

COMBINE: COMpilation and Backend-INdependent vEctorization for Multi-Party Computation

Benjamin Levy¹, Muhammad Ishaq², Ben Sherman^{*3}, Lindsey Kennard^{†4}, Ana Milanova⁵,
and Vassilis Zikas⁶

^{1, 3, 4, 5}Rensselaer Polytechnic Institute (RPI), Troy, NY

^{2,6}Purdue University, West Lafayette, IN

¹levyb3@rpi.edu, ^{2,6}{ishaqm, vzikas}@purdue.edu, ³benjamin@bensherman.io,

⁴kennal@alum.rpi.edu, ⁵milanova@cs.rpi.edu

September 26, 2023

Abstract

Recent years have witnessed significant advances in programming technology for multi-party computation (MPC), bringing MPC closer to practice and wider applicability. Typical MPC programming frameworks focus on either *front-end* language design (e.g., Wysteria, Viaduct, SPDZ), or *back-end* protocol implementation (e.g., ABY, MOTION, SPDZ).

We propose a methodology for an MPC compilation toolchain, which by mimicking the compilation methodology of classical compilers enables *middle-end* (i.e., *machine-independent*) optimizations, yielding significant improvements. We advance an intermediate language, which we call *MPC-IR* that can be viewed as the analogue of (enriched) Static Single Assignment (SSA) form. MPC-IR enables backend-independent optimizations in a close analogy to machine-independent optimizations in classical compilers. To demonstrate our approach, we focus on a specific backend-independent optimization, SIMD-vectorization: We devise a novel classical-compiler-inspired automatic SIMD-vectorization on MPC-IR. To demonstrate backend independence and quality of our optimization, we evaluate our approach with two mainstream backend frameworks that support multiple types of MPC protocols, namely MOTION and MP-SPDZ, and show significant improvements across the board.

Contents

1	Introduction	3
1.1	Our Contribution	4
1.2	Outline	5
2	Overview of Methodology	5
2.1	IMP-Source as MPC Source Code	6
2.2	MPC-IR and Schedule Cost	6
2.3	Vectorized MPC-IR and Schedule Cost	6

*Work done while the author was a graduate student at RPI.

†Work done while the author was a graduate student at RPI.

3	Related Work	6
3.1	MPC Compilers	6
3.2	Intermediate representations for MPC	7
3.3	Automatic vectorization in HPC	7
4	Analytical (Parallel) Cost Model	8
4.1	Scheduling in MPC	8
4.2	(Intractability of) Optimal MPC Scheduling	8
5	Compiler Frontend	10
5.1	Overview	10
5.2	Syntax and Semantic Restrictions	11
5.3	From IMP Source to E-SSA	12
5.4	From SSA to MPC-IR	12
5.5	Base MPC-IR Syntax and Taint Types	13
6	Backend-Independent Vectorization	15
6.1	Dependence Analysis	15
6.2	Scalar and Array Expansion	15
6.3	Basic Vectorization Algorithm	16
6.4	Correctness Argument	18
6.5	Extension with Array Writes	20
7	Compiler Backends	22
7.1	Taint Analysis	22
7.2	From (Optimized) MPC-IR to MOTION	22
7.2.1	Variable Declarations	23
7.2.2	Code Generation	23
7.2.3	Vectorization and SIMD Operations	23
7.2.4	Reshaping and Raising Dimensions	24
7.2.5	Upcasting from Plaintext to Shared	24
8	Evaluation and Analysis	25
8.1	Benchmarks	25
8.2	MOTION Experiments	26
8.2.1	Experiment Setup	26
8.2.2	Results and Analysis	26
8.2.3	Analysis: Communication Size Reduction	29
8.2.4	Comparison with MOTION-native Inner Product	31
8.3	MP-SPDZ Experiments	33
8.3.1	Experiment Setup	33
8.3.2	Results and Analysis	35
9	Conclusion and Future Work	35
10	Acknowledgements	35
A	Background	39
A.1	Static Single Assignment (SSA) Form	39
A.2	Automatic Vectorization	40

1 Introduction

Multi-party computation (MPC) allows N parties P_1, \dots, P_N to perform a computation on their private inputs securely. Informally, security means that the secure computation protocol computes the correct output (correctness) and it does not leak any information about the individual party inputs beyond what can be deduced from the output (privacy).

MPC theory dates back to the early 1980s [Yao82; GMW87; BGW88; CCD88]. Long in the realm of theoretical cryptography, MPC has seen significant advances in application in recent years. New tools and compilers bring MPC closer to practice and wider applicability, e.g., [Bog+09; BG11; MZ17; MR18]. The goal is to enable programmers to write *secure* and *efficient* programs without commanding extensive knowledge of cryptographic primitives.

Recent advances in MPC programming technology tend to focus on either frontend language design (e.g., Wysteria [RHH14a], Wys* [RSH19], and Viaduct [Aca+21]) or backend circuit/protocol design and implementation (e.g., SPDZ family [KOS16; Ara+18; Kel20], MOTION [Bra+22]). The former, frontend-focused thread devised high-level constructs to express multiple parties, computation by different parties, and information flow from one party to another [RHH14a; RSH19; Aca+21]. The latter, backend-focused thread devised cryptographic protocols, typically at the circuit-level [DSZ15; Ara+18; Bra+22; Pat+21; KOS16].

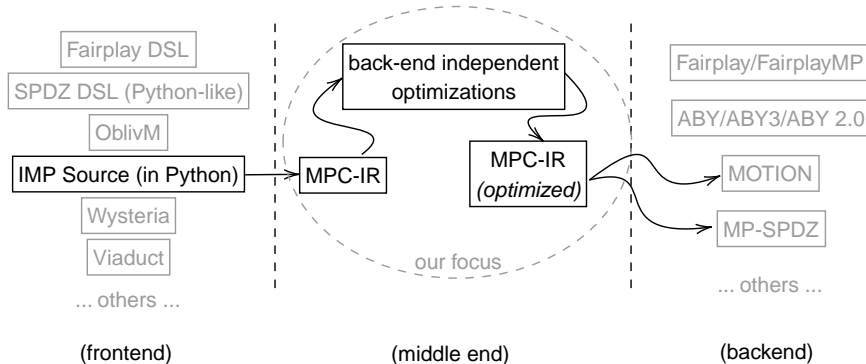


Figure 1: Our focus is middle-end of the compiler stack.

In this work we focus on the *middle end*. We formalize an intermediate representation (IR) tailored to MPC, called *MPC-IR*, and focus on what we call *backend-independent* optimizations, a close analogue to *machine-independent* optimization in the classical compiler. Fig. 1 depicts our position in the compiler stack. We envision different front ends compiling into MPC-IR — the frontend we use in this work is *IMP Source*, an easy to use Python-like language.

MPC-IR exposes the *linear structure* of MPC programs, which simplifies program analysis; this is in contrast to IMP source, which has branching constructs. At the same time, MPC-IR is sufficiently “high-level” to support analysis and optimizations that take into account control and data flow in a specific program. As an added benefit, MPC-IR facilitates simple and abstract modeling of (amortized) cost associated to different operations; this makes it suitable for defining optimizations that take advantage of amortization at the circuit level such as SIMD-vectorization and protocol mixing [Büs+18; IMZ19; Fan+22; Esc+20; Che+23]. Importantly, our formalization enables reasoning about optimizations over MPC-IR and correctness proofs of the optimizations and transformations. We demonstrate the utility of the IR and optimizations on two MPC backends, MOTION [Bra+22] and MP-SPDZ [Kel20]. We include a more detailed exposition and comparison with related work in §3.

1.1 Our Contribution

We describe an end-to-end compiler framework that takes a Python-like routine (frontend) and produces optimized backend code. More concretely we present (a) our (Python-like) IMP Source language, its syntax and semantic restrictions; (b) formal specification of MPC-IR; (c) a translation of IMP Source into MPC-IR, (d) a specific backend-independent optimization: novel SIMD-vectorization on MPC-IR; and (e) translation of MPC-IR into MOTION and MP-SPDZ code for a large set of benchmarks.

We also provide an analytical model for cost estimation of amortized schedules. Our model simplifies the problem of cost estimations by abstracting away several of the complexities. We note in passing that such cost modeling is important as it drives not only vectorization but also optimizations such as protocol mixing and scheduling [IMZ19; Fan+22; Che+23]. Given the restrictions on MPC programs that are imposed by privacy requirements—e.g., the linear program structure and bounded loops—one might think that this problem is simpler for MPC than the classical scheduling problem (which is known to be NP-hard). One might ask: *Is cost estimation for amortized schedules in MPC (in this model) tractable?* We answer this question to the negative: We show that even in our cost model, the problem remains NP-Hard; we show this via a reduction to the Shortest Common Supersequence (SCS) problem.

The above demonstrates that optimized MPC scheduling is an interesting problem. In this work, we provide an optimization by utilizing our cost model and IR and taking inspiration from the classical area of high-performance computing (HPC): a common technique there is *vectorization*, aka “SIMDification”. Informally, a Single Instruction, Multiple Data (in short, SIMD) operation works with vectors of data instead of scalars, and replaces N operations on scalars with a single operation on vectors of size N .

In more detail, we provide a novel technique for automatic vectorization of MPC-IR programs. Vectorization not only reduces running time, but also reduces communication by enabling better packing. If the backend MPC framework supports SIMD gates—most mainstream backends do, in particular the ones chosen here—this results in smaller circuits and reduces both the time and memory footprint of circuit generation. We note in passing that although much of the applied MPC literature has focused on run-time improvements, for the technology to become mainstream it is imperative to also improve the compilation time. Our work demonstrates improvements in both.

We demonstrate the wide applicability of our framework (and the expressivity of the source language) by running the compiler on 15 programs with interleaved if- and for-statements. We further demonstrate the utility and the backend-independent nature of our optimization by generating iterative and vectorized code for two well-known backend frameworks, MOTION [Bra+22] and MP-SPDZ [Kel20], taking advantage of the respective SIMD API in vectorized code; each backend then translates the code into circuits and runs the circuits. We chose these two frameworks as they are not only among the more prominent in the relevant literature, but they are also highly parameterizable as they support several MPC paradigms. They can generate and evaluate Boolean or Arithmetic circuits with GMW-style protocols [GMW87] or Boolean circuits with BMR-style protocols [BMR90]—for brevity we refer to the corresponding protocols as (Boolean/Binary or Arithmetic) GMW and BMR. This not only allows us to showcase our optimization in a wide range of benchmarks—across different MPC paradigms (see below)—but also opens our methodology to what we view as a major application area for our optimizer, namely MPC mixing.

In more detail, for MOTION we generate code for Boolean GMW and BMR—we do not include Arithmetic GMW, as MOTION does not support all operations in the Arithmetic GMW protocol, e.g., MUX, and we do not yet support protocol mixing. Towards evaluation we run experiments with two parties (2PC) and three parties (3PC). *Circuit evaluation time* for vectorized code improves over iterative code up to 30x for 2PC (resp. 40x for 3PC setting) in GMW and up to 45x for 2PC (resp. 55x for 3PC) in BMR. For the operations that do not depend on number of parties, *communication size* reduces by up to 13x in GMW and 3x in BMR. Similarly, *circuit generation time* and *number of gates* reduce, respectively, by up to 200x and 480x in GMW, and 80x and 450x in BMR.

For MP-SPDZ we generate Python code then compile and execute in the Arithmetic setting and in the Binary setting. We run experiments with 2PC and demonstrate reduction in compilation time of up to 25x in the Binary setting.

Full source code and benchmarks are available at <https://github.com/milana2/>

```

1 def biometric(C: shared[list[int]], D: int,
2   S: shared[list[int]], N: int) ->
3   shared[tuple[int,int]]:
4   min_sum : int = MAX_INT
5   min_idx : int = 0
6   for i in range(N):
7     sum : int = 0
8     for j in range(D):
9       # d = S[i,j] - C[j]
10      d : int = S[i * D + j] - C[j]
11      p : int = d * d
12      sum = sum + p
13      if sum < min_sum:
14        min_sum : int = sum
15        min_idx : int = i
16  return (min_sum, min_idx)

```

(a) IMP Source

```

1 min_sum!1 = MAX_INT
2 min_idx!1 = 0
3 for i in range(0, N):
4   min_sum!2 = PHI(min_sum!1, min_sum!4)
5   min_idx!2 = PHI(min_idx!1, min_idx!4)
6   sum!2 = 0
7   for j in range(0, D):
8     sum!3 = PHI(sum!2, sum!4)
9     d = SUB(S[(i * D) + j], C[j])
10    p = MUL(d,d)
11    sum!4 = ADD(sum!3,p)
12    t = CMP(sum!3,min_sum!2)
13    min_sum!3 = sum!3
14    min_idx!3 = i
15    min_sum!4 = MUX(t, min_sum!3, min_sum!2)
16    min_idx!4 = MUX(t, min_idx!3, min_idx!2)
17  return (min_sum!2, min_idx!2)

```

(b) MPC-IR

```

1 min_sum!1 = MAX_INT
2 min_idx!1 = 0
3 # S^ is same as S. C^ replicates C N times:
4 S^ = raise_dim(S, ((i * D) + j), (i:N,j:D)) #S^ [i,j] = S[i,j]
5 C^ = raise_dim(C, j, (i:N,j:D)) #C^ [i,j] = C[j]
6
7 sum!2[l] = [0,...,0]
8 # computes _all_ "at once"
9 d[l,J] = SUB_SIMD(S^ [l,J], C^ [l,J])
10 p[l,J] = MUL_SIMD(d[l,J], d[l,J])
11
12 for j in range(0, D):
13   # sum!2[l], sum!3[l], sum!4[l] are size-N vectors
14   # computes N intermediate sums "at once"
15   sum!3[l] = PHI(sum!2[l], sum!4[l])
16   sum!4[l] = ADD_SIMD(sum!3[l], p[l,j])
17
18 min_idx!3[l] = [0,1,...,N-1]
19 for i in range(0, N):
20   min_sum!2 = PHI(min_sum!1, min_sum!4)
21   t[i] = CMP(sum!3[i], min_sum!2)
22   min_sum!4 = MUX(t[i], sum!3[i], min_sum!2)
23   for i in range(0, N):
24     min_idx!2 = PHI(min_idx!1, min_idx!4)
25     min_idx!4 = MUX(t[i], min_idx!3[i], min_idx!2)
26   return (min_sum!2, min_idx!2)

```

(c) Optimized MPC-IR

Table 1: Biometric Matching from IMP Source to Optimized MPC-IR. – MPC-IR is an SSA form without conditionals, therefore conditional on lines 13-15 in (a) turns into linear code on lines 12-16 (b). – In (c), our compiler fully vectorizes the SUB and MUL operations on lines 9 and 10 of (b). The computation of sum (line 11 in (b)) is sequential across the j -dimension, but it is parallel across the i -dimension as the loop on lines 12-16 in (c) illustrates; here $p[l,j]$ refers to the j -th column in p .

[ParallelizationForMPC](https://github.com/milana2/ParallelizationForMPC/tree/gh-pages) where <https://github.com/milana2/ParallelizationForMPC/tree/gh-pages> shows experiments on small input sizes and details the different stages of the compiler.

Our results emphasize the opportunities of backend-independent optimizations; we believe that our work can lead to future work on backend-independent compilation and optimization, ushering new MPC optimizations and combinations of optimizations.

1.2 Outline

The rest of the paper is organized as follows. §2 presents an overview of our techniques. §3 reviews the related work. §4 describes the cost estimation model and argues intractability of optimal scheduling. §5 details the frontend of the compiler, §6 focuses on MPC-IR and backend-independent vectorization, and §7 briefly describes translation into MOTION and MP-SPDZ. §8 presents the experiments and §9 concludes the paper.

2 Overview of Methodology

To demonstrate our main technical contributions, we first provide a high-level overview of our methodology, using the standard MPC benchmark of Biometric matching as our running example. For interested readers unfamiliar with the SSA form or automatic vectorization we provide an overview in the appendix to this paper §A.

2.1 IMP-Source as MPC Source Code

An intuitive (and naive) implementation of Biometric matching is as shown in Listing 1(a). Array C is the feature vector that we wish to match and S is the database of N size- D vectors that we match against.

Our compiler takes essentially standard IMP [NK14] syntax and imposes certain semantic restrictions (details will follow). The programmer writes an iterative program and annotates certain inputs and outputs as *shared*. In the example, arrays C and S are **shared**, meaning that they store shares (secrets), however, the array sizes D and N respectively are plaintext. The code iterates over the S and computes the Euclidean distance of the current entry $S[i]$ and C (its square actually). The program returns the index of the vector that gives the best match and the corresponding sum of squares.

2.2 MPC-IR and Schedule Cost

Our compiler generates MPC-IR, a *linear* Static Single Assignment (SSA) form. Listing 1(b) shows the MPC-IR translation of the code in 1(a).

We turn to our analytical model to compute the *cost* of the iterative program. Assume cost β for a local MPC operation (e.g., XOR in Boolean sharing or ADD in Arithmetic sharing) and cost α for a remote MPC operation (e.g., MUX, CMP, etc.). Assuming that ADD is β and SUB, CMP and MUX are α , the MPC-IR in Listing 1(b) gives rise to an iterative schedule with cost $ND(2\alpha + \beta) + N(3\alpha)$.

2.3 Vectorized MPC-IR and Schedule Cost

We can compute all $N * D$ subtraction operations at line 9 in 1(b) in a single SIMD instruction; similarly we can compute all multiplication operations at line 10 in a single SIMD instruction. Our compiler runs Listing 1(b) through the vectorization optimization to produce 1(c). Note that this is still our IR, Optimized MPC-IR. The compiler turns this code into variables, loops and SIMD primitives (if supported), suitable for the backend to generate the circuit.

