High-precision Secure Computation of Satellite Collision Probabilities[*]

Brett Hemenway[†]   Steve Lu[‡]   Rafail Ostrovsky[§]   William Welser IV[¶]

Abstract

The costs of designing, building, launching and maintaining satellites make satellite operators extremely motivated to protect their on-orbit assets. Unfortunately, privacy concerns present a serious barrier to coordination between different operators. One obstacle to improving safety arises because operators view the trajectories of their satellites as private, and refuse to share this private information with other operators. Without data-sharing, preventing collisions between satellites becomes a challenging task.

A 2014 report from the RAND Corporation proposed using cryptographic tools from the domain of secure Multiparty Computation (MPC) to allow satellite operators to calculate collision probabilities (conjunction analyses) without sharing private information about the trajectories of their satellites.

In this work, we report on the design and implementation of a powerful new MPC framework for high-precision arithmetic on real-valued variables in a two-party setting where, unlike previous works, there is no honest majority, and where the players are not assumed to be semi-honest. We show how to apply this new solution in the domain of securely computing conjunction analyses. Our solution extends existing protocols, in particular the integer-based Goldreich-Micali-Wigderson (GMW) protocol, whereby we use combine and optimize GMW with Garbled Circuits (GC). We prove security of our protocol in the two party, semi-honest setting, assuming only the existence of one-way functions and Oblivious Transfer (the OT-hybrid model). The protocol allows a pair of satellite operators to compute the probability that their satellites will collide without sharing their underlying private orbital information. Techniques developed in this paper would potentially have a wide impact on general secure numerical analysis computations. We also show how to strengthen our construction with standard arithmetic message-authentication-codes (MACs) to enforce honest behavior beyond the semi-honest setting.

Computing a conjunction analysis requires numerically estimating a complex double integral to a high degree of precision. The complexity of the calculation, and the possibility of numeric instability presents many challenges for MPC protocols which typically model calculations as simple (integer) arithmetic or binary circuits.

Our secure numerical integration routines are extremely stable and efficient, and our secure conjunction analysis protocol takes only a few minutes to run on a commodity laptop.

1 Introduction

There are currently more than 1300 active satellites orbiting the earth [UCS15], and this number is growing rapidly. Technological improvements have drastically reduced the barriers to building, launching and maintaining satellites in orbit, and consequently the number of different governments and private corporations maintaining active satellites is growing at an increasing rate (see Figure 1). As the number of satellites and operators grows, the problem of coordinating operations between the different operators becomes more challenging.

Number of Different Satellite Operators


Number of Active Satellites

Figure 1: The number of distinct satellite operators is increasing dramatically. In January 1995 there were only 21 different operators, while in January 2015 there were 357 according to the Union of Concerned Scientists [UCS15]. During the same time frame, the number of active satellites increased from 24 to over 1200.

In February 2009 the telecommunication satellite Iridium-33 collided with the Russian Kosmos 2251 in Low Earth Orbit (LEO). Their relative velocity was over 20,000 miles per hour, and both satellites were instantly destroyed and more than 1000 debris chunks over 4 inches in diameter were created [VO09].

Preventing future collisions requires coordination between a growing number of satellite operators. Unfortunately, privacy concerns present a barrier to cooperation, and in fact satellite operators view the precise trajectories of their on-orbit assets as private information. The Space Surveillance Network (SSN) managed by U.S. Strategic Command (USSTRATCOM) currently tracks more than 20,000 orbital objects with diameters greater than 10 cm. These tracking data, obtained by ground-based telescopes, are a valuable source of information, but they are too low-fidelity to calculate the probability that two active satellites will collide. Calculating collision probabilities (termed a "conjunction analyses") requires high-fidelity data from the satellites' on-board instrumentation – and these high-fidelity data are only available to the satellite's operator.

Thus we are in a situation where satellite operators want to keep their high-fidelity orbital information private, yet these are exactly the data needed to compute conjunction analyses and prevent further collisions. To overcome this obstacle, some operators have contracted the services of a trusted outside party (e.g. AGI). Operators then share their private data with the trusted party, the trusted party performs the conjunction analyses, and issues warnings if the collision probability exceeds a given threshold.

Trusted third parties do not provide a perfect solution, however, as many stakeholders cannot agree on a single trusted party, and even when such a mutually trusted party can be found, they can command a high price for their services. Secure multiparty computation (MPC) provides a cryptographic alternative to a trusted third party. Using MPC to securely compute satellite conjunction analyses was first proposed as a potential solution by Hemenway, Welser and Baiocchi [HWIB14]. Following that proposal, the problem of securely computing conjunction analyses was incorporated into DARPA's PROCEED program as a potential use-case for MPC technology, and received further media attention [HW15]. A three-party honest-majority protocol that could securely perform conjunction analyses was investigated in [KW14a].

Although MPC provides a general framework for computing arbitrary functions securely, efficiency is a primary concern. The four main approaches to MPC use Fully Homomorphic Encryption (FHE) [Gen09], the GMW Protocol [GMW87], the information-theoretic BGW protocol [BOGW88] or Garbled Circuits (GC) [Yao82, Yao86]. All four methods begin by converting the function of interest into a circuit (either boolean or arithmetic) and then evaluating the circuit gate-by-gate.[1] This approach yields polynomial-time algorithms, but without heavy optimizations, these protocols are not practical for any but the simplest calculations. Indeed, one of the primary technical challenges in this area is to design protocols that are efficient enough to be used in practice. Several works in the past [FSW03, CS10, ABZS13, KW14b, PS15] have looked at optimizing MPC for real-valued computations, both for fixed and floating point, though these operate in the honest majority setting, which is a major barrier in the two-party setting for when two satellite operators only want to talk to each other.

1.1 Our Results

In this paper we describe a new design and implementation of a secure two-party computation framework for high-precision arithmetic of real-valued functions that go beyond standard floating point levels of accuracy. As an application, we show that it allows two satellite operators to perform a conjunction analysis securely without the need to share their private orbital information with each other or any outside party. From a theoretical standpoint, our solution is provably secure in the OT-hybrid model assuming only the existence of one-way functions.

Our main contribution is an efficient construction of a new scheme that can securely evaluate complex numerical calculation, in particular, the numerical integration of calculating the probability of collision in a conjunction analysis computation via "dynamic" fixed-point integer calculations. Our new scheme extends the integer-arithmetic GMW protocol [GMW87] and Garbled Circuits [Yao82, Yao86], and we show how to do it securely in the two-party setting where there is no honest majority. This scheme can evaluate not only arithmetic gates $(+, \times)$, but also augmented functionality gates which include comparison (which will be optimized as a less-than-zero gate), shift-by-constant, which then allows us to perform higher order functions

---

[1]An alternative approach using Garbled RAM [LO13, GHL$^+$14] avoids the problem of converting the calculation to a circuit. Practical implementations of GRAM is an interesting area to be explored.

such as integer division, square root, exp, erf, and numerical integration. As part of our new construction, we also introduce several optimizations and secure computation tools along the way, which may be of independent interest. Furthermore, our framework is sufficiently general that it could be of interest to other domains that require the secure computation of numerical analysis. Finally, we provide an extension to the scheme by deploying the arithmetic MAC technique of BeDOZa [BDOZ11] and SPDZ [DPSZ12] to provide security guarantees against a larger class of adversaries.

Our secure numerical integration routines are extremely stable and efficient. Using commercial hardware (a Dell laptop), we are able to compute a conjunction analysis in just a few minutes. We envision a deployment, where operators use public orbital information to identify satellites that are at risk of collision. Then, for each of these "close" encounters the two owners would engage in a secure two-party computation to calculate the true collision probability. After this public prefiltering, each operator will only have have a small number of high-precision conjunction analysis calculations that need to be performed securely. Conjunction analyses can be performed for near approaches up to five days in advance [HAO10], so a computation time on the order of minutes should be efficient enough for practical applications.

2 Background

2.1 Secure Computation

We use the standard real/ideal paradigm for defining security of a Multiparty Computation (MPC) protocol. We let $\stackrel{\text{comp}}{\approx}$ denote computational indistinguishability of probability distributions, *i.e.*, no PPT algorithm can distinguish them with non-negligible probability.

Definition 2.1. *We say that a (two party) protocol π securely computes (in the semi-honest model) a deterministic functionality $\mathcal{F}$ if for every PPT algorithm A there is a PPT algorithm Sim such that*

$$\text{IDEAL}^{\mathcal{F}}_{\textsf{Sim}} \stackrel{\text{comp}}{\approx} \text{REAL}^{\pi}_{A}$$

Where $\text{IDEAL}^{\mathcal{F}}_{\textsf{Sim}}$ is the probability distribution of the output of the simulator Sim interacting with the ideal functionality (i.e., Sim gets the input and output of the corrupted party) and REAL^{π}_{A} is the probability distribution of the view (protocol messages and internal randomness) of the corrupted party A during the execution of the real protocol.

