be of dimension $1 \times d$, and the second matrix need only be of dimension $d \times (d+2)$. Also, the last matrix can be of dimension $(d+2) \times 3$. And finally, each of the remaining intermediate matrices need only be of dimension $(d+2) \times (d+2)$. Putting these observations together, the total number of encodings in the ciphertext is

$$\begin{aligned} M &= d + d(d+2) + 3(d+2) + (\kappa-3)(d+2)^2 \\ &= d^2(\kappa-2) + (d+1)(4\kappa-6). \end{aligned} \quad (2)$$

The matrix-compressed variant (MC-variant). Note, however, that if we do not apply the matrix pre-multiplication optimization, but instead apply dimension reduction directly to the comparison state machine associated with the normal (not interleaved) ordering of the input digits, then the first matrix is of dimension $1 \times d$, the second matrix is of dimension $d \times 3$, and all other $\kappa-2$ matrices are of dimension either $3 \times (d+2)$ or $(d+2) \times 3$. Putting this together, we have $\kappa = 2n$ and

$$\begin{aligned} M &= d + 3d + 3(\kappa-2)(d+2) \\ &= 3(\kappa-2)(d+2) + 4d. \end{aligned} \quad (3)$$

Concretely, for a domain of size $N = 10^{12}$, if we choose to represent the inputs in base $d = 5$, then this implies $n = 18$ (since $12 \leq \log_{10}(5^{18}) < 13$), and hence with the matrix pre-multiplication optimization along with the Cryptol optimization for dimension reduction, we have $\kappa = 19$ and $M = 845$. Without applying matrix pre-multiplication and only using dimension reduction with base $d = 10$, we have $n = 12$, $\kappa = 24$, and $M = 832$.

Experimentally for the mmmaps we tested, we found that the matrix pre-multiplication optimization produces shorter

ciphertexts than applying dimension reduction without matrix pre-multiplication. However, we only tested this for an input domain of $N = 10^{12}$ and security parameter $\lambda = 80$. As N grows larger, depending on the asymptotic behavior of encoding sizes as κ increases and λ varies, future implementations of the comparison state machine may find that one can produce shorter ciphertexts when applying dimension reduction without matrix pre-multiplication.

6.2 Order-Revealing Encryption

To implement order-revealing encryption, we set our plaintext domain to the numbers in the range $[N]$. By taking $N = 10^{12}$, we found that selecting the base representation $d = 5$ and applying the matrix pre-multiplication optimization resulted in using only $\kappa = 19$ levels of the underlying mmap, which achieved the shortest ciphertexts for this domain. In fact, this construction yields the shortest known ciphertexts for ORE on a domain of size 10^{12} , as explained below.

An alternative (basic) construction. The closest competitor to our ORE construction in terms of ciphertext size and overall efficiency is a construction due to Lewi and Wu [33], which we refer to as the “basic” ORE scheme, described below.

Let $[N]$ be the message space. Let $F : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ be a secure pseudorandom function (PRF) and $H : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}$ be a hash function (modeled as a random oracle). Let CMP be the comparison function, defined as $\text{CMP}(x, y) = 1$ if $x < y$ and $\text{CMP}(x, y) = 0$ if $x > y$. The basic ORE scheme Π_{ore} is defined as follows.

- **keygen** $(1^\lambda) \rightarrow (\text{pp}, \text{sk})$. The algorithm samples a PRF key $k \xleftarrow{\mathcal{R}} \{0, 1\}^\lambda$ for F , and a random permutation $\pi : [N] \rightarrow [N]$. The secret key sk is the pair (k, π) , and there are no public parameters.
- **encrypt** $(\text{sk}, i, x) \rightarrow \text{ct}$. Write sk as (k, π) . If $i = 1$, the ciphertext output is simply $\text{ct} = (F(k, \pi(x)), \pi(x))$. If $i = 2$, then the encryption algorithm samples a nonce $r \xleftarrow{\mathcal{R}} \{0, 1\}^\lambda$, and for $j \in [N]$, it computes $v_j = \text{CMP}(\pi^{-1}(j), y) \oplus H(F(k, j), r)$. Finally, it outputs $\text{ct} = (r, v_1, v_2, \dots, v_N)$.
- **eval** $(\text{pp}, \text{ct}_1, \text{ct}_2) \rightarrow \{0, 1\}$. The compare algorithm first parses $\text{ct}_1 = (k', h)$ and $\text{ct}_2 = (r, v_1, v_2, \dots, v_n)$, then outputs $v_h \oplus H(k', r)$.