In MPC backends, executing n operations “at once” in a single SIMD operation costs less than executing those n operations one by one. This is particularly important for interactive gates, since it allows many 1-bit values to be sent at once. We consider that each operation has a *fixed* portion that benefits from amortization and a *variable* portion that does not benefit from amortization: $\alpha = \alpha_{fix} + \alpha_{var}$. This gives rise to the following formula for amortized cost: $f(n) = \alpha_{fix} + n\alpha_{var}$, as opposed to unamortized cost $g(n) = n\alpha_{fix} + n\alpha_{var}$. (We extend the same reasoning to β -instructions.)

Thus, the fixed cost of the vectorized program amounts to $2\alpha_{fix} + D\beta_{fix} + N(3\alpha_{fix})$. The variable cost is the same in both the vectorized and non-vectorized programs. The first term in the sum corresponds to the vectorized subtraction and multiplication (lines 9-10 in (c)), the second term corresponds to the for-loop on j (lines 12-16) and the third one corresponds to the remaining for-loops on i (lines 19-25). Clearly, $2\alpha_{fix} + D\beta_{fix} + N(3\alpha_{fix}) \ll ND(2\alpha_{fix} + \beta_{fix}) + N3\alpha_{fix}$. Empirically, we observe orders of magnitude improvement e.g., for Biometric Matching evaluation time, 10x and 23x in GMW and BMR respectively in 2PC, and 12x and 28x in 3PC. Additionally, the un-vectorized version runs out of memory for $N = 256$, while the vectorized one runs with the standard maximal input size $N = 4,096$.

3 Related Work

3.1 MPC Compilers

The early MPC compilers—Fairplay [BNP08] and Sharemind [BLW08]—demonstrated that MPC can be brought to applications. This led to the development of a diverse landscape of implementations/compiler for specific MPC protocols, e.g., PICCO [ZSB13], Obliv-C [ZE15], TinyGarble [Son+15], Wystiria [RHH14a], Frigate [Moo+16], SPDZ/SPDZ-2 [Dam+12; Dam+13; KSS13], SCALE-MAMBA [COS19] and others. A new generation of state-of-the art (and actively developed) MPC backends includes MP-SPDZ [Kel20] and

the ABY/HyCC/MOTION [DSZ15; B us+18; Bra+22] frameworks. These frameworks are highly parameterizable and allow the use of different MPC protocols/paradigms. Another recent development is Viaduct, a language and compiler that supports a range of secure computation frameworks, including MPC and ZKP. For a (by now slightly outdated) review of MPC compiler frameworks please see [Has+19]. In contrast to the above works, we focus on backend-independent optimizations available at the “higher-level” MPC-IR. Our key goal is to demonstrate provably correct compiler optimization for MPC, not a language frontend. Indeed, replacing our frontend with a more functional frontend, such as Viaduct’s, is an interesting future direction.

Obliv-C [ZE15], Wysteria [RHH14b] and Viaduct [Aca+21] focus on higher-level language design. OblivM [Liu+15] has similar goals to ours but our works are complementary in the sense that while OblivM relies on programmer annotations such as map-reduce constructs, we automatically detect opportunities for optimization at an intermediate level of representation. Similarly, GraphCG [Nay+15] introduces high-level programming abstractions and novel parallel oblivious algorithms; programmers can make use of the abstractions to write highly parallel and efficient secure algorithms.

HyCC [B us+18] is a compiler from C Source into ABY circuits. It does source-to-source compilation with the goal to decompose the program into modules and then assign protocols to modules. In contrast, we focus on MPC-IR-level optimizations, specifically vectorization, although we envision future optimizations as well. HyCC, similarly to Buscher et al. [BK15] and [B us18], uses an off-the-shelf source-to-source polyhedral compiler ¹ to perform vectorization at the level of source code. The disadvantage of using an off-the-shelf source-to-source compiler is that it solves a more general problem than what MPC presents and may forgo optimization because polyhedral compilation does not work well with conditionals (see [Ben+10]). In addition, the MPC-IR produced by our vectorization algorithm can serve as input to protocol mixing algorithms, such as, for example, OPA [IMZ19], which requires vectorized input and appears to use ad-hoc and manual vectorization.

3.2 Intermediate representations for MPC

Recent work makes progress on intermediate representations for MPC. Ozdemir et al. [OBW20] develop CirC, an IR with backends into ZKP and SMT primitives. MPC-IR is a higher level of representation than CirC; specifically, it does not unroll loops, leading to more scalable analysis. Heldmann et al. [Hel+21] present an LLVM-IR based toolchain for compilation into circuits. We take a different approach — rather than reverse engineer LLVM-IR, which is rich and complex, we propagate necessary information from the frontend to the middle-end IR; in addition, we formalize MPC-IR to enable reasoning about correctness of transformations. Another recent IR, developed concurrently and independently from ours, is FUSE [Bra+23]. It enables optimizations such as vectorization, in a way similar to the Section 6.1 optimization in MP-SPDZ [Kel20] (to our understanding). MPC-IR is a higher level IR and is orthogonal to FUSE. One can compile MPC-IR into FUSE and take advantage of the optimizations available at this level; our results with MP-SPDZ demonstrate that the two IRs and optimizations complement each other. In addition, we present syntax and semantics of MPC-IR, which enables reasoning about code transformations, a feature not present in [Bra+23; Hel+21].

3.3 Automatic vectorization in HPC

Automatic vectorization is a longstanding problem in high-performance computing (HPC). We present a vectorization algorithm for MPC-IR that builds upon classical loop vectorization by Allen and Kennedy [AK87]. We view Karrenberg’s work on Whole function vectorization [Kar15] as most closely related to our work — it linearizes the program and vectorizes both branches of a conditional applying masking to avoid execution of the branch-not-taken code, and selection (similar to MUX). We argue that vectorization over linear MPC-IR is a problem that warrants a new look, while drawing from results in HPC. Since both branches of the conditional and the multiplexer *always* execute, not only can we apply aggressive vectorization on linear code,

¹To our understanding, HyCC uses Par4All (<https://github.com/Par4All/par4all>), however, does not appear to be included with HyCC’s publicly available distribution.

but (perhaps more importantly) we can also build analytical models that accurately predict execution time. The models can drive optimizations such as vectorization and protocol mixing; the optimizations interact in non-trivial ways.

Polyhedral parallelization [Ben+10] is another rich area. It considers a higher-level source (typically AST) representation, while, in contrast, our work takes advantage of linear MPC-IR and SSA, and the corresponding dependence analysis. Karrenber’s work [Kar15] is rare in that space, in the sense that it considers vectorization over SSA, which has similarities with analysis over MPC-IR. We differ in the array representation, notion of dependence, and reasoning about dependence, as we specifically target MPC-IR.

4 Analytical (Parallel) Cost Model

Next, we introduce our model for cost estimation of the MPC schedules and prove that optimal schedule (of MPC) is NP-Hard.

4.1 Scheduling in MPC

We work in a single protocol setting i.e., all MPC tasks are evaluated in a single protocol from start to finish. In addition, we abstract common features of MPC execution, in the following assumptions:

- (1) There are two types of MPC instructions, local and remote. A local instruction (i.e., ADD or XOR) has cost β and a remote instruction (i.e. MUL) has cost α , where $\alpha \gg \beta$. We assume that all remote instructions have the same cost α and all local ones have the same cost β .
- (2) In MPC frameworks, executing n operations “at once” in a single SIMD operation costs a lot less than executing those n operations one by one. Following Amdahl’s law, we write $\alpha = \frac{1}{s}p\alpha + (1-p)\alpha$, where p is the fraction of execution time that benefits from amortization and $(1-p)$ is the fraction that does not, and s is the available resource. Thus, $n\alpha = \frac{n}{s}p\alpha + n(1-p)\alpha$. For the purpose of the model we assume that s is large enough and the term $\frac{n}{s}p\alpha$ amounts to a *fixed cost* incurred regardless of whether n is 10,000 or just 1. (This models the cost of preparing and sending a packet from party A to party B for example.) Therefore, amortized execution of n operations is $f(n) = \alpha_{fix} + n\alpha_{var}$ in contrast to unamortized execution $g(n) = n\alpha_{fix} + n\alpha_{var}$. We have $\alpha_{fix} \ll n\alpha_{fix}$ and since fixed cost dominates variable cost (particularly for remote operations), we have $f(n) \ll g(n)$.
- (3) MPC instructions scheduled in parallel benefit from amortization *only if* they are the same instruction. Given our previous assumption, 2 MUL instructions can be amortized in a single SIMD instruction that costs $\alpha_{fix} + 2\alpha_{var}$, however a MUL and a MUX instruction still cost $2\alpha_{fix} + 2\alpha_{var}$ even when scheduled “in parallel”.²

4.2 (Intractability of) Optimal MPC Scheduling

Given a serial schedule (a linear graph) of an MPC program i.e. a sequence of instructions $S := (S_1; \dots; S_n)$, where S_i is an instruction, and a def-use dependency graph $G(V, E)$ corresponding to S , our task is to construct a parallel schedule (another linear graph) $P := (P_1; \dots; P_m)$ observing the following conditions:

- (1) All P_i ’s consist of instructions of the same kind.
- (2) Def-use dependencies of the graph $G(V, E)$ are preserved i.e. if instructions $S_i, S_j, i < j$ form a def-use i.e. an edge exists from S_i to S_j in G , then they can only be mapped to $P_{i'}, P_{j'}$ such that $i' < j'$.

Correctness of P follows due to the preservation of def-use *dependencies*. One can easily argue by induction on the length of schedule S that the computed function is the same in both S and P .

The cost of schedule S is

$$cost(S) = \sum_{i=1}^n cost(S_i) = L_\alpha \alpha_{fix} + L_\beta \beta_{fix} + L_\alpha \alpha_{var} + L_\beta \beta_{var} \tag{1}$$

²This is not strictly true, but assuming it, e.g. as in [IMZ19; DSZ15; MR18], helps simplify the problem.

where L_α is the number of α -instructions and L_β is the number of β ones. (We used this formula to compute the cost of the unrolled MPC Source program in §2.) The cost of schedule P is more interesting:

$$\text{cost}(P) = \sum_{i=1}^m \text{cost}(P_i) \quad (2)$$

Each P_i may contain multiple instructions, and $\text{cost}(P_i)$ is amortized. Thus, according to our model $\text{cost}(P_i) = \alpha_{\text{fix}} + |P_i|\alpha_{\text{var}}$ if P_i stores $|P_i|$ α -instructions, or $\text{cost}(P_i) = \beta_{\text{fix}} + |P_i|\beta_{\text{var}}$ if it stores β -instructions. (Similarly, we used this formula to compute the cost of the Optimized MPC Source program in §2.)

Our goal is to construct a parallel schedule P that reduces the program cost (when compared to cost of S). One would hope that simpler MPC program structure would make optimal schedule tractable. Intuitively, the problem is to combine multiple independent schedules (or sequences of instructions) into a single schedule where same instructions are scheduled into a SIMD-instruction P_i . This amounts to finding a Shortest Common Supersequence for the independent schedules. We formalize the argument below and show that the scheduling problem is NP-hard via a reduction to the Shortest Common Supersequence problem [Vaz10].

Proof. To prove that optimal scheduling is an NP-Hard problem, we consider the following convenient representation. An MPC program is represented as a set of sequences $\{s_1, \dots, s_n\}$ of operations. In each sequence s_i operations depend on previous operations via a def-use i.e. $s_i[j], j > 1$ depends on $s_i[j - 1]$.

As an example, consider the MPC program consisting of the following three sequences, all made up of two distinct α -instructions M_1 and M_2 , e.g., M_1 is MUL and M_2 is MUX. The right arrow indicates a def-use *dependence*, meaning that the source node must execute before the target node:

1. $M_1 \rightarrow M_2 \rightarrow M_1$
2. $M_1 \rightarrow M_1 \rightarrow M_1$
3. $M_2 \rightarrow M_1 \rightarrow M_2$

The problem is to find a schedule P with *minimal cost*. For example, one such schedule for the sequences above is

$$M_1(1)||M_1(2); M_1(2); M_2(1)||M_2(3); M_1(1)||M_1(2)||M_1(3); M_2(3)$$

The parentheses above indicate the sequence where the instruction comes from: (1), (2), or (3). Cost of schedule P is computed using Eq. (2) and it amounts to $5\alpha_{\text{fix}} + 9\alpha_{\text{var}}$.

The problem of finding a schedule P with a minimal $\text{cost}(P)$ is shown to be NP-Hard problem, as it can be reduced to the problem of finding a *shortest common supersequence*, a known NP-Hard problem [Vaz10]. The shortest common supersequence problem is as follows: *given two or more sequences find the the shortest sequence that contains all of the original sequences*. This can be solved in $O(n^k)$ time, where n is the cardinality of the longest sequence and k is the number of sequences. We can see that the optimal schedule is the shortest schedule, since the shortest schedule minimizes the fixed cost while the variable cost remains the same.

To formalize the reduction, suppose P is a schedule with minimal cost (computed by a black-box algorithm). Clearly P is a supersequence of each sequence s_i , i.e., P is a common supersequence of $s_1 \dots s_n$. It is also a shortest common supersequence. The cost of $\text{cost}(P) = L\alpha_{\text{fix}} + N\alpha_{\text{var}}$ where L is the length of P and N is the total number of instructions across all sequences. Now suppose, there exist a shorter common supersequence P' of length L' . $\text{cost}(P') < \text{cost}(P)$ since $L'\alpha_{\text{var}} + N\alpha_{\text{var}} < L\alpha_{\text{var}} + N\alpha_{\text{var}}$, contradicting the assumption that P has the lowest cost. \square

Finally, we remark on the (necessity of) assumptions. The assumption that “there are only two types of instructions, α and β ” comes in the proof only, and it comes w.l.o.g. as to prove NP-hardness for an arbitrary circuit (that uses higher level gates), one can first reduce this new circuit to another one that only uses MUL and ADD gates and then use our proof. Since the reduction is polynomial time, this proof will imply NP-hardness of optimal scheduling for the new (lower level) circuit. We stress that this assumption is for the




Figure 2: Compiler Framework.

NP-hardness proof only. The vectorization optimization does not take into account costs—it vectorizes every operation it can—so there is no negative impact there. Notwithstanding, since this assumption is w.l.o.g., for optimizations that actually use the cost tables, e.g., mixing, one can use a refined model and it will not affect the proof.

The assumption that “only same-type instructions benefit from vectorization” is a simplifying assumption; it is not strictly true, but assuming it, as in e.g., [IMZ19; DSZ15; MR18], helps simplify the problem. One might be able to get additional improvements by identifying parallelizable heterogeneous instructions, but it is unclear how to do this in an automated compiler without unrolling loops (unrolling loops would go a level of abstraction below MPC-IR and would slow down the compilation process and increase the memory needed by the compiler to store the unrolled loop information).

5 Compiler Frontend

This section presents an overview of our compiler, followed by our source syntax and semantic restrictions.

5.1 Overview

Fig. 2 shows the phases of our compiler. We present a systematic translation of high-level source, which includes if-then-else statements, into a linear sequence of primitive MPC instructions. We start with an Abstract Syntax Tree (AST) source syntax, then convert it into a Control-flow graph (CFG), as Cytron’s classical SSA algorithm [Cyt+91] is defined over a CFG. SSA is the natural means for converting if-then-else statements into MUX primitives, as there is correspondence between ϕ -nodes and MUX primitives. Our translation takes advantage of the properties of Cytron’s SSA, particularly a *minimal number*³ of phi-nodes that results in a minimal number of MUX primitives. We then translate into MPC-IR, which is most conveniently described with an AST syntax (cf. §5.5 and §6.4).

We also note that we settled on this process after the straightforward approach failed. Initially, we attempted to reuse existing implementations of SSA intermediate representations such as Soot Shimple (<http://soot-oss.github.io/soot/>) or LLVM IR (<https://llvm.org/>). The problem was that these were CFG representations and they had lost connection between the ϕ -node and the conditional that triggered the ϕ -node. Specifically, given $x = \phi(y, z)$, there is no information what conditional triggered the ϕ -node and whether y corresponds to the true or false branch of the conditional. Moreover, ϕ -nodes with 3 or more arguments are common in standard SSA. Furthermore, existing SSA IRs do not handle arrays, while handling of arrays is important for vectorization. As a remark, while it is possible to reconstruct the missing information via a form of control dependence analysis, and it is possible to add arrays, this proves very difficult due to the complexity of the IRs (these IRs are designed to handle richer and more complex syntax than MPC). A clean state solution, where we start from the AST and retain all necessary information in the CFG i.e. enhanced-CFG (or E-CFG in Fig. 2) and construct enhanced-SSA (or E-SSA) that handles arrays,

³A naive SSA translation will place a phi-node for every variable at every control merge node, but Cytron’s SSA places phi-nodes only where they are needed, at the Dominance Frontier of each variable assignment.

proves the correct choice and drives our progress on the compiler. The E-SSA form naturally gives rise to MPC-IR, where conditional ϕ -nodes translate to MUX nodes, and loop ϕ -nodes translate into what we call pseudo PHI-nodes (shown in lines 4, 5 and 8 in Listing 1(b)).

In the remainder of this section we first describe the source syntax and semantic restrictions (in §5.2), we then outline the source-to-SSA translation (in §5.3) followed by the SSA-to-MPC-IR translation (in §5.4). §5.5 describes MPC-IR syntax and the simple taint type system we define on top of this syntax.

5.2 Syntax and Semantic Restrictions

Source syntax is standard IMP syntax but with for-loops:

$e ::= e \text{ op } e \mid x \mid \text{const} \mid \mathbf{A}[e]$	<i>expression</i>
$s ::= s; s \mid$	<i>sequence</i>
$x = e \mid \mathbf{A}[e] = e \mid$	<i>assignment stmt</i>
for i in $\text{range}(I) : s \mid$	<i>for stmt</i>
if $e : s$ else : s	<i>if stmt</i>

The syntax allows for array accesses, arbitrarily nested loops, and if-then-else control flow.