We briefly outline some of the relevant technology below.

2.1.1 Garbled Circuits

Garbled Circuits (GC) were originally proposed by Andrew Yao in oral presentations [Yao82, Yao86]. The first formal proof of security for Yao's GC protocol was given by Lindell and Pinkas [LP09], and later formalized as a standalone cryptographic primitive by Bellare, Hoang and Rogaway [BHR12b, BHR12a].

Although there have been many improvements and variants on Yao's original idea for garbled circuits, the important features of GC-based MPC protocols are: (1) the protocol can be done in two rounds of communication, independent of the size of the circuit. (2) Each garbled gate is encrypted using a symmetric-key cryptosystem (e.g. AES) so evaluating the garbled circuit requires roughly a number of AES operations proportional to the size of the circuit. (3) Transferring the input tokens requires OT, a public-key operation, and the size of this public-key computation is roughly proportional to the size of the secret inputs.

More formally, a circuit garbling scheme as a triple of algorithms (G, GI, GE) where $G(1^k, C)$ takes as input a security parameter k and circuit C and outputs some garbled circuit Γ and garbling key gsk. $X \leftarrow GI(x, gsk)$ converts an input x and a gsk into a garbled input X, and $y \leftarrow GE(\Gamma, X)$ evaluates a garbled circuit Γ on an garbled input X. We only consider schemes that are projective, namely $X \leftarrow GI(x, gsk)$ can be decomposed into garbling the bits of x, namely $X \leftarrow GI(0, x_0, gsk), \dots, GI(n, x_n, gsk)$. Correctness means that y should be equal to $C(x)$. The security property is that, given Γ and X, nothing is revealed about x beyond what is revealed by the output. More formally,

Correctness For correctness, we require that for any circuit C and input x we have that that:

$$\Pr\left[C(x) = GE(\Gamma, X)\right] = 1$$

where $(\Gamma, gsk) \leftarrow G(1^k, C)$ and $X \leftarrow GI(x, gsk)$.

Security For security, we require that there is a PPT simulator CircSim such that for any C, x, we have that:

$$(\Gamma, X) \stackrel{\text{comp}}{\approx} \text{CircSim}\left(1^k, C, C(x)\right)$$

where $(\Gamma, gsk) \leftarrow G(1^k, C)$ and $X \leftarrow GI(x, gsk)$.

In recent years, there have been many extremely efficient implementations of garbled circuits including FastGC [HEKM11], VMCrypt [Mal11], the billion-gate compiler [KSS12], JustGarble [BHKR13], TASTY [HKS+10], ABY [DSZ15], OblivM [LWN+15] and Frigate [MGC+16].

2.1.2 Secret-sharing based Protocols

Secret-sharing based protocols were first introduced by works of GMW [GMW87], or BGW [BOGW88] and CCD [CCD88]. Although these three protocols are all based on secret-shared computations, each uses different mechanisms and has different security models (honest vs. no honest majority, the existence of Oblivious Transfer, etc.). In these protocols, each player begins by secret-sharing [Sha84] her private inputs among all the players. The players then engage in a protocol to compute the given circuit, gate-by-gate, on the shares. The GMW protocol requires OT for each multiplication gate, while the BGW protocol is information-theoretic and computing each gate requires only linear algebra (along with the assumption that a strict majority of players are honest). The important features of secret-sharing based protocols are: (1) the round complexity of the protocol is proportional to the depth of the circuit. (2) The GMW protocol requires a number of OTs roughly proportional to the number of multiplication gates in the circuit.

There have been many practical implementations of GMW-based MPC protocols including VIFF [DGKN09], TinyOT [NNOB12] and Wysteria [RHH14]. The Sharemind platform [BLW08] provides a general-purpose platform for secure computation based on the BGW protocol. The Sharemind platform has been used to perform conjunction analyses [KW14a]. A more detailed comparison of our work with the Sharemind implementation can be found in Section 6.2.

Archer et al. provide a survey on the state of practical MPC protocols [ABPP15].

2.1.3 Online/Offline Model of Secure Computation

Our construction works in the online/offline model of computation, and we briefly review the model here. The idea behind the online/offline model is that certain amounts of cryptographic material can be computed *independently* of the input, such that it can be stored and then recalled during the live computation when the inputs are available. To improve efficiency, we split our secure computation into two phases: the offline phase where input-independent data is precomputed and the online phase where input-dependent computations occur. Early research in this area was done by Beaver [Bea95, Bea97], and the power of this model continues to be demonstrated in works such as Ishai et al. [IKM+13] and have found use in implementations as well [BDOZ11, DPSZ12]

Th offline phase may involve communication between the parties – as long as that communication is independent of their private inputs. In some situations, we can use the aid of an *offline dealer* that only participates during the offline phase and contributes no inputs nor receives any outputs during the online phase, but instead distributes correlated randomness to the parties. Secondly, we can talk about two kinds of pre-computed data: those that remain persistent across multiple online invocations (e.g. public keys and parameters being sent in advance), and those that are used and consumed (e.g. one-time pads being sent in advance). Then, during the online phase, cryptographic material from the offline phase is used in conjunction with the inputs in order to do the live computation. We will see many instantiations of this in the forthcoming sections.

Figure 2: The basic functionalities of Oblivious Transfer (OT), and oblivious multiplication

2.2 Oblivious Transfer

Oblivious transfer (OT) is a cryptographic primitive introduced by Rabin [Rab05], and we use the 1-out-of-2 variant introduced by Even-Goldreich-Lempel [EGL82]. In this variant, a Sender holds two string values x_0 and x_1 and the Receiver holds a choice bit b. The Receiver should get x_b without learning anything about x_{1-b} and without the Sender learning anything about b. This can be viewed as a secure protocol for the functionality $\mathsf{OT}((x_0, x_1); b) = (\bot; x_b)$, where $\bot$ denotes the empty message. We write OT_m to indicate that the strings x_0 and x_1 are of length m. Oblivious transfer can be implemented under a variety of standard assumptions (cf. Goldreich's book [Gol01, Gol04] for details).

2.3 Shared Arithmetic Triples (Oblivious Linear-function Evaluation)

Oblivious linear function evaluation is a natural extension of OT, where one party (e.g. the Sender) holding two values $a, b \in \mathbb{F}$ and another party (e.g. the Receiver) holding some value $x \in \mathbb{F}$. The goal is to have the Receiver get $ax + b$ without learning a, b and without the Sender learning x.

We consider two alternative, equivalent functionalities: one of shared-input-shared-output multiplication (SISO-M), and one of just shared-output multiplication (SO-M). Due to the symmetric nature of these notions, we refer to the two parties as Alice and Bob instead of the Sender and Receiver.

For SISO-M, Alice and Bob hold shares of two field elements, x, y and they would like to obtain shares of the product. In particular, Alice holds $x_0, y_0 \in \mathbb{F}$ and Bob holds $x_1, y_1 \in \mathbb{F}$ and the goal is to have Alice obtain a random $z_0 \in \mathbb{F}$ and Bob obtain z_1 such that $(z_0 + z_1) = (x_0 + x_1) \cdot (y_0 + y_1)$. Note that since each of the two outputs is uniformly distributed over the field, each party learns no additional information about the other party's input. In the shared-output version, Alice holds x, Bob holds y and they want to obtain z_0 and z_1 such that $(z_0 + z_1) = x \cdot y$. See Figure 2.

2.4 Conjunction Analysis Calculations

Before we describe our secure conjunction analysis computation, we review the problem of computing it in the clear. Our secure computation solution is based on Alfano's method [Alf05].

Each satellite is modelled as a spherical object, and thus its dimensions are completely captured by a single parameter, its radius. Although the radius is not particularly sensitive, our solution will hide the

Figure 3: The hardbody of radius $R_a + R_b$ traces out a collision tube through the combined density ellipsoid. The probability of collision is the probability mass of the density function inside the collision tube.


Cross section of pdf

Cross section of collision tube

Figure 4: The Encounter Plane

radius as well the trajectory information. Each satellite operator's private input has four parts

> **Position:** $\mathbf{p}_a, \mathbf{p}_b$ in $\mathbb{R}^3$
> **Velocity:** $\mathbf{v}_a, \mathbf{v}_b$ in $\mathbb{R}^3$
> **Error:** Covariance matrices $\mathbf{C}_a, \mathbf{C}_b$ in $\mathbb{R}^{3 \times 3}$
> **Radius:** R_a, R_b in $\mathbb{R}$

Private inputs from each satellite operator.