Note that a single ciphertext from this scheme is precisely $N + 2\lambda + \lceil \log_2(N) \rceil$ bits long. For $N = 10^{12}$ and $\lambda = 80$, this amounts to ciphertexts of length 116.42 GB.⁵

Choosing the best optimizations. Our goal is to construct an ORE scheme which achieves shorter ciphertexts than the above construction, without compromising security. To do this, we use our MIFE implementation for the comparison function, and we apply our optimizations to make the ciphertext as short as possible.

We compare the ciphertext sizes for four different ORE constructions, obtained from either using the GGHLite or CLT mmap, and by applying either the DC-variant or MC-variant optimizations. For each of these options, we fix the input domain size $N = 10^{12}$ and vary the input base representation $d \in [2, 25]$. Using Equations (2) and (3), we can compute the estimated ciphertext size as a function of

⁵Clearly, increasing λ has a relatively unnoticeable effect on the overall ciphertext size for the settings of N we consider.

Figure 6.1: Estimates of the ciphertext size (in GB) for ORE with best-possible semantic security at $\lambda = 80$, for domain size $N = 10^{12}$ and for bases $d \in [2, 25]$. We compare GGHLite and CLT, with the DC-variant and MC-variant optimizations.

Figure 6.2: Estimates of the ciphertext size (in GB) for ORE with best-possible semantic security at $\lambda = 80$, for varying domain sizes. The exponent e on the x -axis denotes support for plaintexts in the range from 1 to $N = 10^e$. We compare GGHLite map (DC-variant), the CLT map, and the basic construction Π_{ore} (described in §6.2).

d (since κ is determined by the choice of d and N). See Figure 6.1 for the results. We find that, for $N = 10^{12}$, the shortest ciphertexts for ORE from GGHLite are obtained when $d = 5$ using the DC-variant optimization, and the shortest ciphertexts for ORE from CLT are obtained when $d = 6$ using the DC-variant optimization as well.

Under these settings, the DC-variant optimization for GGHLite reads the inputs in base 5, requiring $\kappa = 19$, to produce a total of 845 encodings per ciphertext, for a total size of 91.4 GB. For CLT, the DC-variant optimization reads in the inputs in base 6, requiring $\kappa = 19$, to produce a total of 974 encodings, for a total size of 6.81 GB.

We also measure the ciphertext size as we vary the domain size; see Figure 6.2. We measure the estimated ciphertext size for various domain sizes when using GGHLite, CLT, and the Π_{ore} construction described above. The results for GGHLite and CLT are using the optimal bases as detailed in Figure 6.1. We find that for $N = 10^{11}$ and $N = 10^{12}$, ORE using the CLT mmap and GGHLite mmap, respectively, produces a smaller ciphertext than Π_{ore} . This demonstrates that for certain domain sizes, our ORE construction produces the smallest known ciphertexts (versus ORE schemes that do not require mmaps).

7. PROGRAM OBFUSCATION

A *program obfuscator* [8, 24] is a compiler that aims to make a program “unintelligible” while preserving its functionality. Formally, an obfuscator $\mathcal{O}$ is a tuple of algorithms written as $\mathcal{O} = (\text{obf}, \text{eval})$, where the obfuscation algorithm obf takes as input a program P (e.g., expressed as a Boolean circuit), and outputs an obfuscated program $\text{obf}(P)$, and the evaluation algorithm takes an obfuscated program $\text{obf}(P)$ and produces an output. An obfuscator is *correct* for a program P if, for all valid inputs x accepted by P , we have that $\text{eval}(\text{obf}(P))(x) = P(x)$.