The prototype version of the compiler assumes the following semantic restrictions on source programs. Currently, the compiler does not enforce the restrictions, however, they can be easily encoded as rules in a syntax-directed translation [ASU86; Sco09] over the syntax above. The reason why we do not implement the rules is because the majority of these restrictions are implementation restrictions that can be lifted in future versions of our compiler.

- Loops are of the form $0 \leq i < I$ and bounds are fixed at compile time. It is a standard restriction in MPC that the bounds must be known at circuit-generation time.
- Arrays are one-dimensional. N -dimensional arrays are linearized and accessed in row-major order.
- Array subscripts are plaintext. We will lift this restriction in future work by applying the standard linear scan [Liu+15; Ara+18] when the subscript is a secret-shared value.
- The subscript e is a function of the indices of the enclosing loops. For read access, the compiler allows an arbitrary such function. However, it restricts write access to *canonical writes*, i.e., $\mathbf{A}[i, j, k] = \dots$ where i, j and k are the indices of the *outermost* loops enclosing the array write statement. These indices loop over the three dimensions of \mathbf{A} and all write accesses to \mathbf{A} follow this restriction. We note that this is a restriction on the vectorization optimization; there is no reason to restrict arbitrary writes in the code, they just are not optimized. We state this restriction upfront as it simplifies vectorization and its exposition.
- The final restriction disallows array writes from within if-then-else statements. This is to ensure that arguments of MUX in the MPC-IR translation are base types, i.e., just int or bool. In our experience, this causes a minor inconvenience to the programmer as they may not write

```
1 if e: A[i] = val
```

Instead they write

```
1 if not(e): val = A[i]
2 A[i] = val
```

In addition, our compiler defines and implements a taint type system at the level of MPC-IR. We define the base MPC-IR syntax and the type system in §5.5. We note that while the programmer writes annotations at the level of IMP Source (as in Listing 1(a)), the annotations propagate through the transformations; annotations are inferred and checked with a standard taint analysis (based on the type system) at the level of MPC-IR. The only required annotations are on the input arguments.

For the rest of this paper we write i, j, k to denote the loop nesting: i is the outermost loop, j , is immediately nested in i , and so on until k and we use I, J, K to denote the corresponding upper bounds. We write $\mathbf{A}[i, j, k]$ to denote *canonical access* to an element, either array element or a scalar expanded to its loop nest i, j, k . To simplify the presentation we describe our algorithms in terms of three-element tuples i, j, k , however, discussion generalizes to arbitrarily large loop nests.

5.3 From IMP Source to E-SSA

Our compiler translates from Source to E-SSA as follows:

Parsing: Use Python’s `ast` module to parse the input source code to a Python AST.

Syntax checking: Ensure that the AST matches the restricted subset defined in §5.2. This step outputs an instance of the `restricted_ast.Function` class, which represents our restricted subset of the Python AST.

3-address E-CFG conversion: Convert the restricted-syntax AST to a three-address enhanced control-flow graph. To do this, first, add an empty basic block to the CFG and mark it as current. Next, for each statement in the restricted AST’s function body, process the statement. Statements can either be for-loops, if-statements, or assignments (as in §5.2). Rules for processing each kind of statement are given below:

- **For-loops:** Create new basic blocks for the loop condition (the *condition-block*), the loop body (the *body-block*), and the code after the loop (the *after-block*). Insert a jump from the end of the current block to the condition-block. Then, mark the condition-block as the current block. Insert a for-instruction at the end of the current block with the loop counter variable and bounds from the AST. Next, add an edge from the current block to the after-block labeled “FALSE” and an edge from the current block to the body-block labeled “TRUE”. Then, set the body-block to be the current block and process all statements in the AST’s loop body. Finally, insert a jump to the condition-block and set the after-block as current.
- **If-statements:** Create new basic blocks for the “then” statements of the if-statement (the *then-block*), the “else” statements of the if-statement (the *else-block*), and the code after the if-statement (the *after-block*). At the end of the current block, insert a conditional jump to the then-block or else-block depending on the if-statement condition in the AST. Next, mark the then-block as current, process all then-statements, and add a jump to the after-block. Similarly, mark the else-block as current, process all else-statements, and add a jump to the after-block. Finally, set the after-block to be the current block, and give it a *merge condition* property equal to the condition of the if-statement.
- **Assignments:** In the restricted-syntax AST, the left-hand side of assignments can be a variable or an array subscript. If it is an array subscript, e.g., $A[i] = x$, change the statement to $A = \text{Update}(A, i, x)$. If the statement is not already three-address code, for each sub-expression in the right-hand side of the assignment, insert an assignment to a temporary variable.

SSA conversion: Convert the 3-address CFG to SSA with Cytron’s algorithm.

5.4 From SSA to MPC-IR

Once the compiler converts the code to SSA, it transforms ϕ -nodes that correspond to if-statements into MUX nodes. From the 3-address CFG conversion step, ϕ -nodes corresponding to if-statements will be in a basic block with the merge condition property. For example, if $X!3 = \phi(X!1, X!2)$ is in a block with merge condition C , the compiler transforms it into $X!3 = \text{MUX}(C, X!1, X!2)$. Next, the compiler runs the dead code elimination algorithm from Cytron’s SSA paper.

Next, the control-flow graph is *linearized* into MPC-IR, which has loops but no if-then-else-statements. This means that both branches of all if-statements are executed, and the MUX nodes determine whether to use results from the then-block or from the else-block. The compiler linearizes the control-flow graph with a variation of depth-first search. Blocks with the “merge condition” property are only considered the second time they are visited, since that will be after both branches of the if-statement are visited. (The Python AST naturally gives rise to a translation where each conditional has exactly two targets, and each “merge condition” block has exactly two incoming edges, a TRUE and a FALSE edge. Thus, each ϕ -node has exactly two multiplexer arguments, which dovetails into MUX. This is in contrast with Cytron’s algorithm which

operates at the level of the CFG and allows for ϕ -nodes with multiple arguments.) Each time the compiler visits a block, it adds the block's statements to the MPC-IR. If the block ends in a for-instruction, the compiler recursively converts the body and code after the loop to MPC-IR and adds the for-loop and code after the loop to the main MPC-IR. If the block does not end in a for-instruction, the compiler recursively converts all successor branches to MPC-IR and appends these to the main MPC-IR.

```

{ Step 1: Replace  $\phi$ -nodes with MUX nodes }
for each basic block block in the control-flow graph do
  if block has the merge condition property then
    merge_cond  $\leftarrow$  merge condition variable of block
    for each  $\phi$ -node phi =  $\phi(v_1, v_2)$  in block do
      Replace phi with MUX(merge_cond,  $v_1$ ,  $v_2$ ) in block
    end for
  end if
end for
{ Step 2: Linearize E-CFG into MPC-IR }
visited  $\leftarrow$  empty set
merge_visited  $\leftarrow$  empty set
Define search(block):
  if block has the merge condition property and block is not in merge_visited then
    Add block to merge_visited
    return empty list
  end if
  Add block to visited
  if block is a for-loop header then
    cfg_body  $\leftarrow$  successor of block containing the beginning of the for-loop body
    cfg_after  $\leftarrow$  successor of block containing the beginning of the code after the for-loop
    mpc_body  $\leftarrow$  list of the  $\phi$ -functions in block concatenated with search(cfg_body)
    loop  $\leftarrow$  MPC-IR for-loop statement with the same counter variable and bounds as block and with
    mpc_body as its body
    return loop prepended to search(cfg_after)
  else
    result  $\leftarrow$  empty list
    for each successor successor of block do
      if successor is not in visited then
        result  $\leftarrow$  result concatenated with search(successor)
      end if
    end for
    return result
  end if
return search(entry block of the control-flow graph)

```

Now, the remaining ϕ -nodes in MPC-IR are the loop header nodes. These are the *pseudo* ϕ -nodes and we write PHI in MPC-IR. A pseudo PHI-node $x!1 = \text{PHI}(x!0, x!2)$ in a loop header is evaluated during circuit generation. If it is the 0-th iteration, then the PHI-node evaluates to $x!0$, otherwise, it evaluates to $x!2$.

5.5 Base MPC-IR Syntax and Taint Types

The syntax of the MPC-IR program produced by the above section is essentially IMP syntax. (In §6 we extended the base syntax to account for vectorization.) Most notably, there is no if-then-else statement but

there are MUX expressions:

$$\begin{array}{ll}
e ::= e \text{ op } e \mid x \mid \text{const} \mid \mathbf{A}[e] \mid \text{MUX}(e, e, e) & \text{expression} \\
s ::= s; s \mid & \text{sequence} \\
x = \text{PHI}(x, x) \mid x = e \mid \mathbf{A}[e] = e \mid & \text{assignment stmt} \\
\text{for } i \text{ in range}(I) : s & \text{for stmt}
\end{array}$$

Expressions are typed $\langle q \ \tau \rangle$, where q and τ are as follows:

$$\begin{array}{ll}
\tau ::= \text{int} \mid \text{bool} \mid \text{list}[\text{int}] \mid \text{list}[\text{bool}] & \text{base types} \\
q ::= \text{shared} \mid \text{plain} & \text{qualifiers}
\end{array}$$

The type system is standard, and in our experience, a sweet spot between readability and expressivity. The **shared** qualifier denotes shared values, i.e., ones shared among the parties and computed upon under secure computation protocols; the **plain** qualifier denotes plaintext values. Subtyping is **plain** $<$: **shared**, meaning that we can convert a plaintext value into a shared one, but not vice versa. Subtyping on qualified types is again as expected, it is covariant in the qualifier and invariant in the type: $\langle q_1 \ \tau_1 \rangle <: \langle q_2 \ \tau_2 \rangle$ iff $q_1 <: q_2$ and $\tau_1 = \tau_2$.

The typing rules for non-trivial expressions are as follows:

$$\begin{array}{c}
\frac{\Gamma \vdash e_1 : \langle q_1 \ \tau \rangle \quad \Gamma \vdash e_2 : \langle q_2 \ \tau \rangle \quad \tau \in \{\text{int}, \text{bool}\}}{\Gamma \vdash e_1 \text{ op } e_2 : \langle q_1 \vee q_2 \ \tau \rangle} \text{(BINARY OP)} \\
\frac{\Gamma \vdash e : \langle \text{plain int} \rangle \quad \Gamma \vdash \mathbf{A} : \langle q \ \text{list}[\tau] \rangle \quad \tau \in \{\text{int}, \text{bool}\}}{\Gamma \vdash \mathbf{A}[e] : \langle q \ \tau \rangle} \text{(ARRAY ACCESS)} \\
\frac{\Gamma \vdash e_1 : \langle q_1 \ \text{bool} \rangle \quad \Gamma \vdash e_2 : \langle q_2 \ \tau \rangle \quad \Gamma \vdash e_3 : \langle q_3 \ \tau \rangle \quad \tau \in \{\text{int}, \text{bool}\}}{\Gamma \vdash \text{MUX}(e_1, e_2, e_3) : \langle q_1 \vee q_2 \vee q_3 \ \tau \rangle} \text{(MUX)}
\end{array}$$

Similarly, the typing rules for statements are as follows. The constraints are standard: the right-hand side of an assignment is a subtype of the left-hand side.

$$\begin{array}{c}
\frac{\Gamma \vdash x_1 : \langle q_1 \ \tau \rangle \quad \Gamma \vdash x_2 : \langle q_2 \ \tau \rangle \quad \Gamma \vdash x_3 : \langle q_3 \ \tau \rangle \quad q_2 \vee q_3 <: q_1}{\Gamma \vdash x_1 = \text{PHI}(x_2, x_3) : OK} \text{(PHI ASSIGN)} \\
\frac{\Gamma \vdash x : \langle q_1 \ \tau \rangle \quad \Gamma \vdash e : \langle q_2 \ \tau \rangle \quad q_2 <: q_1 \quad \tau \in \{\text{int}, \text{bool}\}}{\Gamma \vdash x = e : OK} \text{(VAR ASSIGN)} \\
\frac{\Gamma \vdash i : \langle \text{plain int} \rangle \quad \Gamma \vdash I : \langle \text{plain int} \rangle \quad \Gamma \vdash s : OK}{\Gamma \vdash \text{for } i \text{ in range}(I) : s : OK} \text{(FOR STMT)}
\end{array}$$

As mentioned earlier, the only annotations the program need provide is on program inputs. The compiler infers the rest of the annotations. The type system has two purposes (1) it imposes restrictions, and (2) it enables code generation, specifically, it informs the backend on weather a statement operates on shared variables or plaintext ones, and the backend generates appropriate MOTION code.

6 Backend-Independent Vectorization

This section describes our vectorization algorithm. While vectorization is a longstanding problem, and we build upon existing work on scalar expansion and classical loop vectorization [AK87], our algorithm is unique as it works on the MPC-IR SSA-form representation. We posit that vectorization over MPC-IR is a problem that warrants a fresh look, in part because of MPC’s unique linear structure and in part because vectorization interacts with other MPC-specific optimizations in non-trivial ways (other works have explored manual vectorization and protocol mixing in an ad-hoc way, e.g., [DSZ15; Büs+18; IMZ19]).

6.1 Dependence Analysis

We build a dependence graph where the nodes are the MPC-IR statements and the edges represent the def-use relations. Since MPC-IR is an SSA form, def-use edges $X \rightarrow Y$ are explicit. We distinguish between *forward* edges where X appears before Y in the linear MPC-IR and *backward* edges where Y appears before X .

Def-use edges We classify def-use edges as follows:

- same-level edge $X \rightarrow Y$ where X and Y are in the same loop nest, say i, j, k . E.g., the def-use edge 9 to 10 in the Biometric MPC-IR in Listing 1 is a same-level edge. A same-level edge can be a backward edge in which case a PHI-node is the target of the edge. E.g., 15 to 4 in Biometric is a same-level backward edge.
- outer-to-inner $X \rightarrow Y$ where X is in an outer loop nest, say i , and Y is in an inner one, say j, k . E.g., 1 to 4 in Biometric forms is an outer-to-inner edge.
- inner-to-outer $X \rightarrow Y$ where X is a PHI-node in an inner loop nest, k , and Y is in the enclosing loop nest i, j . E.g., the def-use from 8 to 12 gives rise to an inner-to-outer edge. An inner-to-outer edge can be a backward edge as well, in which case both X and Y are ϕ -nodes with the source X in a loop nested into Y ’s loop (not necessarily immediately).
- mixed forward edge $X \rightarrow Y$. X is in some loop i, j, k and Y is in a loop nested into i, j, k' . We transform mixed forward edges as follows. Let x be the variable defined at X . We add a variable and assignment $x' = x$ immediately after the i, j, k loop. Then we replace the use of x at Y with x' . This transforms a mixed forward edge into an "inner-to-outer" forward edge followed by an outer-to-inner forward edge. Thus, Basic Vectorization handles one of "same-level", "inner-to-outer", or "outer-to-inner" def-use edges.

Closures We define $closure(n)$ where n is a PHI-node. Intuitively, it computes the set of nodes (i.e., statements) that form a dependence cycle with n . The closure of n is defined as follows:

- n is in $closure(n)$
- X is in $closure(n)$ if there is a same-level path from n to X , and $X \rightarrow n$ is a same-level back-edge.
- Y is in $closure(n)$ if there is a same-level path from n to Y and there is a same-level path from Y to some X in $closure(n)$.

6.2 Scalar and Array Expansion

An important component of our algorithm is the scalar expansion to the corresponding loop dimensionality, which is necessary to expose opportunities for vectorization. In the Biometric example, $d = S[i*D+j]-C[j]$ equiv. to $d = S[i,j]-C[j]$, which gave rise to $N * D$ subtraction operations in the sequential schedule, is lifted. The argument arrays S and C are lifted and the scalar d is lifted: $d[i,j] = S[i,j]-C[i,j]$. The algorithm then detects that the statement can be vectorized.

Raise dimension The *raise_dim* function expands a scalar (or array). There are two versions of *raise_dim*. One reshapes an arbitrary access into a canonical read access in the corresponding loop. It takes the original array, the original access pattern function $f(i, j, k)$ in loop nest i, j, k and the loop bounds $((i:I), (j:J), (k:K))$ (cf. 5.2):

$$\text{raise_dim}(A, f(i, j, k), ((i:I), (j:J), (k:K)))$$

It produces a new 3-dimensional array A' by iterating over i, j, k and setting each element of A' as follows:

$$A'[i, j, k] = A[f(i, j, k)]$$

The end result is that uses of $A[f(i, j, k)]$ in loop nest i, j, k are replaced with canonical read-accesses to $A'[i, j, k]$ that can be vectorized. In the running Biometric example, $C' = \text{raise_dim}(C, j, (i:N, j:D))$ lifts the 1-dimensional array C into a 2-dimensional array. The i, j loop now accesses C' in the canonical way, $C'[i, j]$.

The other version of *raise_dim* lifts a lower-dimension array into a higher-dimension for access in a nested loop. It is necessary when processing outer-to-inner dependences. Here A is an i -array and raise dimension adds two additional dimensions; this version reduces to the above version by adding the access pattern function, which is just i :

$$\text{raise_dim}(A, i, (i:I, j:J, k:K))$$

Drop dimension *drop_dim* is carried out when an expanded scalar (or array) written in an inner loop is used in an enclosing loop. It takes a higher dimensional array, say i, j, k and removes trailing dimensions, say j, k :

$$\text{drop_dim}(A, (j:J, k:K))$$

It iterates over i and takes the result at the maximal index of j and k , i.e., the result at the last iterations of j and k :

$$A'[i] = A[i, J-1, K-1]$$

Arrays Conceptually, we treat all variables as arrays. There are three kinds of arrays.