Each satellite is assumed to deviate from its position, $\mathbf{p}$, and these deviations are assumed to be normally distributed with covariance matrix[2] $\mathbf{C}$. Thus each satellite's physical location is given by $\mathbf{p} + \mathcal{N}(\mathbf{C})$, These normal distributions are truncated at eight standard distributions [Alf07] resulting in a "density ellipsoid." Although each satellite is on an elliptical path, in any short time window, the satellite's trajectory is almost linear, and during the course of the conjunction analysis, the two satellites are assumed to have linear relative velocities. These simplifying assumptions were not introduced to facilitate a secure computation, but instead are all part of the routine (insecure) conjunction analyses being performed on a daily basis.

Because the positional errors on the two satellites are assumed to be independent, we can shift all the errors onto one body, and imagine a "hardbody" of radius $R_a + R_b$ passing through a density ellipsoid with covariance matrix $\mathbf{C}_a + \mathbf{C}_b$. This spherical hardbody traces a "collision tube" through the combined density ellipsoid, and the probability of collision is then simply the probability mass of the density ellipsoid within this collision tube (See Figure 3).

To simplify the calculation further, the three dimensional pdf is sliced, perpendicular to the relative velocity, at the point of nearest approach (*i.e.,* at the point where the density is largest). This defines the "encounter plane", and the cross-section of the ellipsoid in the encounter plane is a density ellipse (See Figure 4).

Thus the final probability is calculated by integrating the two dimensional density ellipse in the region given by the cross-section of the hardbody. Given the combined radius $R = R_a + R_b$, the center of the

---

[2]These covariance matrices are usually assumed to be diagonal, *i.e.,* the variances along the three principal axes are independent. This assumption does not significantly affect the computation.

hardbody in the encounter plane, (x_m, y_m), and the lengths of the semi-principle axes (σ_x, σ_y) of the density ellipse, we calculate the probability of collision, p.

$$p = \frac{1}{2\pi\sigma_x\sigma_y} \int_{-R}^{R} \int_{-\sqrt{R^2-x^2}}^{\sqrt{R^2-x^2}} f(x,y)dydx \tag{1}$$

where the integrand is

$$f(x,y) = \exp\left[\frac{-1}{2}\left[\left(\frac{x-x_m}{\sigma_x}\right)^2 + \left(\frac{y-y_m}{\sigma_y}\right)^2\right]\right]$$

Changing variables, this becomes

$$p = \frac{3}{\sqrt{8\pi}\sigma_x} \int_{-R}^{R} g(x)dx \tag{2}$$

where the integrand is

$$g(x) = \left[\text{erf}\left(\frac{y_m + \sqrt{R^2-x^2}}{\sqrt{2}\sigma_y}\right) + \text{erf}\left(\frac{-y_m + \sqrt{R^2-x^2}}{\sqrt{2}\sigma_y}\right)\right] \exp\left(\frac{-(x+x_m)^2}{2\sigma_x^2}\right)$$

This integral does not have a closed form, and Alfano suggests approximating this integral using Simpson's Rule (*i.e.*, approximating the integral using arcs of parabola). The complete calculation is described in Algorithm 1. A more thorough discussion of the mathematics involved is in Appendix A. We also apply a change-of-variables $z = x/R$ so that the square root inside the integral and the limits of integration does not depend on inputs, which allows for greater flexibility in hardwiring constants into the circuit.

Algorithm 1 The conjunction analysis calculation

1: **Inputs:** $\{\mathbf{v}_i, \mathbf{C}_i, \mathbf{p}_i, R_i\}_{i\in a,b}$
2: $\mathbf{v}_r \leftarrow \mathbf{v}_b - \mathbf{v}_a$
3: $\mathbf{i} \leftarrow \frac{\mathbf{v}_r}{|\mathbf{v}_r|}, \quad \mathbf{j} \leftarrow \frac{\mathbf{v}_b \times \mathbf{v}_a}{|\mathbf{v}_b \times \mathbf{v}_a|}, \quad \mathbf{k} \leftarrow \mathbf{i} \times \mathbf{j}$
4: $\mathbf{Q} \leftarrow \begin{bmatrix} \mathbf{j} & \mathbf{k} \end{bmatrix}$
5: $\mathbf{C} \leftarrow \mathbf{Q}^T(\mathbf{C}_a + \mathbf{C}_b)\mathbf{Q}$
6: $(\mathbf{u}, \mathbf{v}) \leftarrow \text{Eigenvectors}(\mathbf{C})$
7: $(\sigma_x^2, \sigma_y^2) \leftarrow \text{Eigenvalues}(\mathbf{C})$
8: $\sigma_x \leftarrow \sqrt{\sigma_x^2}, \quad \sigma_y \leftarrow \sqrt{\sigma_y^2}$
9: $\mathbf{u} \leftarrow \frac{\mathbf{u}}{|\mathbf{u}|}, \quad \mathbf{v} \leftarrow \frac{\mathbf{v}}{|\mathbf{v}|}$
10: $\mathbf{U} \leftarrow \begin{bmatrix} \mathbf{u} & \mathbf{v} \end{bmatrix}$
11: $\begin{bmatrix} x_m \\ y_m \end{bmatrix} \leftarrow \mathbf{U}^T\mathbf{Q}^T(\mathbf{p}_b - \mathbf{p}_a)$
12:

$$p \leftarrow \frac{1}{2\pi\sigma_x\sigma_y} \int_{-R}^{R} \int_{-\sqrt{R^2-x^2}}^{\sqrt{R^2-x^2}} f(x,y)dydx$$

13: Where

$$f(x,y) = \exp\left[\frac{-1}{2}\left[\left(\frac{x-x_m}{\sigma_x}\right)^2 + \left(\frac{y-y_m}{\sigma_y}\right)^2\right]\right]$$

14: **Return:** p

```
G1: −1          //G1 is the constant value −1
G2: * G1 I1     //G2 is −1 * I1
G3: + I2 G2     //G3 is I2 − I1
G4: < 0 G3      //G4 is 1 if I2 − I1 < 0, else 0
G5: * G4 I1     //G5 is I1 if I2 − I1 < 0, else 0
G6: * G4 G1     //G6 is −1 if I2 − I1 < 0, else 0
G7: 1           //G7 is 1
G8: + G7 G6     //G8 is 1 if I2 − I1 ≥ 0, else 0
G9: * G8 I2     //G9 is I2 if I2 − I1 ≥ 0, else 0
O1: + G9 G5     //output gate 1 is max(I1, I2).
```

Figure 5: A circuit that takes two inputs (I1 and I2) and returns the maximum of the two.

3   Our Techniques

Our construction uses an "augmented" arithmetic circuit, consisting of (integer) addition, multiplication and division gates, with special gates for comparisons and bit-shifts, and we securely compute this circuit gate-by-gate.

3.1   Basic Gates

The basic gates used in our construction are as follows:

- **input**: Input gates represent inputs in $\mathbb{Z}$, and are represented with the prefix I, thus I3 is input gate 3.

- **output**: Output gates represent outputs in $\mathbb{Z}$, and are represented with the prefix O, thus O2 is output gate 2.

- **const**: Constant gates output a fixed constant. The constants are in $\mathbb{Z}$ (though we will later see that they are within some bound), which means in particular that they may include negative integers. Thus G20: −7 sets the value of gate 20 to −7.

- **+**: Addition gates take two integers (in $\mathbb{Z}$) as inputs and output their sum. Note that subtraction gates can be implemented using addition and multiplication gates since we are working over the integers, thus for the remainder of the paper we eliminate subtraction gates from our discussion.

- *****: Multiplication gates take two integers (in $\mathbb{Z}$) as inputs and output their product. For example, G100: * G50 G51 represents that the output of gate G100 is the product of outputs of two gates G50 and G51.

- **< 0**: This is a less-than-zero gate. It takes an integer (in $\mathbb{Z}$) as input and outputs 1 if the integer is strictly less than 0. In particular, G100: < 0 G50 represents that the output of gate G50 is set to be 1 if and only if G50 < 0.

- **>> c**: This gate represents a shift-by-constant gate, though as we shall see, it will actually represent truncation (via an implicit exponent reduction). The gate takes shifts an input by a public constant that is the output of a **const** gate. For example, G100: >> c G50 G51 represents that the output of gate G100 is G50 shifted to the right by G51.

3.2 Integer Representation and Implicit Denominator

Conceptually, the conjunction analysis computation works over the real numbers, and thus a few steps need to be taken in order to represent them as finite-sized integers. One approach is to represent real numbers using the standard IEEE floating point specification and perform arithmetic operations in that fashion. This is the approach that was taken in [KW14a].