VBB and pseudo-VBB security. An obfuscator $\mathcal{O}$ is *virtual black-box (VBB) secure*⁶ for a program P if for any efficient adversary $\mathcal{A}$, there exists an efficient simulator $\mathcal{S}$, given only oracle access to $P(\cdot)$, for which the quantity

$$\left| \Pr \left[\mathcal{A}(1^\lambda, \text{obf}(P)) = 1 \right] - \Pr \left[\mathcal{S}^{P(\cdot)}(1^\lambda) = 1 \right] \right|$$

is negligible. We say that an obfuscator $\mathcal{O}$ is *pseudo-VBB secure* (pVBB) for program P and obfuscator $\mathcal{O}' = (\text{obf}', \text{eval}')$ if $\mathcal{O}'$ is both VBB secure and for every efficient adversary $\mathcal{A}$, there exists an efficient adversary $\mathcal{B}$ for which the quantity

$$\left| \Pr \left[\mathcal{A}(1^\lambda, \text{obf}(P)) = 1 \right] - \Pr \left[\mathcal{B}(1^\lambda, \text{obf}'(P)) = 1 \right] \right|$$

is negligible. In other words, if an obfuscator $\mathcal{O}$ is pVBB secure for a program P and obfuscator $\mathcal{O}'$, then any efficient attack on the security of $\mathcal{O}$ translates directly to an efficient attack on the VBB security of $\mathcal{O}'$ for the program P .

In our work, we construct a point function obfuscator that is pVBB secure for the Sahai-Zhandry obfuscator [39]. Our obfuscator operates identically to the Sahai-Zhandry obfuscator, which is VBB secure, except that we discard half of the ciphertext that corresponds to the second input in the “dual-input” branching programs that obfuscator uses. Effectively, our obfuscator can be seen as operating on “single-input” branching programs, which do not obtain VBB security, but do obtain pseudo-VBB security. We emphasize that this distinction in security is purely definitional from an attacker’s point of view, as any attack on our obfuscator immediately results in an attack on the Sahai-Zhandry obfuscator.

In this section we show how we use `cryfsm` and `libmmap` to build such a program obfuscator. Apon et al. [4] gave the first implementation of program obfuscation, using the CLT mmap [19] and a program compiler based on the approaches of Barak et al. [7] and Ananth et al. [3]. We extend this codebase in the following ways:

- **Multilinear maps.** We integrate in `libmmap` to support both the CLT and GGHLite mmmaps.
- **Program compilers.** We support MBPs output by `cryfsm`, using the Sahai-Zhandry obfuscator [39].

Point function obfuscation. We evaluated our implementation by obfuscating point functions, namely, functions that output 0 on a single (secret) input, and 1 otherwise. Previous work [4] also evaluated obfuscation for point functions, but was only able to successfully obfuscate 14-bit point functions with an mmap security parameter of $\lambda = 60$. As

⁶The reason we consider VBB versus *indistinguishability* obfuscation is that we consider point functions, for which VBB obfuscators are believed to exist.

Figure 7.1: Estimates for the ciphertext size (in GB) for point function obfuscation, for domain sizes $N = 2^{80} = 2^\lambda$ and $N = 2^{40} = 2^\lambda$. In the case of $\lambda = 80$, the minimums are achieved at $d = 19$ for GGHLite and $d = 8$ for CLT. In the case of $\lambda = 40$, the minimums are achieved at $d = 9$ for GGHLite and $d = 6$ for CLT.

noted by Bernstein et al. [9], the secret input of an n -bit point function can be recovered by simply enumerating over all 2^n possible inputs. In our experiments, we set $n = \lambda$, and consider point function obfuscation for 40-bit and 80-bit inputs.

The MBP for a λ -bit point function is of length λ and consists of a total of 2λ matrices, each of dimension 2×2 . As a small optimization, we can apply dimension reduction to obtain a branching program where the first pair of matrices need only be of dimension 1×2 . The more significant optimization comes by condensing the input representation through increasing the input base d .

The total number of encodings that we must publish in the obfuscation of a λ -bit point function can be computed as $M = 2 + 4 \cdot d \cdot \lceil \lambda / \log_2(d) \rceil$. We estimate the ciphertext size for various choices of bases in Figure 7.1, which incorporates our estimations for the size of a single encoding in GGHLite and CLT for $\lambda = 40$ and $\lambda = 80$.