- Scalars: We expand scalars into arrays for the purposes of vectorization. For those, all writes are canonical writes and all reads are canonical reads. We will *raise dimension* when a scalar gives rise to an outer-to-inner dependence edge (e.g., `sum!2` in line 6 of the MPC-IR code will be raised to a 2-dimensional array since `sum!2` is used in the inner j -loop). We will *drop dimension* when a scalar gives rise to an inner-to-outer dependence edge (e.g., `sum!3` for which the lifted inner loop computes D values per iteration of the outer loop, but the outer loop iteration only needs the last one.)
- Read-only input arrays: There are no writes, while we may have non-canonical reads, $f(i, j, k)$. Vectorization adds raise dimension operations at the beginning of the function to lift these arrays to the dimensionality of the loop where they are used, possibly *reshaping* the arrays.
- Read-write arrays: Writes are canonical (by restriction) but reads can be non-canonical. We may apply both raise and drop dimension, however, they respect the fixed dimensionality of the output array. The array cannot be raised to a dimension lower than its canonical (fixed) dimensionality and it cannot be dropped to lower dimension. The restriction to canonical writes essentially reduces the case of arrays to the case of scalars, simplifying vectorization and correctness reasoning.

6.3 Basic Vectorization Algorithm

There are two key phases of the algorithm. Phase 1 inserts raise dimension and drop dimension operations according to def-uses. E.g., if there is an inner-to-outer dependence, it inserts *raise_dim*, and similarly, if there is an outer-to-inner dependence, it inserts *drop_dim*. After this phase operations work on arrays of the corresponding loop dimensionality and we optimistically vectorize all arrays.

Phase 2 proceeds from the inner-most towards the outer-most loop. For each loop it anchors dependence cycles (closures) around pseudo PHI nodes then removes vectorization from the dimension of that loop.

There are two important points in this phase. First, it may break a loop into smaller loops which could allow vectorization in intermediate statements in the loop. Second, it creates opportunities for vectorization in the presence of write arrays, even though Cytron’s SSA adds a backward edge to the array PHI-node, thus killing vectorization of statements that access the array.

The code in blue color in the algorithm below highlights the extension with array writes. We advise the reader to omit the extension for now and consider just read-only arrays. We explain the extension in §6.5. (As many of our benchmarks include write arrays, it plays an important role.)

Phases 3 cleans up local arrays of references. This is an optional phase and our current implementation does not include it; thus, we elide it from this presentation.

```

{ Phase 1: Raise/drop dimension of scalars to corresponding loop nest. We traverse stmts linearly in MPC-IR. }
for each MPC stmt:  $x = Op(y_1, y_2)$  in loop  $i, j, k$  do
  for each argument  $y_n$  do
    case stmt'(def of  $y_n$ )  $\rightarrow$  stmt(def of  $x$ ) of
      same-level:  $y'_n$  is  $y_n$ 
      outer-to-inner: add  $y'_n[i, j, k] = raise\_dim(y_n)$  at stmt' (more precisely, right after stmt')
      inner-to-outer: add  $y'_n[i, j, k] = drop\_dim(y_n)$  at stmt (more precisely, in loop of stmt right after loop of stmt')
    end for
  { Optimistically vectorize all.  $I$  means vectorized dimension. }
  change to  $x[I, J, K] = Op(y'_1[I, J, K], y'_2[I, J, K])$ 
end for
{ Phase 2: Recreating for-loops for cycles; vectorizable stmts hoisted up. }
for each dimension  $d$  from highest to 0 do
  for each PHI-node  $n$  in loop  $i_1, \dots, i_d$  do
    compute closure( $n$ )
  end for
  {  $cl_1$  and  $cl_2$  intersect if they have common statement or update same array; "intersect" definition can be expanded }
  while there are closure  $cl_1$  and  $cl_2$  that intersect do
    merge  $cl_1$  and  $cl_2$ 
  end while
  for each closure  $cl$  (after merge) do
    create for  $i_d$  in ... loop
    add PHI-nodes in  $cl$  to header block
    add target-less PHI-node for  $A$  if  $cl$  updates array  $A$ 
    add statements in  $cl$  to loop in order of dependences
    { Dimension is not vectorizable: }
    change  $I_d$  to  $i_d$  in all statements in loop
    treat for-loop as monolith node for def-uses: e.g., some def-use edges become same-level.
  end for
  for each target-less PHI-node  $A_1 = PHI(A_0, A_k)$  do
    in vectorizable stmts, replace use of  $A_1$  with  $A_0$ 
    discard PHI-node if not used in any  $cl$ , replacing  $A_1$  with  $A_0$  or  $A_k$  as necessary
  end for
end for
{ Phase 3: Remove unnecessary dimensionality.}

```

Consider our running example in Listing 1(B). Phase 1 will raise dimensions of `min_sum!1` to a 1-dimensional array as it is defined outside of the loop but is used inside the i -loop. It will expand `C` into a 2-dimensional (i, j) -array. Phase 1 will also add `drop_dim` to drop the dimension of `sum!3`, which is defined in the inner loop and is of dimension (i, j) , but is used in the outer i -loop and needs to align to that loop

dimensionality.

Phase 2 starts with the inner j -loop. There are no dependences for the SUB and MUL statements (lines 9-10 in Listing 1(B)) and they are moved outside of the loop. The ADD is part of a cycle and it remains enclosed in a j -loop. Moving up to the outer i -loop, the addition j -loop is not part of a cycle in i and Phase 2 moves that loop outside vectorizing the i dimension of the summation (this results in the loop in lines 12-16 in Listing 1(C)). The MUX computations are part of cycles and they remain in i -loops.

6.4 Correctness Argument

We build a correctness argument as follows. First, we define the MPC-IR syntax. We then define the *linearization* of an MPC-IR program as an *interpretation* over the syntax. The linearization is a *schedule* as defined in §4. We prove a theorem that states that the Basic vectorization algorithm preserves the def-use relations, or in other words, linearization of the vectorized MPC-IR program gives rise to the exact same set of def-use pairs as linearization of the original program does. It follows easily that the schedule corresponding to the vectorized program computes the same result as the schedule corresponding to the original program.

MPC-IR Syntax Fig. 3 states the syntax and linearization semantics of MPC-IR. Although notation is heavy, the linearization simply produces schedules as discussed in §2 and §4. The iterative MPC-IR gives rise to what we called sequential schedule where loops are unrolled and MPC-IR with vectorized dimensions gives rise to what we called parallel schedule. For simplicity, we consider only scalars and read-only arrays, however, the treatment extends to write arrays as well (with our restriction on array writes to canonical writes). $x[i, J, k]$ denotes the value of scalar variable x at loop nest i, j, k . Upper case J denotes a vectorized dimension and lower case i, k denote iterative dimensions. There are semantic restrictions over the syntax: (1) x is a 3-dimensional array and (2) $x[i, J, k]$ is enclosed in for-loops on non-vectorized dimensions i and k :

```

1 for i in range(I):
2   ...
3   for k in range(K):
4     ... x[i,J,k] ...

```

Linearization Linearization is the concretization operation, which, as we mentioned earlier computes a schedule. The concretization function γ is defined as an interpretation of MPC-IR syntax, as is standard. It is shown in the middle column of Fig. 3. The concretization of an `op_SIMD` statement expands the vectorized dimension(s) into *parallel* statements; `||` introduces SIMD (parallel) execution. The concretization of the `for i in range(I) : s` statement simply unrolls the loop substituting i with 0, 1, etc.; here `;` denotes *sequential* execution.

s	$::= s_1; s_2$ $ x[i, J, k] = \text{op_SIMD}(y_1[i, J, k], y_2[i, J, k])$	$\gamma(s) = \gamma(s_1) ; \gamma(s_2)$ $\gamma(x[i, J, k] = \text{op_SIMD}(y_1[i, J, k], y_2[i, J, k])) =$ $x[i, 0, k] = y_1[i, 0, k] \text{ op } y_2[i, 0, k] \ $ $x[i, 1, k] = y_1[i, 1, k] \text{ op } y_2[i, 1, k] \ \dots \ $ $x[i, J-1, k] = y_1[i, J-1, k] \text{ op } y_2[i, J-1, k]$	<i>sequence</i> <i>operation</i>
	$ x[i, J, k] = \text{const}$ $ x[i, J, k] = \text{PHI}(x_1[i, J, k], x_2[i, J, k-1])$ $ x[i, J, k] = \text{raise_dim}(x'[i], (j: J, k: K))$ $ x[i, J] = \text{drop_dim}(x'[i, J, k], k)$ $ \text{for } i \text{ in range}(I) : s$	<i>analogous</i> $\gamma(\text{for } i \text{ in range}(I) : s) =$ $\gamma(s)[0/i] ; \gamma(s)[1/i] ; \dots ; \gamma(s)[I-1/i]$	<i>constant</i> <i>pseudo PHI</i> <i>raise dimension(s)</i> <i>drop dimension(s)</i> <i>loop</i>

Figure 3: MPC-IR Syntax and Semantics. γ defines the semantics of MPC-IR which is a linearization of input MPC-IR. A SIMD operation parallelizes operations across the vectorized J dimension. `||` denotes parallel execution, which is standard. γ of a for loop unrolls the loop. `;` denotes sequential execution. Iterative MPC-IR trivially extends to non-vectorized dimensions over the enclosing loops.

As an example, consider the vectorized MPC-IR from our running example. All variables are two dimensional arrays and the loop is vectorized in I but iterative in j :

```

1 for j in range(0, D):
2   sum!3[I,j] = PHI(sum!2[I,j], sum!4[I,j-1])
3   sum!4[I,j] = ADD(sum!3[I,j], p[I,j])

```

Assuming $D = 2$ and $I = 2$ for simplicity, linearization produces the following schedule:

```

1 sum!3[0,0] = PHI(sum!2[0,0], sum!4[0,-1]) ||
2               sum!3[1,0] = PHI(sum!2[1,0], sum!4[1,-1])
3 ;
4 sum!4[0,0] = ADD(sum!3[0,0], p[0,0]) ||
5               sum!4[1,0] = ADD(sum!3[1,0], p[1,0])
6 ;
7 sum!3[0,1] = PHI(sum!2[0,1], sum!4[0,0]) ||
8               sum!3[1,1] = PHI(sum!2[1,1], sum!4[1,0])
9 ;
10 sum!4[0,1] = ADD(sum!3[0,1], p[0,1]) ||
11               sum!4[1,1] = ADD(sum!3[1,1], p[1,1])

```

Note that by definition of the pseudo PHI function, $\text{PHI}(\text{sum!2}[0,0], \text{sum!4}[0,-1])$ evaluates to $\text{sum!2}[0,0]$ and therefore, the -1 index in the second argument does not matter.

Statements and def-uses over MPC-IR Let a be an MPC-IR program. Since MPC-IR is an SSA form, def-use edges in a are explicit (as in §6.1): if $s_0 \in a$ defines variable x , e.g., $x = \dots$, $s_1 \in a$ uses x , e.g., $\dots = \dots x$, then there is a def-use edge from s_0 to s_1 . We write $s_0[i, j, k]$ for statement s_0 enclosed in loop nest i, j, k .

Let a_0, a_1 be two MPC-IR programs. Two statements, $s_0 \in a_0$ and $s_1 \in a_1$ are *same*, written $s_0 \equiv s_1$ if they are of the same operation and they operate on the same variables: same variable name and same dimensionality. Recall that dimensions in MPC-IR are either iterative, lower case, or vectorized, upper case. Two statements are same even if one operates on an iterative dimension and the other one operates on a vectorized one, e.g., $s_0[i, j, k] \equiv s_1[I, j, K]$.

Statements and def-uses over linearized schedule An *atomic* statement is a statement produced by linearization. We write \underline{s}_0 to denote statements in the concrete schedule as well as $\underline{s}_0[\underline{i}, \underline{j}, \underline{k}]$ to denote fully instantiated values of i, j , and k , such as for example $\underline{s}_0[0, 1, 0]$. Clearly, the linearization of same statements produces the same set of atomic statements in the linearized schedule.

A def-use pair of atomic statements, denoted $\underline{s}_0 \rightarrow \underline{s}_1$ (indexing implicit), is defined in the standard way as well: \underline{s}_0 writes a location, say $x[\underline{i}, \underline{j}, \underline{k}]$, and \underline{s}_1 reads the same location.

Formal treatment Property P defined below relates the linearized schedule of iterative MPC-IR program a_0 to the linearized schedule of the vectorized program a_1 . More precisely, a_0 is the MPC-IR program augmented with raise and drop dimension statements, i.e., Phase 1 without optimistic vectorization of all dimensions. a_1 is produced from a_0 by Phase 2 of the Basic Vectorization algorithm.

Definition 1. We say that $\gamma(a_0) \equiv \gamma(a_1)$ iff (1) atomic statement $\underline{s}[\underline{i}, \underline{j}, \underline{k}] \in \gamma(a_0)$ iff $\underline{s}[\underline{i}, \underline{j}, \underline{k}] \in \gamma(a_1)$ and (2) $\underline{s}_0 \rightarrow \underline{s}_1 \in \gamma(a_0)$ iff $\underline{s}_0 \rightarrow \underline{s}_1 \in \gamma(a_1)$ (indexing implicit).

Let us first prove the following lemma which states that Basic vectorization preserves statements and def-use edges in the original MPC-IR.

Lemma 1. For each statement s in a_0 , there is same statement s' in a_1 , and vice versa. For each def-use edge e in a_0 , there is a same edge e' in a_1 , and vice versa.

Proof. Proof sketch of Lemma 1. Phase 2 of Basic Vectorization does not introduce any new statements in the code, it just vectorizes dimensions. Similarly, reordering of statements preserves exactly the def-use edges in the original MPC-IR. \square

The main theorem below states that Basic vectorization preserves def-use edges.

Theorem 1. $\gamma(a_0) \equiv \gamma(a_1)$.

Proof. Proof sketch of Theorem 1. The first condition of property P follows directly from Lemma 1. The proof of the second condition is by analysis of the def-use edges in $\gamma(a_0)$ and the corresponding edges in $\gamma(a_1)$; the key is that Basic vectorization preserves the def-uses in a_0 .

A forward edge $s_0 \rightarrow s_1 \in a_0$ remains a forward edge in a_1 . Without loss of generality, let us assume an outer loop i and a nested loop j . The forward edge entails the following ordering in linearization $\gamma(a_0)$:

$$\begin{array}{ll} \underline{s_0}[\underline{i}] ; \underline{s_1}[\underline{i}, \underline{j}] & \text{outer-to-inner edge} \\ \underline{s_0}[\underline{i}, \underline{j}] ; \underline{s_1}[\underline{i}, \underline{j}] & \text{same-level edge} \\ \underline{s_0}[\underline{i}, \underline{j}] ; \underline{s_1}[\underline{i}] & \text{inner-to-outer edge} \end{array}$$

meaning that for a fixed \underline{i} , def $\underline{s_0}[\underline{i}]$ is scheduled *before* use $\underline{s_1}[\underline{i}]$. Due to the preservation of the edge in a_1 , the above ordering holds in $\gamma(a_1)$ as well.

Consider a backward edge $s_0 \rightarrow s_1 \in a_0$. We have that s_1 is a PHI-node in some loop, say i . There are two cases: (1) there is a path of forward edges from s_1 to s_0 , and (2) there is no such path. In case (1) Basic vectorization detects a cycle (closure) around s_1 , and therefore, $s_0 \rightarrow s_1$ remains a backward edge in a_1 . The linearization of the backward edge imposes ordering $\underline{s_0}[\underline{i} - 1] ; \underline{s_1}[\underline{i}]$ and due to preservation of the backward edge in a_1 , the ordering holds in $\gamma(a_1)$ as well. In case (2) Basic vectorization may turn the backward edge into a forward one, however, it preserves the $\underline{s_0}[\underline{i} - 1] ; \underline{s_1}[\underline{i}]$ ordering constraint by construction. \square

The key argument is that the Basic vectorization algorithm preserves def-uses when it transforms a_0 into a_1 . This leads to preservation of concrete edges in $\gamma(a_0)$ into $\gamma(a_1)$. A corollary of the main theorem follows:

Corollary 1.1. $\gamma(a_0)$ and $\gamma(a_1)$ produce same result, or more precisely, for every location $\times[\underline{i}, \underline{j}, \underline{k}]$, $\gamma(a_0)$ and $\gamma(a_1)$ compute the same result.