To leverage the power of arithmetic circuits, a natural choice would be approximate real numbers as rationals. Thus we could approximate every $r \approx a/b$ and carry around two integers to represent each value. Using this representation addition is somewhat complex, although division is essentially "free."

After unsatisfactory performance using the rational representation, the representation that we use is akin to that of fixed-width arithmetic. We use an implicit rational representation where the denominator is some power of 2. This power is computed based on a static analysis of the input ranges and the circuit itself, and is attached to each wire as public "metadata". Using this representation, we can obtain as many decimal places of accuracy as desired at the cost of the numerator growing larger. Because the denominator is public, during a secure computation, each party can *locally* adjust their shares in order to arrive at a common denominator. For example, if Alice holds $x_0/2^n$ and $y_0/2^m$ and Bob holds $x_1/2^n$ and $y_1/2^m$, then they can locally convert the x shares to have denominator $\max(2^n, 2^m)$ by multiplying by $2^{|n-m|}$ to the shares with the smaller denominator.

The numerators are represented as m-bit two's-complement signed integers, and m can be chosen large enough that during the course of evaluation it will never overflow. To further ensure non-overflow, we use these shift gates as accuracy truncation rather than an actual shift: namely anytime we shift the numerator, we also implicitly shift the denominator, hence not changing the overall number. Therefore, shift gates can be thought of as truncation gates.

3.3 Combining GC with Arithmetic GMW

We compute (integer) addition and multiplications natively using GMW. To compute comparison and shift gates, we represent them as *Boolean* circuits and then evaluate them using GC. The GC must take two secret-inputs and compute a secret sharing of the output of the gate. Since the inputs are arithmetically secret shared, one must first convert them to bits before inputting them into the Boolean circuit that computes an augmented gate, and then convert them back into arithmetic shares. We explain how two parties can perform these computations in Algorithms 2 and 3. Share conversion of this nature has been investigated in previous works, e.g. Yu and Yang [YY12].

Algorithm 2 Share-converted Less-Than-Zero

1: **Hardwired:** A modulus $M = 2^m$
2: **Inputs:** Alice holds x_0, Bob holds x_1. Alice additionally provides a random R
3: $x \leftarrow x_0 + x_1 \pmod{M}$ using standard m-bit add-with-carry circuit
4: $b \leftarrow sgn(x)$
5: **Return:** $z_1 = b + R \pmod{M}$ to Bob. Alice sets $z_0 = -R \pmod{M}$ herself.

Algorithm 3 Share-converted shift-right-by-constant

1: **Hardwired:** A modulus $M = 2^m$ and a shift constant c.
2: **Inputs:** Alice holds x_0, Bob holds x_1. Alice additionally provides a random R
3: $x \leftarrow x_0 + x_1 \pmod{M}$ using standard m-bit add-with-carry circuit
4: $y \leftarrow x >> c$ by duplicating the sign bit wire and dropping c rightmost wires
5: **Return:** $z_1 = y + R \pmod{M}$ to Bob. Alice sets $z_0 = -R \pmod{M}$ herself.

Representing Mathematical Functions as Circuits: Evaluating the integral given in Equation 2 requires division, $\exp(\cdot)$, $\sqrt{\cdot}$, and $\text{erf}(\cdot)$. We explain how we chose to implement these functions using our circuits.

Circuit Representation for Division: We implement integer division using repeated subtraction. Using known bounds on the inputs, we can track maximum and minimum values for each gate in the circuit, and using this (public) meta information, we can bound the number of subtractions necessary for each division in the circuit.

Circuit Representation for $\exp(\cdot)$: We represent the function $\exp(\cdot)$ using a degree-24 Taylor series. We hard-code the Taylor coefficients as constants in the circuit.

Circuit Representation for $\sqrt{\cdot}$: To approximate a square root, we use the iterative Babylonian Algorithm. Given an input S, and an initial estimate x_0, the Babylonian Algorithm computes

$$x_{n+1} = \frac{1}{2}\left(x_n + \frac{S}{x_n}\right)$$

For increased efficiency, we compute multiple steps at once, *i.e.*, computing x_{n+4} as a ratio of two degree 16 polynomials in S and x_0. For our calculations, we do 6 batches of degree 16 each to ensure a sufficient degree of accuracy.

Circuit Representation for $\mathrm{erf}(\cdot)$: Recall

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

Unfortunately, erf has no closed-form solution, so we need a method for approximating erf using arithmetic circuits. Taylor expansions of erf fare poorly outside of a very restricted domain (see Appendix D). Instead of the Taylor expansion, we approximate $1 - \mathrm{erf}(x)$ using the degree 96 rational function:

$$\frac{1}{(1 + a_1 x + a_2 x^2 + a_3 x^3 + a_4 x^4 + a_5 x^5 + a_6 x^6)^{16}}$$

Where

a_1	$= .3275911$	a_4	$= .0001520143$
a_2	$= .254829592$	a_5	$= .0002765672$
a_3	$= .0092705272$	a_6	$= .0000430638$

When $x > 0$ this approximation has strong error bounds that are uniform for all values x. In particular, this approximation has an absolute error less than 10^{-7} across the entire range [AS65].

4 Main Construction

We now describe our main construction, which is the semi-honest two-party secure computation of the conjunction analysis functionality in the online/offline model. We let π_{CA} denote the protocol we are about to describe. We first describe how to precompute the cryptographic resources during the offline phase.

4.1 Precomputation of Cryptographic Resources

In the offline preprocessing phase, we generate three cryptographic "resources" for later use.

4.1.1 Pregenerated Random OTs

Our computation requires executing a huge number of OTs. Using a now standard trick due to Beaver [Bea95], we can generate random OTs during the precomputation phase and then later "consume" these random OTs during the online computation.

To pregenerate a random OT, the Sender holds *random* (r_0, r_1) and the Receiver holds a *random* bit z. After running OT on these random values, so the Receiver gets r_z. We now describe how to use this random OT to perform an actual OT in the online phase.

Now suppose they want to run OT on actual inputs (s_0, s_1) for the sender and b for the Receiver. To begin, the Receiver sends $w = z \oplus b$, then the Sender sends to the Receiver $(q_0, q_1) = (s_0 \oplus r_{0 \oplus w}, s_1 \oplus r_{1 \oplus w})$. Finally the Receiver outputs $t = q_b \oplus r_z$.

There are two methods by which the participants can generate the necessary random OT instances during the online phase. (1) The participants can generate a small number of OT instances and then use OT extension [IKNP03] to extend these to a huge number of OTs. (2) They can make use of a trusted dealer who simply provides correlated randomness to the two players, *i.e.,* the dealer will generate three random values (r_0, r_1, z) and provide r_0, r_1 to Alice and z, r_z to Bob.

4.1.2 Pregenerated Random Triples

Similar to OTs, for shared arithmetic triples, we can generate random arithmetic triples in the precomputation phase. We outline the technique on how to actually generate shared-output versions of these, which is due to the work of Ishai et al. [IPS09]. This technique works in batches of size t and makes use of OTs.

Suppose Alice holds $a_1, \ldots, a_t$ and Bob holds $b_1, \ldots, b_t$ and they want to compute shares of all $a_i \cdot b_i$. In other words, they want to compute a SO-M (see Figure 2). Let $k = 2t$ and $n = ck$ where c is a constant. Suppose we have t distinct evaluation points ζ_i and n distinct evaluation points θ_i, distinct from the ζs.

Bob: Let $B(x)$ be a random degree $k - 1$ polynomial such that $B(\zeta_i) = b_i$ for $i = 1 \ldots t$. Such a B can be found by interpolation. Sample the polynomial at the points θ_i to get $y_i = B(\theta_i)$ for $i = 1 \ldots n$. Sample $L \subset \{1, \ldots, n\}$ at random of size $t + k - 1$. Set $v_i = y_i$ if $i \in L$, and v_i to be random if $i \notin L$. Send the v_is to Alice.

Alice: Let $A(x)$ be the unique degree $t - 1$ polynomial such that $A(\zeta_i) = a_i$ for $i = 1 \ldots t$. Such an A can be found by interpolation. Let Alice choose $R(x)$, a random degree $t + k - 2$ polynomial. Compute $x_i = A(\theta_i)$ and $r_i = R(\theta_i)$ for $i = 1 \ldots n$. Set $w_i = x_i \cdot v_i - r_i$.

Alice and Bob: For each $i = 1 \ldots n$, Bob plays the role of the Receiver in an OT protocol with Alice, who plays the role of the Sender. Bob sets his bit b to be 1 if and only if $i \in L$, and Alice sets her messages $x_1 = w_i$ and x_0 to be random.