- For $\lambda = 40$, we find that the minimal ciphertext size for domain size $N = 2^{40}$ is produced using MBPs under base 9 and length 13 for GGHLite, and base 22 and length 9 for CLT.
- For $\lambda = 80$, we find that the minimal ciphertext size for domain size $N = 2^{80}$ is produced using MBPs under base 19 and length 19 for GGHLite, and base 16 and length 20 for CLT.

Obfuscator implementation. Our implementation is in a mix of Python and C, with Python handling the frontend and with C handling all the computationally expensive portions, and provides interfaces to both obfuscate (`obf`) and evaluate (`eval`) an MBP. We parallelize the encoding of the elements in the MBP by using a threadpool and delegating each encoding operation to a separate thread. Once all the threads for a given matrix in the MBP complete, we then write the (encoded) matrix to disk. Thus, the threadpool approach has a higher RAM usage (due to keeping multiple encodings in memory as we parallelize) than encoding one element at a time and letting the underlying mmap library handle the parallelization, but is more efficient.

Other obfuscators. Our obfuscator is built upon improvements inspired by the Sahai-Zhandry obfuscator, which is

built on the general obfuscator described by Barak et al. [7] and Ananth et al. [3]. In addition to these obfuscators, we also implemented the Zimmerman [40] obfuscator. However, because the Zimmerman obfuscator induces a seemingly unavoidable lower bound on the degree of multilinearity for the inputs we consider, we found that the Zimmerman obfuscator was not competitive with the obfuscator we implemented. More specifically, the Zimmerman obfuscator requires that the degree of multilinearity for the obfuscation of *any* program be at least twice the number of inputs that the circuit accepts—a cost that may be insignificant when obfuscating other programs, but was too high for point functions (even when we tried to increase the input base representation to minimize this cost), and hence unsuitable for our purposes.

8. EXPERIMENTAL ANALYSIS

With one exception, all of our experiments were performed using the Google Compute Engine servers with a 32-core Intel Haswell CPU at 2.5 GHz, 208 GB RAM, and 500 GB disk storage. The exception is with the most expensive obfuscation experiment we ran (obfuscating an 80-bit point function with CLT at security parameter $\lambda = 80$), where we used a Dell PowerEdge R930 server with 4 16-core Intel Xeon CPUs at 2.5 GHz, 2 TB RAM, and 5 TB disk storage.

8.1 MIFE Experiments

We evaluated our multi-input functional encryption constructions with two applications: order-revealing encryption (ORE) (cf. §6.2) and three-input DNF (3DNF) encryption (see the full version). In §6, we showed how we can accurately estimate the ciphertext size from parameters derived from the input size and the security parameter λ , and our experiments confirmed that these parameter estimates are reasonably accurate (all within 1–2% of our reported values).

Additionally, we assessed the performance of the MIFE interface algorithms `keygen`, `encrypt`, and `eval`, along with memory utilization during the `encrypt` computation, which was by far the most costly step. We note that, since the files that we are working with are so large, a non-trivial amount of time was spent in the reading and writing of these files to disk, and so an exact reproduction of our numbers may also need to mimic the disk storage specification we use.

As another sidenote, we reiterate that our primary interest in selecting the parameters for our MIFE applications is to create the most compact ciphertexts possible given our MIFE implementation. As a result, some of our optimizations come with a cost of increased evaluation time, and hence, we believe that it is possible to reduce our evaluation time (potentially at the expense of having larger ciphertexts).

Experimental results. We summarize our experiments for MIFE in Table 8.1. We evaluated the MIFE constructions for ORE with input domain sizes $N = 10^{10}$ and $N = 10^{12}$, testing both GGHLite and CLT as the underlying mmap. For each experiment, we report the computation wall time for `encrypt` and `eval`, the overall ciphertext size `|ct|`, along with the memory usage during the `encrypt` computation. We note that for $\lambda = 80$, running `keygen` for all experiments took about 2–3 hours, and about 5–10 minutes for $\lambda = 40$. The encryption statistics measured were for generating a

λ	experiment	encrypt	eval	ct	RAM
40	GGH (10^{12})	239 m	49 s	9.7 GB	8 GB
	GGH (10^{10})	198 m	43 s	6.4 GB	7 GB
	CLT (10^{12})	35 m	81 s	1.6 GB	20 GB
	CLT (10^{10})	24 m	54 s	1.0 GB	17 GB
80	GGH (10^{12})	151 h	8 m	92 GB	78 GB
	GGH (10^{10})	109 h	6 m	61 GB	72 GB
	CLT (10^{12})	489 m	8 m	7.3 GB	192 GB
	CLT (10^{10})	312 m	5 m	4.6 GB	166 GB