Proof. Proof sketch of Corollary 1.1. This can be established by induction over the length of def-use chains of computation in $\gamma(a_0)$. Assume that for all chains of length $\leq n$ all locations $l[\underline{i}, \underline{j}, \underline{k}]$ hold the same value in $\gamma(a_0)$ and $\gamma(a_1)$. A chain of length $n + 1$ results from the execution of a statement $\times[\underline{i}, \underline{j}, \underline{k}] = y[\underline{i}, \underline{j}, \underline{k}]$ op $z[\underline{i}, \underline{j}, \underline{k}]$. By property P , there is the same statement in $\gamma(a_1)$ and it is scheduled after the definitions of $y[\underline{i}, \underline{j}, \underline{k}]$ and $z[\underline{i}, \underline{j}, \underline{k}]$. By the inductive hypothesis $y[\underline{i}, \underline{j}, \underline{k}]$ and $z[\underline{i}, \underline{j}, \underline{k}]$ hold the same values in $\gamma(a_0)$ as in $\gamma(a_1)$. Therefore, locations $\times[\underline{i}, \underline{j}, \underline{k}]$ hold the same value as well. We remark that due to the SSA form, each location $l[\underline{i}, \underline{j}, \underline{k}]$ is defined at most once. For clarity, we elide PHI nodes and raising and dropping dimensions; extending def-use reasoning is straight forward. \square

6.5 Extension with Array Writes

Array writes may introduce infeasible loop-carried dependencies. Consider an example from [AN88]:

```

1 for i in range(N):
2   A[i] = B[i] + 10;
3   B[i] = A[i] * D[i-1];
4   C[i] = A[i] * D[i-1];
5   D[i] = B[i] * C[i];

```

In Cytron's SSA this code (roughly) translates into

```

1 for i in range(N):
2   A_1 = PHI(A_0, A_2)
3   B_1 = PHI(B_0, B_2)
4   C_1 = PHI(C_0, C_2)
5   D_1 = PHI(D_0, D_2)
6   A_2 = update(A_1, i, B_1[i] + 10);
7   B_2 = update(B_1, i, A_2[i] * D_1[i-1]);
8   C_2 = update(C_1, i, A_2[i] * D_1[i-1]);
9   D_2 = update(D_1, i, B_2[i] * C_2[i]);

```

$B_1 = \text{PHI}(B_0, B_2)$ anchors a cycle that includes statement $A_2 = \text{update}(A_1, i, B_1[i] + 10)$; a naive approach will not vectorize the latter statement even though there is no loop-carried dependency from the write of $B_1[i]$ at 7 to the read of $\dots = B_1[i]$ at 6.

The following algorithm removes certain infeasible loop-carried dependencies that are due to array writes. Consider a loop with index $0 \leq j < J$ nested at i, j, k . Here i is the outermost loop and k is the innermost loop.

```

for each array A written in loop  $j$  do
  { including enclosed loops in  $j$  }
  dep = False
  for each def-of-A:  $A_m[f(i, j, k)] = \dots$  and use-of-A:  $\dots = A_n[f'(i, j, k)]$  in loop  $j$  do
    if  $\exists \underline{i}, \underline{j}, \underline{j}', \underline{k}, \underline{k}'$ , s.t.  $0 \leq \underline{i} < I$ ,  $0 \leq \underline{j}, \underline{j}' < J$ ,  $0 \leq \underline{k}, \underline{k}' < K$ ,  $\underline{j} < \underline{j}'$ , and  $f(\underline{i}, \underline{j}, \underline{k}) = f'(\underline{i}, \underline{j}', \underline{k}')$  then
      dep = True
    end if
  end for
  if dep == False then
    remove back edge into A's  $\phi$ -node in loop  $j$ .
  end if
end for

```

Consider a loop j enclosed in some fixed i . Only if an update (definition) $A_m[f(i, j, k)] = \dots$ at some iteration \underline{j} references the *same* array element as a use $\dots = A_n[f'(i, j, k)]$ at some later iteration \underline{j}' , we may have a loop-carried dependence for A due to this def-use pair. (In contrast, Cytron's algorithm inserts a loop-carried dependency every time there is an array update.) The algorithm above examines all def-use pairs in loop j , including defs and uses in nested loops, searching for values $\underline{i}, \underline{j}, \underline{j}', \underline{k}, \underline{k}'$ that satisfy $f(\underline{i}, \underline{j}, \underline{k}) = f'(\underline{i}, \underline{j}', \underline{k}')$. If such values exist for some def-use pair, then there is a potential loop-carried dependence on A; otherwise there is not and we can remove the spurious backward edge thus "freeing up" statements for vectorization.

We use Z3 [MB08] to check satisfiability of the formula

$$(0 \leq \underline{i} < I) \wedge (0 \leq \underline{j}, \underline{j}' < J) \wedge (0 \leq \underline{k}, \underline{k}' < K) \wedge (\underline{j} < \underline{j}') \wedge f(\underline{i}, \underline{j}, \underline{k}) = f'(\underline{i}, \underline{j}', \underline{k}')$$

Formulas f and f' are simple as loop nests are typically of depth 2-3. Therefore, Z3 completes the process instantly.

Consider the earlier example. There is a single loop, i . Clearly, there is no pair \underline{i} and \underline{i}' , where $\underline{i} < \underline{i}'$ that make $\underline{i} = \underline{i}'$ due to the def-use pairs of A 6-7 and 6-8. Therefore, we remove the backward edge from 6 to the phi-node 2. Analogously, we remove the backward edges from 7 to 3 and from 8 to 4. However, there are many values $\underline{i} < \underline{i}'$ that make $\underline{i} = \underline{i}' - 1$ and the backward edge from 9 to 5 remains (def-use pairs for D). As a result of removing these spurious edges, Vectorization will find that statement 6 is vectorizable. Statements 7, 8 and 9 will correctly appear in the for-loop.

This step renders some array phi-nodes *target-less*, or in other words, these nodes are not targets of any def-use edge. We handle target-less phi-nodes with a minor extension of Basic Vectorization (Phase 2, extension shown in blue). First, we merge closures that update the same array. This simplifies handling of array PHI-nodes: if each closure is turned into a separate loop, each loop will need to have its own array phi-node to account for the update and this would complicate the analysis. Second, we add the target-less node of array A back to the closure that updates A — the intuition is, even if there is no loop-carried dependence from writes to reads on A, A is written and the write (i.e., update) cannot be vectorized due to a different cycle; therefore, the updated array has to carry to the next iteration of the loop. Third, in cases when the phi-node remains target-less, i.e., cases when the array write can be vectorized, we have to properly remove the phi-node replacing uses of the left-hand side of the phi-node with its arguments (the last snippet in blue).

Recall that we restrict array updates to *canonical updates*, that is, an update $A[i, j] = \dots$ is enclosed in loops on i and j . It may be enclosed into a nested loop i, j, k , however, the indices correspond to the outermost loops. This restriction ensures that the array shape does not change and raise dimension and drop dimension can be applied in the same way as in the basic case, thus allowing us to extend correctness reasoning from

the basic case. We will look to relax the restriction in future work, when additional optimizations are overlaid to vectorization (*cf.* §9).

7 Compiler Backends

Translation into the MOTION and MP-SPDZ backends is done by standard interpretation (i.e., syntax-directed translation) over the MPC-IR syntax shown in Fig. 3.

Naturally, translation into C++ code for the MOTION framework presents more challenges. First, MPC-IR requires `shared` qualifiers only on input variables, while MOTION requires all variables to be typed either `shared` or `plain`. To resolve this, we infer qualifiers for all variables by performing *taint analysis* according to the rules of §5.5; this is described in §7.1. Another challenge is dealing with public values e.g., constants. Since there is no support for these, one has to provide such public values as shared input from one of the parties. Input gates are expensive and a naive implementation could introduce a performance hit. Instead, we keep shared copies of plain (public) variables, and update them in lock-step with updates to plain variable. When a plain variable is needed in a shared context, we use the shared copy. Below we present details on the translation into MOTION C++.

The current implementation supports unsigned integers, booleans, tuples, and lists at the frontend level, and it is our intention to extend it with additional types. At the backend level it supports unsigned integers and booleans for secure computation. The MPC protocol parameters e.g., computational/statistical security parameters, are specified independently for each backend. In our evaluation, we use each backend’s default values for them.

7.1 Taint Analysis

The taint analysis works on MPC-IR, which lacks if-then-else control flow. This significantly simplifies treatment as there is no need to handle conditionals and implicit flow. The compiler uses the rules in §5.5, which are standard in positive-negative qualifier systems (here `shared` is the positive qualifier and `plain` is the negative one). In practical terms the taint analysis applies the following simple rules:

1. Loop counters are always `plain`.
2. If any variable on the right-hand side `rhs` of an assignment is shared, then the assigned variable `lhs` is `shared` following subtyping rule `rhs <: lhs`. The analysis applies this rule repeatedly until it reaches a fixpoint.
3. Any variable that cannot be determined as shared via the above rules is `plain`.

In the below snippet `sum!1`, `sum!2` and `sum!3` form a dependency cycle and there is no `shared` value that flows to either one. They are inferred as plaintext.

```
1 plaintext_array = [0, 1, 2, ...]
2 sum!1 = 0
3 for i in range(0, N):
4     sum!2 = PHI(sum!1, sum!3)
5     sum!3 = sum!2 + plaintext_array[i]
```

When converting to MOTION code, any plaintext value used in the right-hand side of a shared assignment is converted to a shared value for that expression.

7.2 From (Optimized) MPC-IR to MOTION

MOTION supports FOR loops and SIMD operations, so translation from MPC-IR to MOTION C++ code is relatively straightforward.

1 A[i] = val	1 A!2 = update(A!1, i, val)	1 A_1[j] = val; 2 A_2 = A_1;
IMP Source	MPC-IR	MOTION Code

Table 2: MOTION translation: Array updates.

1 for i in range(N): 2 tmp = PHI(arr[j], val!0) 3 ...	1 _MPC_PLAINTEXT_i = 0; 2 tmp = arr[_MPC_PLAINTEXT_i]; 3 for (; _MPC_PLAINTEXT_i < _MPC_PLAINTEXT_N; _MPC_PLAINTEXT_i++) { 4 if (_MPC_PLAINTEXT_i != 0) { 5 tmp = val.0; 6 } 7 ... 8 }
MPC-IR	MOTION code

Table 3: MOTION translation: FOR loop with Phi nodes.

7.2.1 Variable Declarations

Our generated C++ uses the following variable-naming scheme: shared variables are named the same as in the MPC-IR with the ! replaced with an underscore (e.g. `sum!2` would be translated to `sum_2`). Plaintext variables follow the same naming convention as shared variables but are prefixed with `_MPC_PLAINTEXT_`. The shared representation of constants are named `_MPC_CONSTANT_` followed by the literal constant (e.g. the shared constant 0 would be named `_MPC_CONSTANT_0`).

The generated MOTION code begins with the declaration of all variables used in the function, including loop counters. If a variable is a vectorized array, it is initialized to a correctly-sized array of empty MOTION shares. Additionally, each plaintext variable and parameter has a shared counterpart declared. Next, all constant values which are used as part of shared expressions are initialized as a shared input from party 0. Finally, plaintext parameters are used as shared inputs from party 0 to initialize their shared counterparts.

7.2.2 Code Generation

Once the function preamble is complete, the MPC-IR is translated into C++ one statement at a time. The linear structure of MPC-IR enables this approach to translation. If there is no vectorization present in a statement, translation to C++ is straightforward: outside of MUX statements and array updates, non-vectorized assignments, expressions, and returns directly translate into their C++ equivalents. Non-vectorized MUX statements are converted to MOTION’s MUX member function on the condition variable. Array updates are translated into two C++ assignments: one to update the value in the original array and one to assign the new array as shown in Listing 2.

MPC FOR loops are converted to C++ FOR loops which iterate the loop counter over the specified range. Pseudo PHI nodes are broken into two components: the “FALSE” branch which assigns the initial value of the PHI node and the “TRUE” branch which assigns the PHI node’s back-edge. The assignment of the “FALSE” branch occurs right before the PHI node’s enclosing loop. As these assignments may rely on the loop counter, the loop counter is initialized before these statements. Inside of the PHI node’s enclosing loop, a C++ `if` statement is inserted to only assign the true branch of the PHI node after the first iteration. Listing 3 illustrates this translation.

7.2.3 Vectorization and SIMD Operations

Vectorization is handled with utility functions to manage accessing and updating slices of arrays. All SIMD values are stored in non-vectorized form as 1-dimensional `std::vectors` in row-major order. Whenever a SIMD value is used in an expression, the utility function `vectorized_access()` takes the multi-dimensional representation of a SIMD value, along with the size of each dimension and the requested slice’s indices,

<pre>1 sum!4[l] = ADD.SIMD(sum!3[l], p[l, j])</pre>	<pre>1 vectorized_assign(sum_4, {_MPC_PLAINTEXT_N}, {true}, {}), 2 vectorized_access(sum_3, {_MPC_PLAINTEXT_N}, {true}, {}) + 3 vectorized_access(p, {_MPC_PLAINTEXT_N}, {_MPC_PLAINTEXT_D}, {true, false}, 4 {_MPC_PLAINTEXT_j});</pre>
MPC-IR	MOTION Code

Table 4: MOTION Translation: Assignment to SIMD value

<pre>1 raise_dim(i + j, (i:N, j:M))</pre>	<pre>1 lift(std::function([&](const std::vector<std::uint32_t > &idxs) {return idxs[0] + idxs[1];}), 2 {_MPC_PLAINTEXT_N, _MPX_PLAINTEXT_M})</pre>
MPC-IR	MOTION Code

Table 5: MOTION translation: Raising dimensions.

and converts that slice to a MOTION SIMD value. Because MOTION supports SIMD operations using the same C++ operators as non-SIMD operations, we do not need to perform any other transformations to the expression. Therefore, once vectorized accesses are inserted the translation of an expression containing SIMD values is identical to that of expressions without SIMD values.

Similarly, the `vectorized_assign()` function assigns a (potentially SIMD) value to a slice of a vectorized array. This operation cannot be done with a simple subscript as SIMD assignments will update a range of values in the underlying array representation.

Updating SIMD arrays is also implemented differently from updating non-vectorized arrays. Instead of separating the array update from the assignment of the new array, these steps are combined with the `vectorized_update()` utility function. This function operates identically to `vectorized_assign()`, however it additionally returns the array after the assignment occurs. This value is then used for the assignment to the new variable. Listing 4 illustrates `vectorized_assign()` and `vectorized_update()` on the Biometric example.

7.2.4 Reshaping and Raising Dimensions

Raising the dimensions of a scalar or array uses the `lift()` utility function which takes a lambda for the raised expression and the dimensions of the output. This function is also used for the scalar expansion of values which have been lifted out of FOR loops as described in §6.2. This function evaluates the expression for each permutation of indices along the dimensions and returns the resulting array in row-major order. The lambda accepts an array of integers representing the index along each of the dimensions being raised, and the translation of the expression which is being raised replaces each of the dimension index variables with the relevant subscript of this array. There is also a special case of the `lift()` function which occurs when we are raising an array. In this case, instead of concatenating the array for each index, we extend the array along all dimensions being raised which are not present in the array already. For example, when raising an array with dimensions $N \times M$ to an array with dimensions $N \times M \times D$, the input array will simply be extended along the D dimension: $A'[n, m, d] = A[n, m]$ for every d . If the input array is already correctly sized it will be returned as-is.

Dropping dimensions use the `drop_dim()` and `drop_dim_monoreturn()` utility functions. They function identically but the latter returns a scalar for the case when the final dimension of an array is dropped. These functions take the non-vectorized representation of an array, along with the dimensions of that array, and return the array with the final dimension dropped.

7.2.5 Upcasting from Plaintext to Shared

Currently, our compiler only supports the `Bmr` and `BooleanGMW` protocols as MOTION does not implement all operations for other protocols. MOTION does not support publicly-known constants for these protocols, so all conversions from plaintext values to shares are performed by providing the plaintext value as a shared input from party 0. Due to this limitation, our translation to MOTION code attempts to minimize the

number of conversions from a plaintext value. This is accomplished by creating a shared copy of each plaintext variable and updating that copy in lock-step with the plaintext variable. Since variables are often initialized to a common constant value (e.g. 0), this approach decreases the number of input gates by only creating a shared input for each initialization constant. Loop counters must still be converted to a shared value on each iteration that they are used, however we only generate this conversion when necessary, i.e., when the counter flows to a shared computation. This is to prevent unnecessary increase in the number of input gates when loop counters are only used as plaintext.

Due to the SSA translation phase as well as the conversions to and from SIMD values which our utility functions perform, our generated vectorized MOTION code often includes multiple copies of arrays and scalar values. These copies do not incur a runtime cost as the arrays simply hold *pointers* to the underlying shares, so no new shares or gates are created as a result of this copying. Cost in MPC programs is dominated by shares and computation on shares.

8 Evaluation and Analysis

We evaluate on the two major MPC backend frameworks, MOTION (commit 6a7c1c7) and MP-SPDZ (commit 640b1a9). Specifically, we compare the iterative (non-vectorized) version of a benchmark against the vectorized one. Note that, since we work at the higher level of MPC-IR, the source-to-source compilation overhead of our compiler is negligible: the slowest benchmark to compile to either MP-SPDZ or MOTION takes $\approx 65\text{ms}$ (average of 100 runs).

8.1 Benchmarks

We compare on a benchmark suite that includes standard MPC benchmarks [Büs+18; IMZ19] as well as new ones from [FN21]. The following is the list of benchmarks:

- (1) *Biometric Matching*: Server has a database S of N records, each record’s dimension is D . Client submits a query C , client and server compute the closest record to C in an MPC. We use $N=128$ for *both* (vectorized and non-vectorized experiment) and $N=4096$ for *vec* (i.e., only the vectorized experiment because the non-vectorized experiment runs out of memory for this input size). D is fixed at 4.
- (2) *Convex Hull*: Given a polygon of N vertices (split between Alice and Bob), Alice and Bob compute the convex hull in an MPC. The benchmark is adapted from [FN21]. We use $N=32$ for *both* and $N=256$ for *vec*.
- (3) *Count 102*: Alice has a string of N symbols, Bob has a regular expression of the form $1(0^*)2$, and they compute the number of substrings that match the regular expression. It is adapted from [FN21]. We use $N=1024$ for *both* and $N=4096$ for *vec*.
- (4) *Count 10*: Same as *Count 102* except now the regular expression is of the form $1(0+)$.
- (5) *Cryptonets Max Pooling*: Given an $R \times C$ -matrix with elements split between Alice and Bob, they compute the max pooling subroutine of the cryptonet benchmark [Dow+16]. We use $R=64$, $C=64$ for the experiments.
- (6) *Database Join*: Given two databases of sizes A and B containing 2-element records, the two parties compute cross join in an MPC. We use $A=B=32$ for *both* and $A=B=64$ for *vec*.
- (7) *Database Variance*: The parties compute variance in a database of N records. $N=512$ for *both* and $N=4096$ for *vec*.
- (8) *Histogram*: Given N ratings (out of 5-stars), the parties compute their histogram. The benchmark is taken from [IMZ19; FN21]. We use $N=512$ for *both* and $N=4096$ for *vec*.
- (9) *Inner Product*: The parties compute inner product of two N -element vectors. $N=512$ for *both* and $N=4096$ for *vec*.