Bob: Bob sets Q to be the unique polynomial of degree $t + k - 2$ with $Q(\theta_i) = w_i$ for $i \in L$, again via interpolation. Bob sets $bob_i = Q(\zeta_i)$, Alice sets $alice_i = R(\zeta_i)$, for $i = 1 \ldots t$. These are the output shares $z_{0,i} = alice_i$ and $z_{1,i} = bob_i$ for Alice and Bob, respectively.

To see why this works, observe that $Q + R = A \cdot B$, and so $bob_i + alice_i = Q(\zeta_i) + R(\zeta_i) = A(\zeta_i) \cdot B(\zeta_i) = a_i \cdot b_i$.

To convert these into shared-input-shared-output triples, we proceed as follows: Suppose Alice holds x_0, y_0 and Bob holds x_1, y_1. They want to compute shares z_0 and z_1 of $(x_0 + x_1)(y_0 + y_1)$. This is just equal to $x_0 y_0 + x_0 y_1 + x_1 y_0 + x_1 y_1$. They can make two calls to the above protocol, once with x_0 and y_1 as inputs, and once with x_1 and y_0 as inputs, which allows us to get shares for the cross-terms $x_0 y_1$ and $x_1 y_0$.

4.1.3 Pregenerated Garbled Circuits

Our protocol will use garbled circuits to compute the shift and comparison gates on secret-shared values. Because the circuit garbling procedure is input-independent, we perform a small precomputation here where we generate all the Γ during this phase.

Our garbled circuits implementation uses a fixed-key blockcipher (AES) as described in JustGarble [BHKR13]. We do not include optimizations like free XOR [KS08, KMR14, App13] or half-gates [ZRE15].

4.2 Online Phase

For the online phase, we use the arithmetic version of the Goldreich-Micali-Wigderson [GMW87] paradigm: each party secret shares his or her inputs, then performs gate-by-gate operations as described below. Alice and Bob's inputs are denoted X_0, Y_0 and X_1, Y_1 respectively.

- For every addition and subtraction gate, Alice and Bob respectively add or subtract their local shares $Z_0 = X_0 \pm Y_0$ and $Z_1 = X_1 \pm Y_1$.

- For every multiplication gate, Alice and Bob consume a shared triple to obtain Z_0 and Z_1 as follows.

  In order to "consume" a precomputed multiplication triple, we do the following steps. Suppose Alice holds X_0, Y_0 and Bob holds X_1, Y_1, and they wish to compute shares Z_0 for Alice and Z_1 for Bob such that $(Z_0 + Z_1) = (X_0 + X_1)(Y_0 + Y_1)$.

Suppose we have a pregenerated triple, *i.e.,* two random values A and B and we let $A \cdot B = C$. Now suppose Alice holds random shares A_0, B_0, C_0 and Bob holds random shares A_1, B_1, C_1 where $A_0 + A_1 = A$, $B_0 + B_1 = B$, $C_0 + C_1 = C$. Then by "consuming" this random triple, Alice and Bob can compute shares of Z as follows:

Alice locally computes $W_0 = X_0 - A_0$, $V_0 = Y_0 - B_0$, Bob locally computes $W_1 = X_1 - A_1$, $V_1 = Y_1 - B_1$.

Alice sends Bob W_0, V_0 and Bob sends Alice W_1, V_1.

Alice locally computes $Z_0 = C_0 + (W_0 + W_1)Y_0 + (V_0 + V_1)X_0 - (W_0 + W_1)(V_0 + V_1)$.

Bob locally computes $Z_1 = C_1 + (W_0 + W_1)Y_1 + (V_0 + V_1)X_1 - (W_0 + W_1)(V_0 + V_1)$.

Observe that $Z_0 + Z_1 = (C_0 + C_1) + (W_0 + W_1)(Y_0 + Y_1) + (V_0 + V_1)(X_0 + X_1) - (W_0 + W_1)(V_0 + V_1)$, which in turn is equal to $C + (X - A)Y + (Y - B)X - (X - A)(Y - B) = C + Z - AY + Z - BX - Z + BX + AY - C = Z$.

- For every shift gate, Alice and Bob's shares are X_0 and X_1 and they want to shift $(X_0 + X_1)$ by some publicly known amount N. This is accomplished using a (precomputed) garbled circuit computing Algorithm 3, where Alice sends Bob the garbled circuit and then Bob uses OT to obtain labels corresponding to his inputs.

- For every comparison gate (optimized as a less-than-zero gate), Alice has her share X_0 and Bob has his share X_1 and they want to see if $(X_0 + X_1)$ is positive or not. This is done via a circuit that computes Algorithm 2. We then use the garbled version of this circuit (that was precomputed earlier) to evaluate it where Alice sends Bob the garbled circuit and then Bob uses OT to obtain labels corresponding to his inputs.

Note that this process can be parallelized across an entire layer of the circuit, so that interaction occurs at each level of the arithmetic circuit rather than at each gate (excluding the free gates). In the end, the output values are shared as O_0 and O_1, whereupon Alice and Bob reveal to each other their shares to obtain the final output.

4.3 Security Proof

We now state our main theorem.

Theorem 4.1. *Assume the existence of an OT functionality (we work in the OT-hybrid model) and a secure garbling scheme (G, GI, GE) (which can be built from any one-way function). Then protocol π_{CA} securely computes the conjunction analysis functionality.*

Proof Sketch. We show how a PPT simulator Sim can efficiently simulate the view of a party given only their input and the output of the CA in the OT-hybrid model. We let CircSim denote the garbled circuit simulator. We refer the reader to [IPS09] for a proof of security of the shared triples generation in the OT-hybrid model. Suppose we have a topological ordering of the M wires $W_0, \ldots, W_M$ (we can associate the output wire(s) of a gate with the gate itself). We go through a standard hybrid argument and define a series of hybrids of views $\mathbf{Hyb}_i$ in which the real gates are replaced with simulated gates, gate by gate. Next, we show how to generate simulated shares for the output of a gate given the simulated input shares (of the corrupted party) of a gate.

For input gates, if the input belongs to the corrupted party, then Sim has the actual value and can secret share it. Otherwise, if the input belongs to the other party, the simulator generates a random share and sets the simulated received share as that random share. For output gates, since Sim knows the output, given the input share, it will claim the received share is the correct output minus its own share. For addition and subtraction gates, since no interaction is needed, input shares are simply added or subtracted to obtain the output shares. For multiplication gates, by the security of the offline phase of triple generation due to [IPS09] and the security of basic GMW protocol, the share of the output wire can be set to be uniformly random. Finally, the less-than-zero and shift gates are calculated via a garbled circuit, and is the only source of asymmetry between the two parties. When Alice is being simulated, since she is the garbled circuit generator and chooses R, the simulator also generates a garbled circuit honestly and chooses a random R and

sets that as Alice's simulated output share. On the other hand, when Bob is being simulated, the simulator first chooses a uniformly random output share R', and then invokes CircSim on R' to obtain a simulated garbled circuit and input labels (X, Γ). Then, in order to generate the view for Bob, it makes the OT-hybrid oracle return X and simulates a message Γ from Alice.

This concludes the description of how the simulator Sim simulates wires. The entire circuit is then simulated gate by gate.

$\square$

5 Extending the Construction

In this section, we show how to extend our construction to provide security against semi-malicious adversaries. In the malicious security model, corrupted parties are allowed to deviate arbitrarily from the prescribed protocol, and the protocol is secure if nothing revealed beyond what is revealed by the output alone.

Our extended construction considers a slightly weaker model, where in the presence of malicious behavior, the honest party will detect such activity and will immediately abort the protocol. The event of abort is information, so the honest player's decision to abort the protocol, may leak information – but we allow this leakage.

The construction in this section achieves security against malicious adversaries *with correlated abort* (see e.g. [IKO+11]), *i.e.,* nothing is revealed except the output, unless in the case of an abort where nothing is learned except that an abort occurred.

In order to extend our solution from the semi-honest model to this stronger setting, we note that our shared triples generation is only secure in the semi-honest model, and thus in order to be able to now securely precompute triples we must resort to another technique. There are several works (e.g. [BDOZ11, DPSZ12] and their follow-up research) that focus on optimizing this offline construction, and our novel contribution is focused on the online phase. Therefore, we employ a secure two-party computation solution with an offline phase that assumes the assistance of a semi-trusted dealer.