Table 8.1: ORE experiments. λ denotes the security parameter of the underlying multilinear map; “experiment” denotes the multilinear map used and the domain size; “encrypt” denotes the running time of encryption; “eval” denotes the running time of evaluation, “|ct|” denotes the size of the ciphertext; and “RAM” denotes the RAM required to encrypt. We use “h” for hours, “m” for minutes, and “s” for seconds.

complete ciphertext, containing all components, as opposed to containing only the left or right (or middle) components.

Since the CLT mmap produces shorter encodings, the encryption and evaluation time for the experiments using CLT were much faster than the corresponding experiments for GGHLite. This is also partly due to the fact that CLT enjoys much more parallelism than GGHLite. However, since GGHLite has been optimized to use less memory (at the expense of increased computation time), the GGHLite experiments did not require as much RAM to run the encryption procedure.

8.2 Program Obfuscation Experiments

To evaluate our program obfuscation implementation, we chose a random secret 40-bit and a random secret 80-bit point, and used `cryfsm` to create the corresponding MBPs for the point functions associated with these points. We selected the input base representation for these programs with the goal of minimizing the total obfuscation size for each obfuscated point function (see §7 for our calculations). Like with MIFE, optimizing for obfuscation or evaluation time could lead to different optimal input base representations.

Experimental results. We tested three settings for point function obfuscation: 40-bit inputs with $\lambda = 40$, 80-bit inputs with $\lambda = 40$, and finally, 80-bit inputs with $\lambda = 80$. We also obfuscated using both CLT and GGHLite for $\lambda = 40$, but only used CLT for $\lambda = 80$, as the GGHLite experiment was too resource-intensive. Our results are summarized in Table 8.2. As we observed in the MIFE experiments, we note that GGHLite performs significantly worse when used in obfuscation compared to CLT. We also note that while obfuscation takes a huge amount of time and resources, evaluation is much less resource-intensive, for both GGHLite and CLT—a consequence of the fact that `eval` only requires multiplying (encoded) matrices, which is highly parallelizable and also much less costly than the encoding operation itself.

However, unlike our observations with the MIFE experiments, here we note that the GGHLite mmap requires more memory than when using CLT. This is because we parallelize the encoding process, whereas our MIFE construction avoids the parallelization during this step in order to reduce the amount of encodings in working memory at any given

λ	experiment	obf	eval	obf	RAM
40	GGH (40-bit)	1.8 h	6.8 s	3.6 GB	82 GB
	CLT (40-bit)	5.9 m	3.6 s	0.4 GB	2.7 GB
	GGH (80-bit)	3.0 h	29.7 s	13.7 GB	116 GB
	CLT (80-bit)	1.3 h	26.0 s	2.5 GB	20.5 GB
80	CLT (80-bit)	4.2 h	99.9 s	11.7 GB	227 GB

Table 8.2: Program obfuscation experiments. λ denotes the security parameter of the underlying multilinear map; “experiment” denotes the multilinear map used and the domain size; “obf” denotes the obfuscation time; “eval” denotes the evaluation time; “|obf|” denotes the obfuscation size; and “RAM” denotes the RAM required to obfuscate (evaluation RAM usage never exceeded 1 GB). Due to the resource requirement for $\lambda = 80$, we ran this experiment on a machine with 2 TB of RAM and 128 cores operating at 2.5 GHz. We use “h” for hours, “m” for minutes, and “s” for seconds.

stage. Since GGHLite encodings are (much) larger than CLT encodings, GGHLite uses much more memory to store these encodings. The tradeoff, of course, is an overall faster obfuscation time for all settings.

These results, while evidently impractical, are a huge improvement over prior work [4], which took 7 hours to obfuscate a 14-bit point function with $\lambda = 60$, resulting in an obfuscation of 31 GB. This improvement mainly come from (1) using a much tighter matrix branching program representation of the program, and (2) operating over different sized bases.