- (10) *k-means Iteration*: Iteration of k-means database clustering [JW05; VC03]. Here L is the size of input data, and N is the number of clusters. We use $L=32$, $N=5$ for *both* and $L=256$, $N=8$ for *vec*.
- (11) *Longest 102*: Same as *Count 102* except that the parties compute the largest substring matching the regular expression. It is adapted from [FN21]. We use the same parameters as in *Count 102*.
- (12) *Max Distance b/w Symbols*: Alice has a string of N symbols and Bob has some symbol 0. They compute the maximum distance between 0s in the string in an MPC. It is adapted from [FN21]. We use $N=1024$ for *both* and $N=2048$ for *vec*.
- (13) *Minimal Points*: Given a set of N points (split between Alice and Bob), Alice and Bob compute a set of minimal points, i.e. there is no other point that has both a lower x and y coordinate. It is adapted from [FN21]. We use $N=32$ for *both* and $N=64$ for *vec*.
- (14) *MNIST ReLU*: Given an input of $O \times I$ elements, the parties compute the MNIST ReLU subroutine. We use $I=512$ for *both* and $I=2048$ for *vec*. O is fixed at 16.
- (15) *Private Set Intersection (PSI)*: Alice holds a set of size A , Bob holds a set of size B and they compute the intersection. We use $A=B=128$ for *both* and $A=B=1024$ for *vec*.

8.2 MOTION Experiments

8.2.1 Experiment Setup

Experiment hardware is generously provided by CloudLab [Dup+19]. For the network, we consider two settings, LAN and WAN. In the LAN setting, we use c6525-25g machines connected via a 10Gbps link with < 1 ms latency. These machines are equipped with 16-core AMD 7302P 3.0GHz processors and 128GB of RAM. For WAN, we use a c6525-25g machine (located in Utah, US) and a c220g1 machine (located in Wisconsin, US). The c220g1 machine is equipped with two Intel E5-2630 8-core 2.40GHz processors and 128GB of RAM. We measured the connection bandwidth between these machines to be 560Mbps and average round trip time (RTT) to be 38ms.

We run 2PC and 3PC experiments on LAN with input datasets that range in size from 2 to 4096. In the WAN setting we only perform 2PC experiments to save time; as evidenced by LAN experiments in §8.2.2, 3PC experiments would only take longer to run. We run all experiments 5 times and report average values of various metrics. The standard deviation, shown as error bar on top of the histogram bars in the graphs, in all observations is at most 4.5% of the mean. Therefore, more runs will not significantly improve results’ accuracy. Tables and graphs in this section are with the largest dataset for which the non-vectorized run completes (it typically runs out of memory while the vectorized runs continue).

8.2.2 Results and Analysis

A summary of the effects of vectorization (MOTION backend) is presented in Table 6. In addition we show graphically circuit evaluation times in Fig. 4 and a variety of other metrics in Fig. 5. In terms of amenability to vectorization, we divide the benchmarks into 3 categories: (1) *High*: these include Convex Hull, Cryptonets (Max Pooling), Minimal Points and Private Set Intersection. These benchmarks are highly parallelizable and see 47x to 23x speedup in BMR, and 33x to 23x in GMW. (2) *Medium*: these include Biometric Matching, DB Variance, Histogram, Inner product, K-means iteration and MNIST ReLU. These benchmarks have non-parallelizable phases e.g. the summing phase of inner product and Biometric Matching.⁴ Still, most computation is parallelizable and it results in speedup from 23x to 5x in BMR, and 23x to 3x in GMW protocol. (3) *Low*: these include the Database Join and the regular expression benchmarks (Count 102, Count 10, Longest 102 and Max Distance between Symbols). Fewer operations in these programs are parallelizable, thus the speedup is lower. We see a speedup from 2x to 1.1x in BMR. In GMW, Database

⁴We remark that summation can be parallelized as a $\log(N)$ -depth tree (e.g., as in [Büs18]). This is an instance of the divide-and-conquer paradigm [FN21] and we have chosen to leave it as future work — divide-and-conquer can be applied more generally and in a principled way in conjunction with vectorization, rather than just on summation.

Table 6: Non-vectorized vs. vectorized comparison in 2PC LAN setting. Times in seconds, Communication in MiB, Numbers in 1000s and rounded to nearest integer; vectorized benchmarks have (V) in name. All metrics are produced by MOTION.

Benchmark	GMW						BMR					
	Online	Setup	# Gates	Circ Gen	# Msgs	Comm.	Online	Setup	# Gates	Circ Gen	# Msgs	Comm.
Biometric Matching	146	16	1,784	119	1,413	140	89	263	1,595	139	2,716	312
Biometric Matching (V)	12	4	34	2	28	14	2	13	30	4	61	130
Convex Hull	48	6	551	40	516	51	28	72	494	39	695	80
Convex Hull (V)	0	1	2	0	1	4	0	2	1	1	2	32
Count 102	79	6	418	35	525	52	15	62	269	33	785	92
Count 102 (V)	71	5	316	24	332	34	11	30	167	16	304	59
Count 10s	79	6	419	35	525	52	14	62	270	33	785	92
Count 10s (V)	71	4	316	24	332	34	11	29	167	16	304	59
Cryptonets (Max Pooling)	50	11	688	46	554	55	36	89	608	51	898	110
Cryptonets (Max Pooling) (V)	1	1	7	1	2	5	2	4	7	2	12	49
Database Join	70	8	433	48	790	80	19	229	458	119	3,518	427
Database Join (V)	54	6	320	35	575	61	16	112	320	57	1,457	285
Database Variance	166	18	2,009	135	1,639	163	95	269	1,708	145	2,795	320
Database Variance (V)	37	6	321	24	334	43	10	30	170	13	178	141
Histogram	94	10	862	68	979	97	27	94	491	51	1,132	135
Histogram (V)	33	5	166	16	164	23	7	17	92	13	154	68
Inner Product	127	15	1,675	108	1,308	130	83	250	1,526	134	2,623	301
Inner Product (V)	16	5	158	12	165	25	6	18	83	7	86	127
k-means	108	12	1,333	88	1,090	108	63	185	1,141	99	1,958	225
k-means (V)	6	3	47	4	43	12	2	11	32	4	54	95
Longest 102	93	7	650	52	713	71	26	93	475	49	1,091	128
Longest 102 (V)	169	6	544	41	519	53	25	60	369	33	605	95
Max. Dist. b/w Symbols	71	8	572	43	576	57	24	69	397	38	748	89
Max. Dist. b/w Symbols (V)	166	7	538	39	512	51	24	57	363	32	589	78
Minimal Points	35	5	458	31	369	37	24	46	401	26	347	40
Minimal Points (V)	0	1	1	0	1	3	0	1	1	0	1	16
MNIST ReLU	132	31	1,843	126	1,483	152	98	247	1,630	135	2,401	298
MNIST ReLU (V)	3	3	25	3	9	17	5	11	25	5	33	136
Private Set Intersection	95	9	558	59	1,049	104	22	186	591	96	2,639	302
Private Set Intersection (V)	1	2	1	2	1	8	1	8	2	4	2	122

Join, Count 102 and Count 10s see speedup from 1.3x to 1.1x. However, Longest 102 and Max distance suffer a slowdown of 0.5x. There is opportunity for vectorization in these benchmarks according to our analytical model, particularly, there is a large EQ operation that is vectorized, although a large portion of the loop cannot be vectorized. We observe that transformation to vectorized code increases multiplicative depth and, the negative effect of increased depth is more noticeable in a round-based protocol like GMW. We conjecture that MOTION performs optimizations over the non-vectorized loop body that decreases the depth; also, EQ is relatively inexpensive in Boolean GMW and BMR compared to ADD and MUL, which also de-emphasizes the benefit of vectorization. We propose a simple heuristic (although we do leave all the benchmarks in the table): if the transformation increases circuit depth beyond some threshold (e.g., more than 10% of the original circuit), we can reject the transformation. In some settings it may still be desirable to vectorize e.g., in data constrained environments, as vectorization reduces communication costs.




Figure 4: 2PC LAN Setting: Circuit evaluation. Time (Setup + Online) is in seconds, error bars show standard deviation.

As shown in Fig. 5a, vectorization reduces communication, up to 12x in GMW and 3x in BMR (see §8.2.3 for details). In summary, vectorization enables better packing and it impacts an interactive protocol like GMW more than a constant round protocol like BMR. Fig. 5c shows that vectorization reduces gates-count up to 480x in GMW, 450x in BMR. Consequently, in the highly vectorizable benchmarks, circuit generation time (see Fig. 5b) for vectorized circuits is a small fraction of the generation time for non-vectorized circuits (up to 200x less in GMW, 80x less in BMR). Online time and setup time are presented in Fig. 5d, and Fig. 5e respectively.

Figures 6, 7a, and 7b zoom into Biometric matching. For input size beyond $N=256$ the memory usage exceeds available memory and prevents circuit generation for the non-vectorized case. Vectorization improves all metrics. In circuit evaluation (see Fig. 6), BMR sees higher speedup (23x faster) compared to GMW (10x faster), while GMW sees faster circuit generation time at 45x lower (see Fig. 7b) compared to BMR’s which is 35x lower. Communication size reduction (see Fig. 7a) is higher for GMW (10x less) compared to BMR (2.5x less). The number of gates, communication size and circuit generation time remain the same in the WAN setting. As expected, Setup and Online times increase, as shown in Fig. 8.

For the 3PC, we observe that, as expected, evaluation time (Fig. 9a) is higher than the 2PC. This is a direct consequence of higher online time (Fig. 9e) for the GMW protocol. Online time for BMR remains

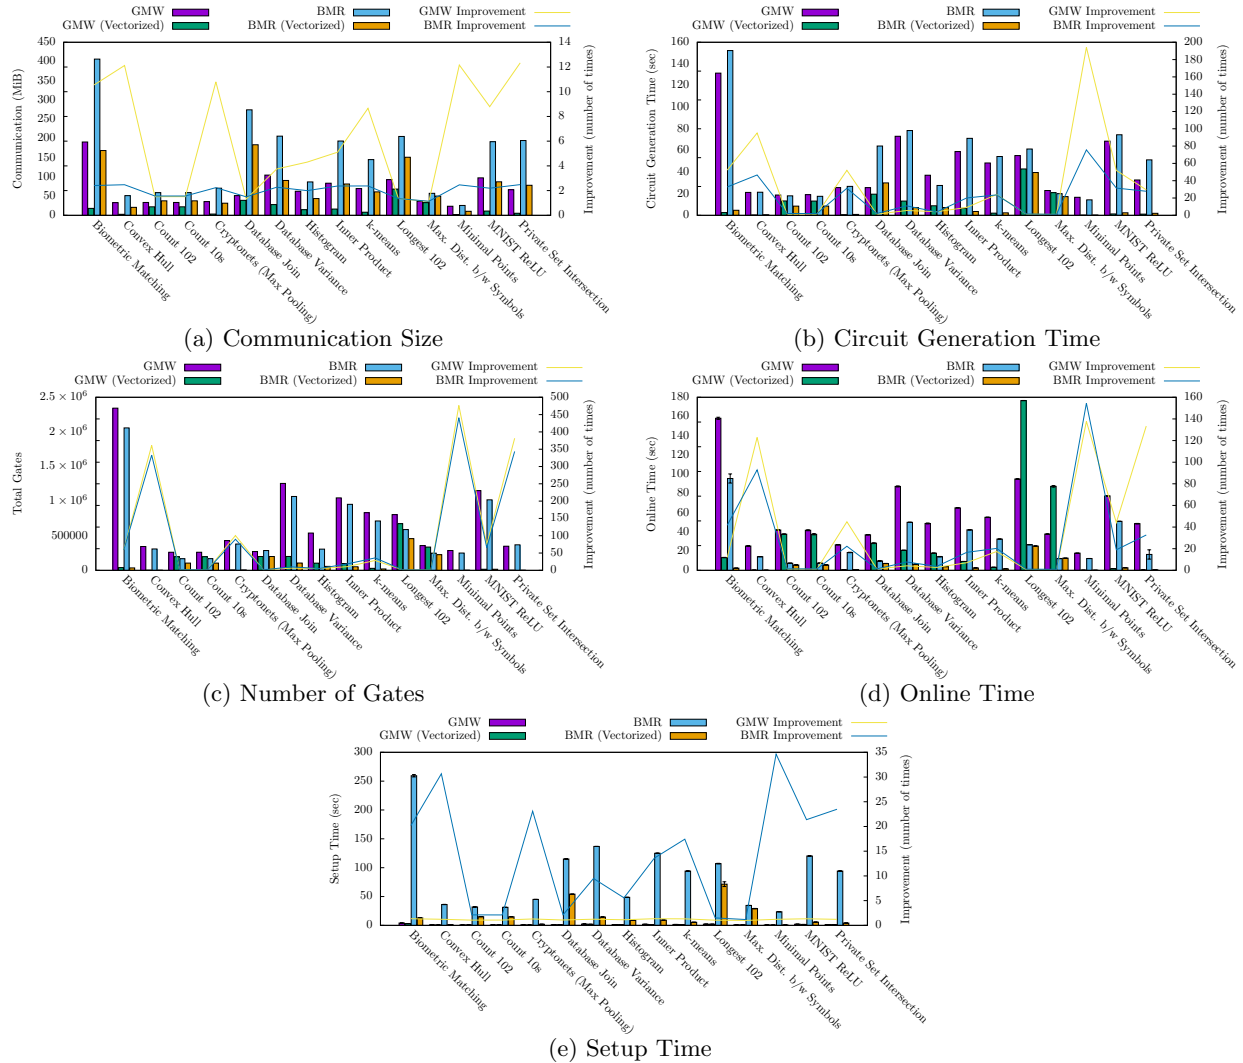


Figure 5: 2PC LAN setting: Various metrics for benchmark set, error bars (where present) show standard deviation.

roughly the same, which is expected because the online phase is essentially local in BMR. BMR suffers a slowdown in the setup phase (Fig. 9f) however. This is due to the circuit for 3 parties requiring more computation. We also include graphs for circuit generation time (Fig. 9c), gates count (Fig. 9d), and communication size (Fig. 9b) for 3PC. Circuit generation sees a slowdown in BMR for the reason mentioned above, communication size per channel and gates count remain the same. The experiments for 3PC confirm that adding more parties to an MPC increases resource requirements.

8.2.3 Analysis: Communication Size Reduction

As shown in communication size graph (Fig. 5a), vectorization results in reduced communication (fewer bits are transferred). This reduction is a result of more efficient data-packing at both (1) the application level (i.e. the MPC backend level), and (2) at the network level. The MPC backend needs to store/send metadata with each primitive/message so that it can be correctly decoded/consumed later. For example, a gate needs an identifier *gid*, a gate type *gtype*, incoming wire identifiers, etc. Say $size(gmeta)$ bits are needed

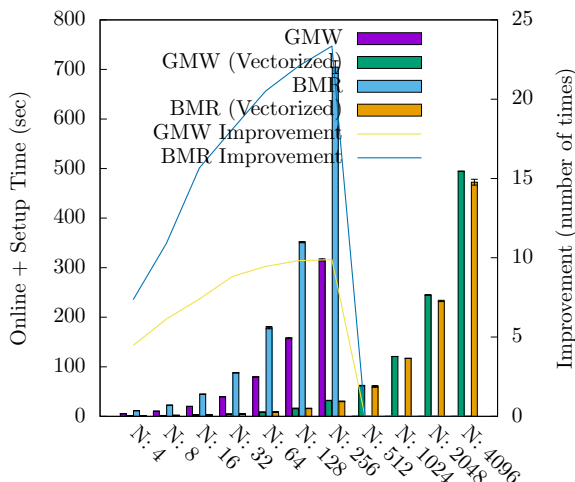


Figure 6: 2PC LAN setting: Biometric Matching evaluation time, the x-axis is database size.




Figure 7: 2PC LAN setting: Biometric Matching, the x-axis shows database size.

to store/send metadata for a single gate. Using one vectorized/SIMD gate instead of $(N + 1)$ non-vectorized gates saves $N \cdot \text{size}(gmeta)$ bits in memory/communication. Similarly, at the network level, each message needs a header h that contains routing and decoding information while the packet is in transit. Say, one (non-vectorized) interactive gate induces a payload p . This means, $\text{size}(h) + \text{size}(p)$ bits are sent to network for each (non-vectorized) interactive gate. Evaluation of N such gates translates to $N \cdot (\text{size}(h) + \text{size}(p))$ bits of communication. On the other hand, a vectorized gate that replaces these N gates is much cheaper, and requires only $\text{size}(h) + N \cdot \text{size}(p)$ bits of communication.