We describe how to securely evaluate an arithmetic circuit $C(x_1, \ldots, x_n, y_1, \ldots, y_n)$ containing $+, -, \times, <$ $0, >> c$ gates over the integers, where < 0 is the "less-than-zero" gate, and $>> c$ is the "arithmetic-shift-right-by-c" gate (we also include "constant" gates). In terms of representation, we bound the total number of bits of any intermediate value in the computation and choose a modulus N that is twice as large as the bound plus a security parameter, and use $[N/2, N/2)$ as the set of representatives for the integers modulo N. We henceforth take all arithmetic to be modulo N, where comparison to zero and shift still makes sense because we are taking the half-interval around zero representation.

Suppose Alice holds the inputs $x_1, \ldots, x_n$ and Bob holds the inputs $y_1, \ldots, y_n$. We will evaluate each gate individually, and thus evaluate the entire circuit by evaluating the "tree" of gates inductively (this is done in parallel to the furthest extent possible, we only wait if one gate depends on another gate). For addition, subtraction, multiplication, and constant gates, we employ an arithmetic MAC style strategy for evaluation (see, e.g. [BDOZ11]) which we describe here.

Let g be a gate, and let L and R be the left and right inputs, respectively, and let $O = g(L, R)$ be the output of the gate. Let $MAC_{\alpha, \beta}(x)$ be $\alpha x + \beta$. We inductively assume that the two inputs are shared and MACced in the following fashion:

- Alice and Bob each privately hold their own global MAC key, α_A and α_B, and they each hold a unique β for each wire in the circuit.

- Let $\beta_A^L, \beta_B^L, \beta_A^R, \beta_B^R$ denote the β for the left and right inputs, for Alice and Bob.

- Alice will then hold random shares x_L and x_R and Bob will hold random shares y_L and y_R subject to $x_L + y_L = L$ and $x_R + y_R = R$.

- Alice will hold $w_A^L = MAC_{\alpha_B, \beta_B^L}(x_L)$ and $w_A^R = MAC_{\alpha_B, \beta_B^R}(x_R)$.

- Bob will hold $w_B^L = MAC_{\alpha_A, \beta_A^L}(y_L)$ and $w_B^R = MAC_{\alpha_A, \beta_A^R}(y_R)$.

- **GOAL:** Obtain x_O, w_A^O, β_A^O for Alice and y_O, w_B^O, β_B^O for Bob such that the inductive invariants $x_O + y_O = O$, $w_A^O = MAC_{\alpha_B, \beta_B^O}(x_O)$, and $w_B^O = MAC_{\alpha_A, \beta_A^O}(y_O)$ hold.

For each gate type, we describe what is required as pregenerated content, and what is done online. As the first step to setup, the dealer generates α_A and α_B and sends them to Alice and Bob respectively.

CONSTANT GATES. For constant gates, if the constant is c, the dealer generates β_A, β_B at random, x at random, and sets $y = c - x$. It then computes $w_A = MAC_{\alpha_B, \beta_B}(x)$ and $w_B = MAC_{\alpha_A, \beta_A}(y)$. It sends x, w_A, β_A to Alice and y, w_B, β_B to Bob to store. During the online phase, Alice and Bob recall these values from storage when needed.

ADDITION AND SUBTRACTION GATES. For addition/subtraction gates, no pregeneration by the dealer is necessary. Indeed, Alice computes $\beta_A^O = \beta_A^L \pm \beta_A^R$, $w_A^O = w_A^L \pm w_A^R$, and $x_O = x_L + x_R$, and Bob performs the analogous computations. Then it is the case that the inductive invariant holds since: $x_O + y_O = x_L \pm x_R + y_L \pm y_R = (x_L + y_L) \pm (x_R + y_R) = L \pm R = O$, and $w_A^O = w_A^L \pm w_A^R = MAC_{\alpha_B, \beta_B^L}(x_L) \pm MAC_{\alpha_B, \beta_B^R}(x_R) = (\alpha_B \cdot x_L + \beta_B^L) \pm (\alpha_B \cdot x_R + \beta_B^R) = \alpha_B(x_L \pm x_R) + (\beta_B^L \pm \beta_B^R) = \alpha_B \cdot x_O + \beta_B^O = MAC_{\alpha_B, \beta_B^O}(x_O)$.

MULTIPLICATION GATES. For multiplication gates, the pregeneration consists of an authenticated "triple". The dealer generates two random numbers a, b and computes $c = a \cdot b$. The dealer then creates authenticated shares for each of these three values as follows. Randomly select $\beta_A^a, \beta_A^b, \beta_A^c, \beta_B^a, \beta_B^b, \beta_B^c$, and x_a, x_b, x_c at random, and set $y_a = a - x_a, y_b = b - x_b, y_c = c - x_c$. It then sets $w_A^a = MAC_{\alpha_B, \beta_B^a}(x_a), w_A^b = MAC_{\alpha_B, \beta_B^b}(x_b), w_A^c = MAC_{\alpha_B, \beta_B^c}(x_c)$ and $w_B^a = MAC_{\alpha_A, \beta_A^a}(y_a), w_B^b = MAC_{\alpha_A, \beta_A^b}(y_b), w_B^c = MAC_{\alpha_A, \beta_A^c}(y_c)$. It sends $x_a, x_b, x_c, w_A^a, w_A^b, w_A^c, \beta_A^a, \beta_A^b, \beta_A^c$ to Alice, and $y_a, y_b, y_c, w_B^a, w_B^b, w_B^c, \beta_B^a, \beta_B^b, \beta_B^c$ to Bob.

During the online step, each party subtracts their a share from their left share (*i.e.*, subtracting w, x, β simultaneously), and the b share from their right share. The resulting differences, call them r and s are mutually revealed by sending each other the resulting MAC and value, and mutually verified. Then each computes their local result as the sum of their c share with the opened r times their right share, the opened s times their left share, and the opened value $r \cdot s$.

For the remaining two gate types, < 0 and $>> c$, we must take additional care because there are also MAC values attached to each wire, so simply running Algorithms 2 and 3 will not suffice. For example, if we are computing a < 0 gate, then we want to compute authenticated shares of $O = (L_A + L_B) < 0$, where O is 0 if it is false and 1 if it is true. The dealer will precompute both possible outputs: an authenticated sharing of 0 and of 1. However, in order to preserve privacy, we cannot reveal to either party which is the sharing of 0 and which is the sharing of 1, otherwise they would learn the output. Thus, we have a "flip" bit f, such that the output is flipped if f is 1. The parties share f: Alice holds f_A, Bob holds f_B such that $f = f_A \oplus f_B$. Thus, if we compute $((L_A + L_B) < 0) \oplus f_A \oplus f_B$ and reveal this Boolean result to both parties, it properly indexes which authenticated share they should each use, and will be a correct sharing of O with a MAC.

We model $LTZ = ((L_A + L_B) < 0) \oplus f_A \oplus f_B$ as a Boolean circuit: the addition is done via a straightforward adder-with-carry circuit, the comparison just looks at the sign bit, and it finishes with two XOR gates. In order to securely evaluate it, we use the garbled circuit methodology with precomputation. The dealer generates a garbled circuit and key $(G, gsk) = GC(LTZ)$. The dealer also generates the garbled 0/1 labels for each input bit by calling the projective input garbler $GI(i, 0, gsk), GI(i, 1, gsk)$ for all $i < |input|$. The dealer sends G to Bob, and the wire labels (garbled input keys) to Alice. It also pregenerates random OTs on behalf of Alice and Bob (i.e. it sends (r_0, r_1) to one party and (b, r_b) to the other for random strings r_0, r_1 and random bit b), so that the keys can be obliviously selected by Bob. During the online phase, Alice selects the keys corresponding to her inputs and sends them to Bob, and uses oblivious transfer to select his keys. Bob then evaluates G and sends back to Alice the result of G, which will then allow them both to select a pregenerated authenticated sharing of the correct bit (generated above).

For shift-by-constant gates, there are multiple output bits, and we treat each bit individually and reconstruct the output via the standard bits-to-integer $\sum a_i 2^i$ transformation, performed on the authenticated shares.

6 Benchmarks

6.1 Internal Testing and Benchmarks

Our implementation takes 20 integer inputs, each representing real numbers as follows: an integer n represents the real number $\frac{n}{2^{20}}$. These 20 inputs correspond to the satellite trajectories of Party 1 and Party 2, namely

Gate	Number
Input	20
Output	1
Constant	353
Addition	36435
Subtraction	95823
Multiplication	101574
Less-than-zero	32270
Shift-by-constant	526
Total	267002

Figure 6: The number of gates in the optimized circuit

the vectors for position (x_i, y_i, z_i), velocity $(v_{x_i}, v_{y_i}, v_{z_i})$, error $(\sigma_{x_i}, \sigma_{y_i}, \sigma_{z_i})$, and a radius R_i for $i = 1, 2$. The output is a single integer representing the probability of collision in the same format.