9. CONCLUSIONS

In this work, we presented **5Gen**, a framework for the prototyping and evaluation of applications that use multilinear maps (mmaps) and matrix branching programs. **5Gen** is built as a multi-layer software stack which offers modularity and easy integration of new constructions for each component type. Our framework offers an optimized compiler that converts programs written in the Cryptol language into matrix branching programs, a representation widely used in mmap-based constructions. **5Gen** includes a library of mmaps available through a common API; we currently support the GGHLite and CLT mmaps, but our library can be easily extended with new candidates. Leveraging the capabilities of our compiler and mmap libraries, we implemented applications from two computing paradigms based on mmaps: multi-input functional encryption (MIFE) and obfuscation.

We measured the efficiency of our MIFE and obfuscation applications with various parameter settings using both the GGHLite and CLT mmaps. While the results show efficiency that is clearly not usable in practice, they provide a useful benchmark for the current efficiency of these techniques. In the case of MIFE, we considered order-revealing encryption (ORE) and 3-input MIFE for DNF formulas that operate on triples of inputs. We chose realistic parameters for ORE for which the implemented mmap construction gives currently the best known secure construction. The same holds true for the 3-input MIFE for DNF formulas. Our experiments show that, for 80 bits of security, ORE ciphertexts over a plaintext domain of size $N = 10^{12}$ are up to 100 GB long with GGHLite and up to 10 GB long with CLT. The ORE

encryption takes hours while decryption completes in minutes.

Likewise, our obfuscation implementation achieves a huge improvement over the only previous work that implemented obfuscation of point functions, primarily due to the optimizations we chose for the matrix branching program. The size of the obfuscation of an 80-bit secret point with 80-bit security is 11.7 GB, with an obfuscation time of roughly four hours, and an evaluation time of the obfuscated program which completes in less than two minutes.

Constructing multilinear maps is an active and rapidly-evolving area of research. Our **5Gen** framework provides an easy-to-use testbed to evaluate new mmap candidates for various applications and is an open-source tool available to other research groups for experimentation.

10. ACKNOWLEDGMENTS

The work of Dan Boneh and Kevin Lewi was supported by NSF, DARPA, a grant from ONR, and the Simons Foundation. The work of Daniel Apon, Jonathan Katz, and Alex J. Malozemoff was supported in part by NSF awards #1111599 and #1223623. The work of Alex J. Malozemoff was conducted in part with Government support through the National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFG 168a, awarded by DoD, Air Force Office of Scientific Research. The work of Brent Carmer and Mariana Raykova was supported by NSF grants CNS-1633282, 1562888, 1565208, and DARPA SafeWare W911NF-15-C-0236.

This material is based upon work supported by the ARO and DARPA under Contract No. W911NF-15-C-0227. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the ARO and DARPA.

11. REFERENCES

- [1] M. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions: Cryptanalysis of some FHE and graded encoding schemes. 2016.
- [2] M. R. Albrecht, C. Cócis, F. Laguillaumie, and A. Langlois. Implementing candidate graded encoding schemes from ideal lattices. In *Asiacrypt*, 2015.
- [3] P. V. Ananth, D. Gupta, Y. Ishai, and A. Sahai. Optimizing obfuscation: Avoiding Barrington’s theorem. In *CCS*, 2014.
- [4] D. Apon, Y. Huang, J. Katz, and A. J. Malozemoff. Implementing cryptographic program obfuscation. Cryptology ePrint Archive, Report 2014/779, 2014.
- [5] B. Applebaum and Z. Brakerski. Obfuscating circuits via composite-order graded encoding. In *TCC*, 2015.
- [6] S. Badrinarayanan, E. Miles, A. Sahai, and M. Zhandry. Post-zeroizing obfuscation: The case of evasive circuits. Cryptology ePrint Archive, Report 2015/167, 2015.
- [7] B. Barak, S. Garg, Y. T. Kalai, O. Paneth, and A. Sahai. Protecting obfuscation against algebraic attacks. In *Eurocrypt*, 2014.
- [8] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *Crypto*, 2001.