Concretely, let us consider an MPC backend implemented on Transport Control Protocol (TCP) over Internet Protocol (IP) i.e., the most common communication stack. Note that, for the sake of communication size comparison, the only difference between UDP and TCP is the smaller header size of 8 bytes⁵ in UDP compared to the at least 20 bytes⁶ in TCP. Both protocols are typically implemented over Internet Protocol and the header size of an IPv4 packet is 20 bytes⁷. In the Arithmetic GMW protocol, multiplication operation

⁵https://en.wikipedia.org/wiki/User_Datagram_Protocol#UDP_datagram_structure

⁶https://en.wikipedia.org/wiki/Transmission_Control_Protocol#TCP_segment_structure

⁷<https://en.wikipedia.org/wiki/IPv4#Header>

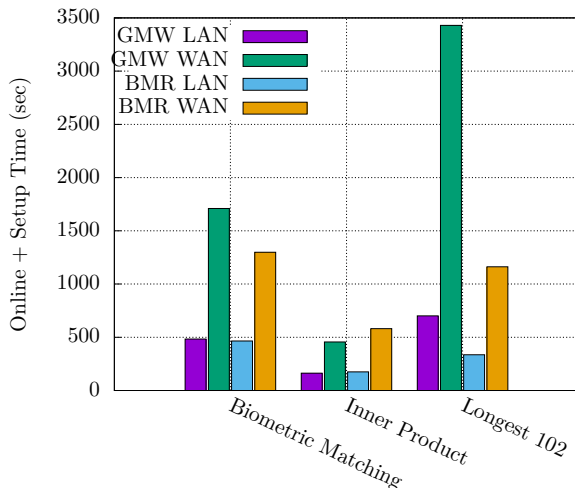


Figure 8: 2PC setting: Circuit evaluation time, LAN vs. WAN.

(MUL) is typically implemented using Beaver’s triples [Bea92]. This means that, in the online phase, all parties need to send 2ℓ bits to each other. If $\ell = 32$ bits, then, TCP payload is $2\ell = 64$ bits or 8 bytes. Considering that the Maximum Transmission Unit (MTU) is typically 1500 bytes, a TCP message may have payload of up to $(1500 - 20 - 20) = 1460$ bytes (the exact value is decided via the Maximum Segment Size (MSS) during TCP stack initialization, specification maximum is 65,496 bytes). Meaning, a vectorized gate replacing $1460/8 \approx 180$ non-vectorized gates could be sent in a single 1500 byte message rather than $182 \cdot (20 + 20 + 8) = 8,640$ bytes otherwise required. Similar reasoning applies to interactive gates in boolean GMW and, while exact improvement depends on the implementation details; packing data reduces both the memory and communication footprint regardless of the underlying MPC backend (as long as it supports vectorized gates).

In the case of BMR, the entire circuit can be packed as one payload and sent using a few TCP packets. Therefore under-utilization of network’s payload-capacity is not an issue. At the application (MPC backend) level however, inefficient packing is still an issue. For example, MOTION uses 64 bits (8 bytes) for gate identifiers. A vectorized gate that replaces 128 non-vectorized gates, requires only one gate identifier i.e. 8 bytes instead 1,024 bytes required for 128 identifiers. Thus, vectorization reduces the size of the circuit. This, in turn, reduces payload for the network and means that fewer TCP packets need to be sent, thereby saving on TCP/IP metadata that would have been needed for additional packets.

8.2.4 Comparison with MOTION-native Inner Product

We compared our results on Inner Product (Fig. 10a, Fig. 10b and Fig. 10c) with the manually SIMD-ified one distributed with MOTION source. We were surprised that we were an order of magnitude slower in Boolean GMW as our circuit ran a significantly larger number of communication rounds. Upon investigation, it turns out that the vectorized multiplications are the same, however, our addition loop incurs significant cost (ADD is non-local and expensive in Boolean GMW). The MOTION-native loop runs

```
1 result += mult_unsimdified[i];
```

while our compiler-generated loop runs

```
1 result[i] = result[i-1] + mult_unsimdified[i];
```

Recall that the scalar expansion is an artifact of our vectorization. We rewrote the accumulation (manually, for testing purposes) and that lead to the same running time as the MOTION native code.




Figure 9: 3PC LAN setting: Various metrics for benchmark set, error bars show standard deviation.

MOTION’s compiler performs analysis that informs circuit generation and the example illustrates the power of the analysis. In the above example, MOTION overloads the `+=` operator to perform divide-and-conquer accumulation in $O(\log(N))$ rounds.

It is unrealistic to expect that MOTION’s static analysis will detect the associative accumulation in the scalar expansion code above. One reason is that MOTION’s analysis is at source or AST level and such static analysis is difficult. However, our investigation showed that not only MOTION does not optimize

```
1 result[i] = result[i-1] + mult_unsimplified[i];
```

it does not optimize the simpler accumulation:

```
1 result = result + mult_unsimplified[i];
```

We conjecture that MPC-IR, a straight-forward representation, will not only enable detection of general associative loops, but also allow for program synthesis to increase opportunities for divide-and-conquer parallelization [FN21]; as the problem is non-trivial, particularly the interaction of divide-and-conquer with vectorization and mixing, we leave it for future work.

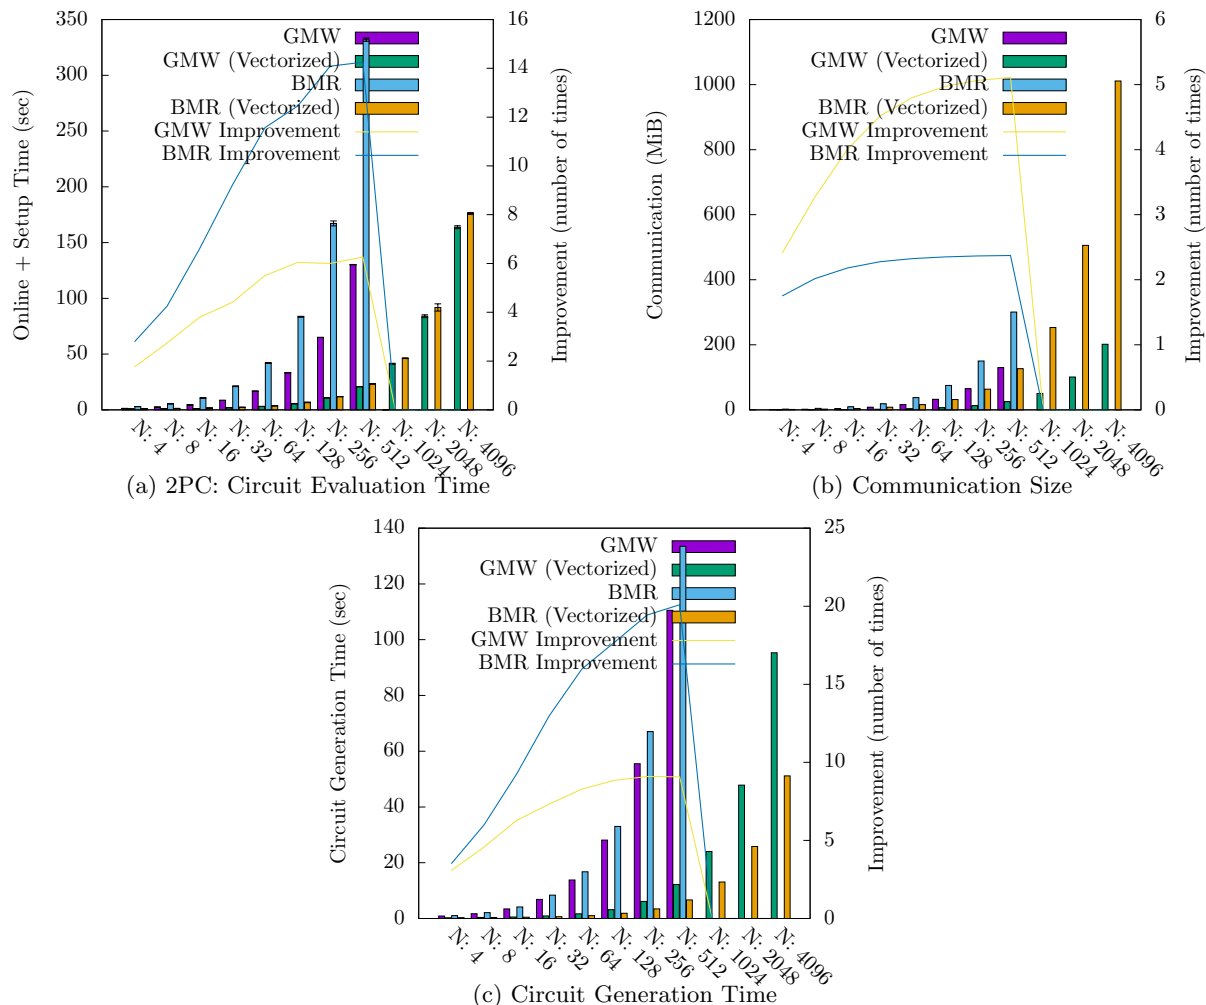


Figure 10: 2PC LAN setting: Inner Product, x-axis shows vector size.

8.3 MP-SPDZ Experiments

The MOTION compiler does not automatically apply vectorization; hence our MOTION experiments compare our vectorization optimization to “no optimization” (the iterative version of the benchmark) and, as expected, demonstrate high gains. Instead, MP-SPDZ does use vectorization but at a lower level than ours [Kel20, Section 6.1] — the MP-SPDZ compiler unrolls the loop and merges (same) instructions from different iterations that may occur in parallel into a vector; the VM runs the vectorized instruction. Thus, it is reasonable to ask if our backend-independent optimization is compatible with MP-SPDZ and whether or not it still yields improvements. Our relevant experiments discussed below answer both these questions to the affirmative.

8.3.1 Experiment Setup

We run a representative subset of our MOTION experiments in the LAN setting (c.f. LAN configuration above). We remark that having already demonstrated the benefits of our optimization on a range of experiments with MOTION, the goal of our MP-SPDZ experiments is to address the above questions. Therefore, we chose a subset of benchmarks to run with MP-SPDZ that give a clear answer to the questions above

Table 7: MP-SPDZ Compilation/execution stats, Security Parameter = 40, Bit length=32, 2PC 10Gbps LAN setting. Compilation and run times are in seconds, Peak memory (Memory) and Communication (Comm.) are in MBs, Vectorized benchmarks have (V) in their name.

Benchmark	Arithmetic										Binary									
	MASCOT					SOHO					Semi-bin					Semi-BMR				
	Compilation	Memory	Int Triples	VM Rounds	Run Time	Comm.	Run Time	Comm.	Compilation	Memory	Bit Triples	VM Rounds	Run Time	Comm.	Run Time	Comm.				
Biometric Matching	15.66	587.25	64,512	4,098	21.06	2,322.25	4.73	37.39	824.90	22,100.97	3,300,864	3,628	1.11	87.45	373.32	92,928				
Biometric Matching (V)	16.16	605.53	64,512	4,098	21.15	2,322.25	4.72	37.39	32.49	1,014.21	3,300,864	3,651	0.67	41.73	369.92	92,928				
Convex Hull	3.77	109.37	1,006,848	75	327.71	36,706.00	66.54	555.19	157.08	4,297.98	536,576	74	0.20	14.37	60.82	15,109				
Convex Hull (V)	3.78	148.31	1,006,848	76	328.75	36,706.00	66.54	555.19	18.64	612.45	665,536	74	0.18	11.85	75.55	18,737				
Count 10s	0.84	44.32	98,816	1,032	45.80	5,054.18	10.31	75.41	35.02	1,058.98	117,760	1,065	0.20	7.67	13.40	3,316				
Count 10s (V)	0.94	51.88	98,816	1,545	45.82	5,054.18	10.27	75.41	26.64	784.60	117,760	2,086	0.24	6.91	13.44	3,326				
Count 102	0.81	42.38	98,816	1,537	45.90	5,054.18	10.32	75.41	34.91	1,040.84	117,760	1,065	0.18	7.66	13.34	3,316				
Count 102 (V)	0.94	51.36	98,816	15,37	45.86	5,054.18	10.33	75.41	26.64	784.74	117,760	2,087	0.24	6.91	13.75	3,316				
Cryptonets (Max Pooling)	1.85	74.26	371,712	25	120.66	13,562.30	24.66	205.80	83.75	2,526.37	291,840	22	0.15	10.46	35.41	8,218				
Cryptonets (Max Pooling) (V)	2.04	88.11	371,712	25	121.37	13,562.30	24.68	205.80	33.01	1,140.89	291,840	22	0.13	7.41	34.30	8,218				
Database Join	8.62	244.82	786,432	12	365.42	40,296.70	78.17	558.06	176.04	5,812.58	774,144	10	0.29	23.97	93.54	21,798				
Database Join (V)	9.56	281.16	786,621	12	365.34	40,316.20	78.36	561.57	150.76	4,973.32	774,237	10	0.27	22.47	91.91	21,801				
Inner Product	0.14	24.18	512	2	0.20	19.54	0.47	5.51	170.52	4,852.62	669,696	551	0.31	22.27	76.38	18,856				
Inner Product (V)	0.25	28.77	512	2	0.21	19.54	0.47	5.51	19.74	620.45	669,696	568	0.30	22.95	75.93	18,856				
Longest 102	16.32	594.25	161,792	4,620	66.87	7,337.38	14.73	107.25	51.51	1,463.79	168,960	4,110	0.36	8.49	18.87	4,757.77				
Longest 102 (V)	16.12	601.64	161,792	2,049	67.28	7,337.38	14.74	107.25	39.16	1,159.94	168,960	5,132	0.39	7.70	19.27	4,757.77				
Max Dist. btw Symbols	15.40	579.75	94,720	4,105	36.33	3,961.44	7.84	59.57	41.07	1,206.88	131,584	3,600	0.33	7.89	14.89	3,707				
Max Dist. btw Symbols (V)	15.21	573.94	94,720	4,616	36.02	3,961.44	7.97	59.57	36.70	1,122.38	131,584	3,600	0.30	7.64	15.00	3,707				
Minimal Points	3.35	103.69	991,360	74	320.47	36,159.60	65.66	545.87	166.72	4,289.05	540,672	73	0.20	14.43	61.61	15,224				
Minimal Points (V)	3.77	158.95	991,360	74	321.88	36,159.60	65.56	545.87	21.21	621.64	667,648	73	0.18	11.79	75.75	18,796				
MNIST ReLU	6.84	212.63	991,232	9	321.28	36,159.70	65.95	546.00	266.69	7,632.11	778,240	8	0.34	24.07	99.70	21,916				
MNIST ReLU (V)	6.08	233.02	991,232	9	320.73	36,159.70	65.47	546.00	94.32	3,163.99	778,240	8	0.26	15.94	101.66	21,916				
Private Set Intersection	11.01	288.23	1,048,704	137	485.46	53,722.30	104.53	745.79	118.31	3,784.08	528,384	135	0.19	14.23	59.39	14,877				
Private Set Intersection (V)	9.36	325.30	1,048,704	137	487.23	53,722.30	104.24	745.79	23.28	757.01	528,384	135	0.13	6.24	59.33	14,877				

while also demonstrating the delicate points of applying our optimizer on a back-end which already vectorizes (MP-SPDZ) vs. one that does not (MOTION). In particular, we focus on 2PC with semi-honest adversary for protocols following all three MPC paradigms (Boolean/Arithmetic GMW and BMR). As an extra point, to demonstrate that our optimization can also be applied to maliciously secure protocols, we run 2PC MP-SPDZ with the MASCOT protocol. Concretely, we run MP-SPDZ in the Arithmetic setting with the MASCOT and SOHO protocols, and in the Binary setting with the Semi-bin and Semi-BMR protocols. We run with datasets that trigger program complexity of at least $O(10^3)$. Our results are discussed below.

8.3.2 Results and Analysis

Table 7 summarizes our experiments. Vectorization has virtually no impact in the Arithmetic setting. In contrast, it has significant impact on compilation in the Binary setting, reducing compilation time from about 11% on Max Distance to 25x on Biometric. It has significant impact on memory footprint as well (22x in Biometric). We conjecture that the above discrepancy is due to the following reason: In the Arithmetic setting vectorization of data (typically integer arrays) through the API is as costly as merging of data and instructions by the MP-SPDZ compiler. Granularity of instructions is larger compared to the Binary setting, leading to fewer nodes and edges and lower complexity of the dependence graph and dependence analysis. In contrast, in the Binary setting’s finer granularity of instructions (e.g., multiplication of two integers is expressed in terms of multiple AND and XOR instructions) leading to a larger and denser graph and significantly more complex dependence analysis. Performing dependence analysis a priori on the very short MPC-IR, which is what our vectorization analysis does, informs the MP-SPDZ compiler and reduces analysis time and memory footprint.

We note in passing, that our experiments demonstrate that vectorization slightly increases the number of rounds on the regular expression programs (e.g., Count 10s). This is consistent with the MOTION results where we observed little improvement or even slowdown. Indeed, our vectorization algorithm breaks the single loop into several loops separating smaller loops with a vectorized operation, typically equality; we conjecture that the vectorized operation serves as a barrier preventing the MP-SPDZ compiler from merging instructions it would otherwise merge in the non-vectorized single-loop program (resulting in improved round complexity). Nevertheless, our LAN experiments show that the above increase in rounds is not reflected as slowdown on running time, so the significant compilation savings make our optimization worthwhile.

9 Conclusion and Future Work

We presented a formalization of the MPC-IR intermediate language followed by a specific backend-independent optimization at the level of MPC-IR: novel SIMD-vectorization. We demonstrated that vectorization has significant impact on performance and showcased the backend-independent nature of our optimization by applying it to two mainstream and parameterizable MPC frameworks, namely MOTION and MP-SPDZ.

We are excited about the opportunities for future work — integration with protocol mixing, divide-and-conquer reasoning and parallelization, as well as inter-procedural context-sensitive analysis at the level of MPC-IR will improve MPC programmability and efficiency.

10 Acknowledgements

We thank the reviewers for their constructive feedback. RPI authors are supported by NSF grants #1814898 and #2232061. Purdue authors are supported by the Algorand Centers of Excellence program managed by Algorand Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Algorand Foundation.