We spent a significant amount of effort optimizing the circuit, and after optimization, the final circuit contained 2.67×10^5 gates (see Figure 6 for a breakdown).

6.1.1 Bounds and Error Tolerance

The goal of a conjunction analysis calculation is to facilitate decision-making and improve space situational awareness. Because our system is providing a numerical approximation to an integral without a closed form, it is important to ensure that the approximations provide sufficient accuracy to inform decision making. As noted in [Alf07], due to the many simplifications used in practice both in mathematical terms (integral approximation, rounding) and physical terms (curvature of the earth, relativistic effects), at some point, the error in the conjunction analysis becomes dominated by these approximations. For example, if the true collision probability is 0.00000001, it may be sufficient to obtain an estimate that has a fairly large relative error. On the other hand, if the true collision probability is .15, then it a much higher (relative) accuracy may be necessary to inform decision-making. The probability region between 10^{-1} and 10^{-7} is called the "operational decision region," [Alf07], and it is most important to maximize the (relative) accuracy for probabilities within this region. Alfano suggests that within this region, estimating the probability to within 2 significant figures (which translates to roughly 1% relative error), should be enough to ensure accurate decision making.

Internally, we evaluated the circuit on 2000 test cases and obtained the following error bounds:

	Absolute error $\lvert approx - true \rvert$	Relative error $\left\lvert \frac{approx - true}{true} \right\rvert$
Min	9.000×10^{-10}	1.240×10^{-8}
Max	8.512×10^{-4}	19 (see remark below)
Avg	7.027×10^{-5}	9.623×10^{-3}

Error tolerance across all tests

Remark: In all cases where the relative error was extremely large, the true probabilities were extremely close to zero, and the large relative error would not impact decision making, e.g. cases where the "true" probability was 10^{-15}, and we estimated $2 \cdot 10^{-14}$.

If we focus on the operational decision region, we obtain much better bounds on the relative error.

	Absolute error $\lvert approx - true \rvert$	Relative error $\left\lvert \frac{approx - true}{true} \right\rvert$
Min	9.000×10^{-10}	2.190×10^{-6}
Max	9.608×10^{-5}	2.543×10^{-3}
Avg	5.759×10^{-7}	1.208×10^{-5}

6.1.2 Benchmarks

We benchmarked our system on a virtual machine running 64bit CentOS 6.4. The machine had 8GB of RAM, 4 processors, and 100GB of hard disk space.

We tested both pregeneration time and online run time between two parties for a single conjunction, as well as disk space used. We present the average runtime over 2000 tests as well as bounds on the maximum and minimum times.

Pregeneration Time	
real	1m30s $\pm$ 15s
user	38s $\pm$ 15s
sys	27s $\pm$ 25s

The pregeneration phase also generates four files (two for each party): `index0, index1, share0, share1`. The files `index0, index1` index the gate information and are 555 kilobytes each. The files `share0` and `share1` are contain the pregenerated garbled circuits, OTs and SISO-M shares needed for the computation. The computation requires a huge number of these shares, and the file `share0` is 5.2 gigabytes in size, while `share1` is 3.8 gigabytes. The file `share0` is larger because this file contains the pre-generated garbled circuits. These share files are needed for each computation, and are all consumed during the online phase.

Once the private inputs are learned, the online phase can be computed in the times below:

Online Time	
real	5m2s $\pm$ 15s
user	4m59s $\pm$ 15s
sys	15s $\pm$ 10s

The fact that the real time is so much larger than the system time indicates that most of the time is due to disk I/O, thus if the share files are stored in memory (say on a more powerful machine), then the evaluation could be performed much faster.

6.2 Comparison with the Sharemind Implementation [KW14a]

Secure computation of conjunction analyses was one of the target use-cases for MPC in DARPA's PROCEED program, and consequently it has been used as a benchmark for other secure computation systems. Our implementation was done in parallel with that of Kamm and Willemson[KW14a], but the design and goals of these projects are fundamentally different. We focused on building a two-party protocol that was custom-tailored to the problem of securely computing conjunction analyses. Kamm and Willemson focused on building an IEEE 754 floating point library for the three-party Sharemind system, and used the conjunction analysis application to exhibit the capabilities of their general platform.

Below, we highlight some of the main architectural differences between our implementation and that of Kamm and Willemson. Our implementation is a two-party system built using GMW and garbled circuits, whereas the Sharemind implementation is a three-party system based on the BGW protocol. Our system uses public-key cryptography to ensure user privacy, whereas the BGW protocol is an information-theoretic protocol that requires that at least two out of the three servers are non-colluding to ensure privacy. Our system provides security against a malicious adversary with correlated abort, whereas the Sharemind system only provides security against semi-honest adversaries. Our system uses fixed precision arithmetic, whereas the Sharemind system implements IEEE 754 floating point arithmetic.

Kamm and Willemson [KW14a] report a computation time of 3m24s using single-precision floating point operations and 4m6s using double-precision floats when distributed across 3 machines, each with 12 cores running at 2.93 GHz and 48 GB of memory.

7 Conclusion

In this work, we described the design and implementation of a custom secure computation protocol to compute the probability of satellite collisions. The underlying collision probability calculation requires numerically estimating a complicated integral – something that was, until recently, beyond the reach of secure computation techniques. We envision that our techniques can extend to other areas that require secure numerical computations as well. In order to improve efficiency, we custom built and optimized an augmented arithmetic circuit to estimate the collision probabilities. We constructed a secure computation protocol then used a combination of GMW and garbled circuits to evaluate these augmented circuits, which allowed us to evaluate our conjunction analysis circuit on the participants' private inputs. Our secure computation works in the offline/online model, where during the offline phase the parties work to generate a large amount of correlated randomness in the form of OTs and shared arithmetic triples. Later, in the online phase, this correlated randomness is consumed to facilitate the secure computation. Because of the sheer quantity of correlated randomness needed, the limiting factor in our computation was disk I/O.

This work provides positive evidence for the fact that MPC technology is now capable of evaluating very complex functions securely and efficiently.

References

[ABPP15] David W Archer, Dan Bogdanov, Benny Pinkas, and Pille Pullonen. Maturity and Performance of Programmable Secure Computation. https://eprint.iacr.org/2015/1039, 2015.

[ABZS13] Mehrdad Aliasgari, Marina Blanton, Yihua Zhang, and Aaron Steele. Secure computation on floating point numbers. In *NDSS 2013*. The Internet Society, February 2013.

[Alf05] Salvatore Alfano. A Numerical Implementation of Spherical Object Collision Probability. *Journal of the Astronautical Sciences*, 53(1):103–109, 2005.

[Alf07] Salvatore Alfano. Review of Conjunction Probability Methods for Short-term Encounters. In *Proceedings of the AAS/AIAA Space Flight Mechanics Meeting*, volume 127 PART 1, pages 719–746, feb 2007.

[App13] Benny Applebaum. Garbling XOR gates "for free" in the standard model. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 162–181. Springer, Heidelberg, March 2013.

[AS65] Milton Abramowitz and Irene A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 1965.

[BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 169–188. Springer, Heidelberg, May 2011.

[Bea95] Donald Beaver. Precomputing oblivious transfer. In Don Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 97–109. Springer, Heidelberg, August 1995.

[Bea97] Donald Beaver. Commodity-based cryptography (extended abstract). In *29th ACM STOC*, pages 446–455. ACM Press, May 1997.

[BHKR13] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient Garbling from a Fixed-Key Blockcipher. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 478–492. Dept. of Comput. Sci. & Eng., Univ. of California, San Diego, La Jolla, CA, USA, IEEE, may 2013.

[BHR12a] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Heidelberg, December 2012.

[BHR12b]   Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 12*, pages 784–796. ACM Press, October 2012.

[BLW08]   Dan Bogdanov, Sven Laur, and Jan Willemson. Sharemind: A framework for fast privacy-preserving computations. In Sushil Jajodia and Javier López, editors, *ESORICS 2008*, volume 5283 of *LNCS*, pages 192–206. Springer, Heidelberg, October 2008.

[BOGW88]   Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988.

[CCD88]   David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (abstract) (informal contribution). In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, page 462. Springer, Heidelberg, August 1988.

[CR08]   Sylvain Chevillard and Nathalie Revol. Computation of the error function erf in arbitrary precision with correct rounding. In J. D. Bruguera and M. Daumas, editors, *In RNC 8, the 8th Conference on Real Numbers and Computers*, pages 27–36, July 2008.

[CS10]   Octavian Catrina and Amitabh Saxena. Secure computation with fixed-point numbers. In Radu Sion, editor, *FC 2010*, volume 6052 of *LNCS*, pages 35–50. Springer, Heidelberg, January 2010.