- [9] D. J. Bernstein, A. Hülsing, T. Lange, and R. Niederhagen. Bad directions in cryptographic hash functions. In *ACISP*, 2015.
- [10] D. Boneh, K. Lewi, M. Raykova, A. Sahai, M. Zhandry, and J. Zimmerman. Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation. In *Eurocrypt*, 2015.
- [11] D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In *TCC*, 2011.
- [12] D. Boneh and A. Silverberg. Applications of multilinear forms to cryptography. Cryptology ePrint Archive, Report 2002/080, 2002.
- [13] D. Boneh, B. Waters, and M. Zhandry. Low overhead broadcast encryption from multilinear maps. Cryptology ePrint Archive, Report 2014/195, 2014.
- [14] D. Boneh, D. J. Wu, and J. Zimmerman. Immunizing multilinear maps against zeroizing attacks. Cryptology ePrint Archive, Report 2014/930, 2014.
- [15] Z. Brakerski and G. N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In *TCC*, 2014.
- [16] J. H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé. Cryptanalysis of the multilinear map over the integers. In *Eurocrypt*, 2015.
- [17] J. H. Cheon, J. Jeong, and C. Lee. An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without an encoding of zero. Cryptology ePrint Archive, Report 2016/139, 2016.
- [18] J.-S. Coron, C. Gentry, S. Halevi, T. Lepoint, H. K. Maji, E. Miles, M. Raykova, A. Sahai, and M. Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In *Crypto*, 2015.
- [19] J.-S. Coron, T. Lepoint, and M. Tibouchi. Practical multilinear maps over the integers. In *Crypto*, 2013.
- [20] J.-S. Coron, T. Lepoint, and M. Tibouchi. New multilinear maps over the integers. In *Crypto*, 2015.
- [21] Cryptol. <http://cryptol.net/>. Accessed: 2016-05-02.
- [22] N. Döttling, S. Garg, D. Gupta, P. Miao, and P. Mukherjee. Obfuscation from low noise multilinear maps. Cryptology ePrint Archive, Report 2016/599, 2016.
- [23] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *Eurocrypt*, 2013.
- [24] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [25] S. Garg, C. Gentry, S. Halevi, and M. Zhandry. Fully secure attribute based encryption from multilinear maps. Cryptology ePrint Archive, Report 2014/622, 2014.
- [26] S. Garg, C. Gentry, A. Sahai, and B. Waters. Witness encryption and its applications. In *STOC*, 2013.
- [27] S. Garg, P. Mukherjee, and A. Srinivasan. Obfuscation without the vulnerabilities of multilinear maps. Cryptology ePrint Archive, Report 2016/390, 2016.
- [28] C. Gentry, S. Gorbunov, and S. Halevi. Graph-induced multilinear maps from lattices. In *TCC*, 2015.
- [29] C. Gentry, A. B. Lewko, A. Sahai, and B. Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In *FOCS*, 2015.
- [30] S. Goldwasser, S. D. Gordon, V. Goyal, A. Jain, J. Katz, F.-H. Liu, A. Sahai, E. Shi, and H.-S. Zhou. Multi-input functional encryption. In *Eurocrypt*, 2014.
- [31] Y. Hu and H. Jia. Cryptanalysis of GGH map. In *Eurocrypt*, 2016.
- [32] A. Langlois, D. Stehlé, and R. Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In *Eurocrypt*, 2014.
- [33] K. Lewi and D. J. Wu. Order-revealing encryption: New constructions, applications, and lower bounds. In *CCS*, 2016.
- [34] H. Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In *Eurocrypt*, 2016.
- [35] H. Lin, R. Pass, K. Seth, and S. Telang. Indistinguishability obfuscation with non-trivial efficiency. In *PKC*, 2016.
- [36] E. Miles, A. Sahai, and M. Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In *Crypto*, 2016.
- [37] E. Miles, A. Sahai, and M. Zhandry. Secure obfuscation in a weak multilinear map model: A simple construction secure against all known attacks. Cryptology ePrint Archive, Report 2016/588, 2016.
- [38] R. Pass, K. Seth, and S. Telang. Indistinguishability obfuscation from semantically-secure multi-linear encodings. Cryptology ePrint Archive, Report 2013/781, 2013.
- [39] A. Sahai and M. Zhandry. Obfuscating low-rank matrix branching programs. Cryptology ePrint Archive, Report 2014/773, 2014.
- [40] J. Zimmerman. How to obfuscate programs directly. In *Eurocrypt*, 2015.