References

- [Aca+21] Cosku Acay et al. “Viaduct: an extensible, optimizing compiler for secure distributed programs.” In: *ACM PLDI 2021*. Ed. by Stephen N. Freund and Eran Yahav. ACM, June 2021, pp. 740–755.
- [AK87] Randy Allen and Ken Kennedy. “Automatic Translation of Fortran Programs to Vector Form.” In: *ACM Trans. Program. Lang. Syst.* 9.4 (1987), pp. 491–542.
- [AN88] Alexander Aiken and Alexandru Nicolau. “Optimal Loop Parallelization.” In: *ACM PLDI 1988*. Ed. by Richard L. Wexelblat. ACM, June 1988, pp. 308–317.
- [Ara+18] Toshinori Araki et al. “Generalizing the SPDZ Compiler For Other Protocols.” In: *ACM CCS 2018*. Ed. by David Lie et al. ACM Press, Oct. 2018, pp. 880–895.
- [ASU86] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley series in computer science / World student series edition. Addison-Wesley, 1986.
- [Bea92] Donald Beaver. “Efficient Multiparty Protocols Using Circuit Randomization.” In: *CRYPTO’91*. Ed. by Joan Feigenbaum. Vol. 576. LNCS. Springer, Heidelberg, Aug. 1992, pp. 420–432.
- [Ben+10] Mohamed-Walid Benabderrahmane et al. “The Polyhedral Model Is More Widely Applicable Than You Think.” In: *Compiler Construction, CC 2010*. Ed. by Rajiv Gupta. Vol. 6011. Springer, 2010, pp. 283–303.
- [BG11] Marina Blanton and Paolo Gasti. “Secure and Efficient Protocols for Iris and Fingerprint Identification.” In: *ESORICS 2011*. Ed. by Vijay Atluri and Claudia Díaz. Vol. 6879. LNCS. Springer, Heidelberg, 2011, pp. 190–209.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. “Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract).” In: *20th ACM STOC*. ACM Press, May 1988, pp. 1–10.
- [BK15] Niklas Büscher and Stefan Katzenbeisser. “Faster Secure Computation through Automatic Parallelization.” In: *24th USENIX Security Symposium (USENIX Security 15)*. USENIX Association, Aug. 2015, pp. 531–546.
- [BLW08] Dan Bogdanov, Sven Laur, and Jan Willemson. “Sharemind: A Framework for Fast Privacy-Preserving Computations.” In: *ESORICS 2008*. Ed. by Sushil Jajodia and Javier López. Vol. 5283. LNCS. Springer, Heidelberg, Oct. 2008, pp. 192–206.
- [BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. “The Round Complexity of Secure Protocols (Extended Abstract).” In: *22nd ACM STOC*. ACM Press, May 1990, pp. 503–513. DOI: [10.1145/100216.100287](https://doi.org/10.1145/100216.100287).
- [BNP08] Assaf Ben-David, Noam Nisan, and Benny Pinkas. “FairplayMP: a system for secure multi-party computation.” In: *ACM CCS 2008*. Ed. by Peng Ning, Paul F. Syverson, and Somesh Jha. ACM Press, Oct. 2008, pp. 257–266.
- [Bog+09] Peter Bogetoft et al. “Secure Multiparty Computation Goes Live.” In: *FC 2009*. Ed. by Roger Dingledine and Philippe Golle. Vol. 5628. LNCS. Springer, Heidelberg, Feb. 2009, pp. 325–343.
- [Bra+22] Lennart Braun et al. “MOTION: A Framework for Mixed-Protocol Multi-Party Computation.” In: *ACM TOPS 25.2* (May 2022), pp. 1–35.
- [Bra+23] Lennart Braun et al. *FUSE – Flexible File Format and Intermediate Representation for Secure Multi-Party Computation*. Cryptology ePrint Archive, Paper 2023/563. 2023.
- [Büs+18] Niklas Büscher et al. “HyCC: Compilation of Hybrid Protocols for Practical Secure Computation.” In: *ACM CCS 2018*. Ed. by David Lie et al. ACM Press, Oct. 2018, pp. 847–861.
- [Büs18] Niklas Büscher. “Compilation for More Practical Secure Multi-Party Computation.” PhD thesis. Darmstadt University of Technology, Germany, 2018.

- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. “Multiparty Unconditionally Secure Protocols (Extended Abstract).” In: *20th ACM STOC*. ACM Press, May 1988, pp. 11–19.
- [Che+23] Edward Chen et al. *Silph: A Framework for Scalable and Accurate Generation of Hybrid MPC Protocols*. Cryptology ePrint Archive, Paper 2023/060. <https://eprint.iacr.org/2023/060>. 2023. DOI: [10.1109/SP46215.2023.00103](https://doi.org/10.1109/SP46215.2023.00103). URL: <https://eprint.iacr.org/2023/060>.
- [COS19] KU Leuven COSIC. *SCALE-MAMBA*. 2019. URL: <https://github.com/KULeuven-COSIC/SCALE-MAMBA>.
- [Cyt+91] Ron Cytron et al. “Efficiently Computing Static Single Assignment Form and the Control Dependence Graph.” In: *ACM Trans. Program. Lang. Syst.* 13.4 (1991), 451?–490. ISSN: 0164-0925.
- [Dam+12] Ivan Damgård et al. “Multiparty Computation from Somewhat Homomorphic Encryption.” In: *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Vol. 7417. Lecture Notes in Computer Science. Springer, 2012, pp. 643–662. DOI: [10.1007/978-3-642-32009-5_38](https://doi.org/10.1007/978-3-642-32009-5_38). URL: https://doi.org/10.1007/978-3-642-32009-5_38.
- [Dam+13] Ivan Damgård et al. “Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits.” In: *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings*. Ed. by Jason Crampton, Sushil Jajodia, and Keith Mayes. Vol. 8134. Lecture Notes in Computer Science. Springer, 2013, pp. 1–18. DOI: [10.1007/978-3-642-40203-6_1](https://doi.org/10.1007/978-3-642-40203-6_1). URL: https://doi.org/10.1007/978-3-642-40203-6_1.
- [Dow+16] Nathan Dowlin et al. “CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy.” In: *ICML 2016*. New York, NY, USA: JMLR.org, June 2016, pp. 201–210.
- [DSZ15] Daniel Demmler, Thomas Schneider, and Michael Zohner. “ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation.” In: *NDSS 2015*. The Internet Society, Feb. 2015.
- [Dup+19] Dmitry Duplyakin et al. “The Design and Operation of CloudLab.” In: *Proceedings of the USENIX Annual Technical Conference (ATC)*. July 2019, pp. 1–14.
- [Esc+20] Daniel Escudero et al. “Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits.” In: *CRYPTO 2020, Part II*. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12171. LNCS. Springer, Heidelberg, Aug. 2020, pp. 823–852.
- [Fan+22] Vivian Fang et al. *CostCO: An automatic cost modeling framework for secure multi-party computation*. Cryptology ePrint Archive, Report 2022/332. <https://eprint.iacr.org/2022/332>. 2022.
- [FN21] Azadeh Farzan and Victor Nicolet. “Phased synthesis of divide and conquer programs.” In: *ACM PLDI 2021*. Ed. by Stephen N. Freund and Eran Yahav. ACM, July 2021, pp. 974–986.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. “How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority.” In: *19th ACM STOC*. Ed. by Alfred Aho. ACM Press, May 1987, pp. 218–229.
- [Has+19] Marcella Hastings et al. “SoK: General Purpose Compilers for Secure Multi-Party Computation.” In: *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2019, pp. 1220–1237.
- [Hel+21] Tim Heldmann et al. “LLVM-Based Circuit Compilation for Practical Secure Computation.” In: *Applied Cryptography and Network Security*. Ed. by Kazue Sako and Nils Ole Tippenhauer. Cham: Springer International Publishing, 2021, pp. 99–121.
- [IMZ19] Muhammad Ishaq, Ana L. Milanova, and Vassilis Zikas. “Efficient MPC via Program Analysis: A Framework for Efficient Optimal Mixing.” In: *ACM CCS 2019*. Ed. by Lorenzo Cavallaro et al. ACM Press, Nov. 2019, pp. 1539–1556.

- [JW05] Geetha Jagannathan and Rebecca N. Wright. “Privacy-Preserving Distributed k-Means Clustering over Arbitrarily Partitioned Data.” In: *ACM CKDDM*. Chicago, IL, USA: ACM, 2005, pp. 593–599.
- [Kar15] Ralf Karrenberg. *Automatic SIMD Vectorization of SSA-based Control Flow Graphs*. Springer, 2015. ISBN: 978-3-658-10112-1.
- [Kel20] Marcel Keller. “MP-SPDZ: A Versatile Framework for Multi-Party Computation.” In: *ACM CCS 2020*. Ed. by Jay Ligatti et al. ACM Press, Nov. 2020, pp. 1575–1590.
- [KOS16] Marcel Keller, Emmanuela Orsini, and Peter Scholl. “MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer.” In: *ACM CCS 2016*. Ed. by Edgar R. Weippl et al. ACM Press, Oct. 2016, pp. 830–842.
- [KSS13] Marcel Keller, Peter Scholl, and Nigel P. Smart. “An architecture for practical actively secure MPC with dishonest majority.” In: *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS’13, Berlin, Germany, November 4-8, 2013*. Ed. by Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung. ACM, 2013, pp. 549–560. DOI: [10.1145/2508859.2516744](https://doi.org/10.1145/2508859.2516744). URL: <https://doi.org/10.1145/2508859.2516744>.
- [Liu+15] Chang Liu et al. “OblivM: A Programming Framework for Secure Computation.” In: *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2015, pp. 359–376.
- [MB08] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. “Z3: An Efficient SMT Solver.” In: *TACAS 2008*. Ed. by C. R. Ramakrishnan and Jakob Rehof. Vol. 4963. Springer, Apr. 2008, pp. 337–340.
- [Moo+16] Benjamin Mood et al. “Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation.” In: *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. 2016, pp. 112–127.
- [MR18] Payman Mohassel and Peter Rindal. “ABY³: A Mixed Protocol Framework for Machine Learning.” In: *ACM CCS 2018*. Ed. by David Lie et al. ACM Press, Oct. 2018, pp. 35–52.
- [MZ17] Payman Mohassel and Yupeng Zhang. “SecureML: A System for Scalable Privacy-Preserving Machine Learning.” In: *2017 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2017, pp. 19–38.
- [Nay+15] Kartik Nayak et al. “GraphSC: Parallel Secure Computation Made Easy.” In: *2015 IEEE Symposium on Security and Privacy*. 2015, pp. 377–394.
- [NK14] Tobias Nipkow and Gerwin Klein. *Concrete Semantics: With Isabelle/HOL*. Heidelberg, Germany: Springer, 2014. ISBN: 3319105418.
- [OBW20] Alex Ozdemir, Fraser Brown, and Riad S. Wahby. *Unifying Compilers for SNARKs, SMT, and More*. Cryptology ePrint Archive, Report 2020/1586. <https://eprint.iacr.org/2020/1586>. 2020.
- [Pat+21] Arpita Patra et al. “ABY2.0: Improved Mixed-Protocol Secure Two-Party Computation.” In: *USENIX Security 2021*. Ed. by Michael Bailey and Rachel Greenstadt. USENIX Association, Aug. 2021, pp. 2165–2182.
- [RHH14a] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. “Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations.” In: *2014 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2014, pp. 655–670.
- [RHH14b] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. “Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations.” In: *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*. 2014, pp. 655–670.
- [RSH19] Aseem Rastogi, Nikhil Swamy, and Michael Hicks. “Wys*: A DSL for Verified Secure Multiparty Computations.” In: *Principles of Security and Trust*. Ed. by Flemming Nielson and David Sands. Cham: Springer International Publishing, 2019, pp. 99–122. ISBN: 978-3-030-17138-4.

- [Sco09] Michael L. Scott. *Programming Language Pragmatics (3. ed.)* Academic Press, 2009.
- [Son+15] Ebrahim M. Songhori et al. “TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits.” In: *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2015, pp. 411–428.
- [Vaz10] Vijay V. Vazirani. *Approximation Algorithms*. Heidelberg, Germany: Springer, 2010. ISBN: 3642084699.
- [VC03] Jaideep Vaidya and Chris Clifton. “Privacy-Preserving k -Means Clustering over Vertically Partitioned Data.” In: Washington, D.C.: ACM, 2003, pp. 206–215.
- [Yao82] Andrew Chi-Chih Yao. “Protocols for Secure Computations (Extended Abstract).” In: *23rd FOCS*. IEEE Computer Society Press, Nov. 1982, pp. 160–164.
- [ZE15] Samee Zahur and David Evans. *Obliv-C: A Language for Extensible Data-Oblivious Computation*. Cryptology ePrint Archive, Report 2015/1153. 2015.
- [ZSB13] Yihua Zhang, Aaron Steele, and Marina Blanton. “PICCO: a general-purpose compiler for private distributed computation.” In: *ACM CCS 2013*. Ed. by Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung. ACM Press, Nov. 2013, pp. 813–826.

A Background

A.1 Static Single Assignment (SSA) Form

Static Single Assignment (SSA) form [Cyt+91] is a compiler technique that transforms the control-flow graph so that there is at most a single (static) definition of each variable. The static single assignment program gives rise to a sparse dependence graph, which facilitates compiler analysis and optimizations. There are three cases: (1) straight-line code, (2) if-then-else statements, and (3) loops.

The following example illustrates straight-line code:

<pre> 1 x = 1 2 y = x 3 x = 2 4 z = x </pre>	<pre> 1 x0 = 1 2 y = x0 3 x1 = 2 4 z = x1 </pre>
--	--

(a) Source (b) SSA form

The transformation indexes each definition of variable x and adjusts each use of x to the corresponding indexed variable. There is no straight-line code in our MPC benchmarks that leads to the introduction of indexed variables as in the above example; essentially all indexing is due to if-then-else and loops.

If-then-else statements and loops give rise to ϕ (PHI) nodes which merge values along different control flow paths. The example illustrates if-then-else:

<pre> 1 if x > 0: 2 y = a 3 else: 4 y = b 5 v = y </pre>	<pre> 1 if x > 0: 2 y0 = a 3 else: 4 y1 = b 5 y2 = PHI(y0,y1) 6 v = y2 </pre>	<pre> 1 c = x > 0: 2 y0 = a 3 y1 = b 4 y2 = MUX(c,y0,y1) 5 v = y2 </pre>
---	--	---

(a) Source (b) SSA form (c) MPC-IR

Again, the compiler indexes the two definitions of y and creates the ϕ -node at the merge point of the if-then-else statement. Variable $y2$ denotes either $y0$ (if control took the then-branch) or $y1$ (if control took the else-branch), that is, v is assigned either $y0$ or $y1$. This example illustrates the usefulness of SSA for transformations into MPC. The ϕ -node roughly corresponds to a MUX and the if-then-else gives rise to the linear MPC code shown in the rightmost column of the table above. Cytron’s algorithm [Cyt+91] for computing SSA form is appealing because of it guarantees a minimal number of ϕ -nodes, and therefore a minimal number of MUX nodes in our client, and it allows for a systematic translation from source to linear MPC. A key difficulty is that Cytron’s algorithm may create ϕ -nodes with 3 or more arguments corresponding to different conditional branches, and it does not keep track of the condition(s) that give rise to the ϕ -node.

This required adapting the algorithm to the setting of MPC, specifically, restricting ϕ nodes to 2-argument nodes, and keeping track of the corresponding condition.

Finally, the example below illustrates loops:

<pre> 1 x = 10 2 for i in range(N): 3 x = x + 1 </pre>	<pre> 1 x0 = 10 2 for i in range(N): 3 x1 = PHI(x0,x2) 4 x2 = x1 + 1 </pre>	<pre> 1 x0 = 10 2 for i in range(N): 3 x1 = PHI(x0,x2) 4 x2 = x1 + 1 </pre>
(a) Source	(b) SSA form	(c) MPC-IR

As in the case for if-then-else there is a ϕ -node for x at the loop header, which is a control merge point. The ϕ -node merges control from the forward edge into the loop and the back edge. In MPC-IR these ϕ -nodes remain as what we call *pseudo-phi* as they disappear when the backend unrolls the loop. They are important as they help retain the compact representation of the circuit into MPC-IR.

A.2 Automatic Vectorization

Automatic vectorization is a longstanding problem in High-Performance Computing and compilers and there is decades of research and thousands of papers in the area. In our work, we define the problem over MPC-IR, which is linear, and therefore our problem is simpler than the general vectorization problem over source code with conditionals.

To introduce the problem, consider the following example from [AK87]:

```

1 for i in range(100):
2   X[i] = X[i] + Y[i]

```

This code can be vectorized and executed “at once”:

```

1 X[1: 100] = X[1 : 100] + Y[1: 100]

```

It can be vectorized because there are no loop-carried dependencies, i.e., no element of X depends on an element of X written in an earlier iteration.

In contrast, the code below cannot be vectorized:

```

1 for i in range(100):
2   X[i+3] = X[i] + Y[i]

```

This is because, say $X[6]$, depends on $X[3]$ which is computed in an earlier iteration.

We apply ideas from classical work on vectorization such as Allen and Kennedy’s work on automatic loop vectorization [AK87]. There are the following steps, starting from the innermost loop and repeating:

1. Compute the dependence graph, most importantly, loop-carried dependencies.
2. Identify statements that do not depend on a previous iteration of the loop and therefore can be computed “at once” rather than one by one.
3. Move those statements in a vectorized assignment outside of the loop.

The key step is the dependence analysis — a more precise analysis will remove more potential loop-carried dependencies and improve opportunities for vectorization. There is a significant body of work on dependence analysis and vectorization in the high performance literature. In our work we devise an analysis specific to the MPC-IR and its syntax and semantic restrictions. In future work we will relax restrictions and study more precise dependence analysis.