[DGKN09]   Ivan Damgård, Martin Geisler, Mikkel Krøigaard, and Jesper Buus Nielsen. Asynchronous multiparty computation: Theory and implementation. In Stanislaw Jarecki and Gene Tsudik, editors, *PKC 2009*, volume 5443 of *LNCS*, pages 160–179. Springer, Heidelberg, March 2009.

[DPSZ12]   Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662. Springer, Heidelberg, August 2012.

[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, February 2015.

[EGL82]   Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 205–210. Plenum Press, New York, USA, 1982.

[FSW03]   Pierre-Alain Fouque, Jacques Stern, and Jan-Geert Wackers. Cryptocomputing with rationals. In Matt Blaze, editor, *FC 2002*, volume 2357 of *LNCS*, pages 136–146. Springer, Heidelberg, March 2003.

[Gen09]   Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.

[GHL+14]   Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled RAM revisited. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 405–422. Springer, Heidelberg, May 2014.

[GMW87]   Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

[Gol01]   Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001.

[Gol04]   Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.

[HAO10] Robert Hall, Salvatore Alfano, and Alan Ocampo. Advances in Satellite Conjunction Analysis. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2010.

[HEKM11] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster Secure Two-Party Computation Using Garbled Circuits. In *In USENIX Security Symposium*, 2011.

[HKS+10] Wilko Henecka, Stefan Kögl, Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. TASTY: tool for automating secure two-party computations. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 10*, pages 451–462. ACM Press, October 2010.

[HW15] Brett Hemenway and William Welser. Cryptographers Could Prevent Satellite Collisions. *Scientific American*, (February):28–29, feb 2015.

[HWIB14] Brett Hemenway, William Welser IV, and Dave Baiocchi. Achieving Higher-Fidelity Conjunction Analyses Using Cryptography to Improve Information Sharing. Technical report, RAND Corporation, 2014.

[IKM+13] Yuval Ishai, Eyal Kushilevitz, Sigurd Meldgaard, Claudio Orlandi, and Anat Paskin-Cherniavsky. On the power of correlated randomness in secure computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 600–620. Springer, Heidelberg, March 2013.

[IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 145–161. Springer, Heidelberg, August 2003.

[IKO+11] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 406–425. Springer, Heidelberg, May 2011.

[IPS09] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Secure arithmetic computation with no honest majority. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 294–314. Springer, Heidelberg, March 2009.

[KMR14] Vladimir Kolesnikov, Payman Mohassel, and Mike Rosulek. FleXOR: Flexible garbling for XOR gates that beats free-XOR. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 440–457. Springer, Heidelberg, August 2014.

[KS08] Vladimir Kolesnikov and Thomas Schneider. Improved Garbled Circuit: Free XOR Gates and Applications. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part II*, ICALP '08, pages 486–498, Berlin, Heidelberg, 2008. Springer-Verlag.

[KSS12] Benjamin Kreuter, Abhi Shelat, and Chih H Shen. Billion-gate secure computation with malicious adversaries. In *Proceedings of the 21st USENIX conference on Security symposium*, Security'12, page 14, Berkeley, CA, USA, 2012. USENIX Association.

[KW14a] Liina Kamm and Jan Willemson. Secure floating point arithmetic and private satellite collision analysis. *International Journal of Information Security*, pages 1–18, 2014.

[KW14b] Toomas Krips and Jan Willemson. Hybrid model of fixed and floating point numbers in secure multiparty computations. In Sherman S. M. Chow, Jan Camenisch, Lucas Chi Kwong Hui, and Siu-Ming Yiu, editors, *ISC 2014*, volume 8783 of *LNCS*, pages 179–197. Springer, Heidelberg, October 2014.

[LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 719–734. Springer, Heidelberg, May 2013.

[LP09]   Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009.

[LWN+15]   Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. ObliVM: A programming framework for secure computation. In *2015 IEEE Symposium on Security and Privacy*, pages 359–376. IEEE Computer Society Press, May 2015.

[Mal11]   Lior Malka. VMCrypt: modular software architecture for scalable secure computation. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 11*, pages 715–724. ACM Press, October 2011.

[MGC+16]   Benjamin Mood, Debayan Gupta, Henry Carter, Kevin Butler, and Patrick Traynor. Frigate: A Validated, Extensible, and Efficient Compiler and Interpreter for Secure Computation. In *Proceedings of the IEEE European Symposium on Security and Privacy (Euro S&P)*, 2016.

[NNOB12]   Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A New Approach to Practical Active-Secure Two-Party Computation. *CRYPTO*, 7417 LNCS:681–700, 2012.

[PS15]   Pille Pullonen and Sander Siim. Combining secret sharing and garbled circuits for efficient private IEEE 754 floating-point computations. In Michael Brenner, Nicolas Christin, Benjamin Johnson, and Kurt Rohloff, editors, *FC 2015 Workshops*, volume 8976 of *LNCS*, pages 172–183. Springer, Heidelberg, January 2015.

[Rab05]   Michael O. Rabin. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187, 2005. `http://eprint.iacr.org/2005/187`.

[RHH14]   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*, pages 655–670. IEEE Computer Society Press, May 2014.

[Sha84]   Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 47–53. Springer, Heidelberg, August 1984.

[UCS15]   Union of concerned scientists. http://www.ucsusa.org/, 2015. Accessed 2015-11-09.

[VO09]   Associated Press Veronika Oleksyn. What a mess! experts ponder space junk problem. *USA Today*, February 2009.

[Yao82]   Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, November 1982.

[Yao86]   Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.

[YY12]   Ching-Hua Yu and Bo-Yin Yang. Probabilistically correct secure arithmetic computation for modular conversion, zero test, comparison, MOD and exponentiation. In Ivan Visconti and Roberto De Prisco, editors, *SCN 12*, volume 7485 of *LNCS*, pages 426–444. Springer, Heidelberg, September 2012.

[ZRE15]   Samee Zahur, Mike Rosulek, and David Evans. Two Halves Make a Whole. In *CRYPTO*, volume 9057, pages 220–250, 2015.

Appendix

A   Details on Projecting into the Encounter Plane

Calculating Encounter Plane:

Recall the inputs to the calculation are, two velocity vectors $\mathbf{v}_a,\mathbf{v}_b$ in $\mathbb{R}^3$, two covariance matrices $\mathbf{C}_a,\mathbf{C}_b$ in $\mathbb{R}^{3\times3}$, two position vectors $\mathbf{p}_a, \mathbf{p}_b$ in $\mathbb{R}^3$ and two radii R_a, R_b in $\mathbb{R}$.

The encounter plane is defined to be the plane perpendicular to the relative velocity of the objects, containing the point $\mathbf{p}_a$. The point $\mathbf{p}_a$ will be taken as the origin in the encounter plane. Calculating a basis for the encounter plane can then be done as follows:

- Define the relative velocity $\mathbf{v}_r = \mathbf{v}_b - \mathbf{v}_a$

- Define the encounter coordinates

$$\mathbf{i} = \frac{\mathbf{v}_r}{|\mathbf{v}_r|}, \quad \mathbf{j} = \frac{\mathbf{v}_b \times \mathbf{v}_a}{|\mathbf{v}_b \times \mathbf{v}_a|}, \quad \mathbf{k} = \mathbf{i} \times \mathbf{j}$$

  Thus the encounter plane is the $\mathbf{j}$-$\mathbf{k}$ plane, and $\mathbf{j}, \mathbf{k}$ form an orthornormal basis for this plane.

- Let

$$\mathbf{P} = \begin{bmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \end{bmatrix}$$

  Then $\mathbf{P}$ is the change-of-basis matrix that takes the standard basis to the $(\mathbf{i}, \mathbf{j}, \mathbf{k})$-basis, and the matrix $\mathbf{P}$ is orthogonal, i.e., $\mathbf{P}^T\mathbf{P} = I$.

Projecting Into the Encounter Plane:

Once the encounter plane is defined, we project the entire ellipsoid into the encounter plane.

- Let $\mathbf{Q} = \begin{bmatrix} \mathbf{j} & \mathbf{k} \end{bmatrix}$. Then $\mathbf{Q}$ is the 3×2 matrix that takes vectors in the encounter plane (written in the $\mathbf{j}$-$\mathbf{k}$ basis) to vectors in $\mathbb{R}^3$ in the standard basis. The 2×3 matrix $\mathbf{Q}^T$ projects vectors in $\mathbb{R}^3$ to the encounter plane (relative to the $\mathbf{j} - \mathbf{k}$ basis).

- Then the covariance matrix, projected into the encounter plane becomes

$$\mathbf{C} = \mathbf{Q}^T(\mathbf{C}_a + \mathbf{C}_b)\mathbf{Q}$$

  This $\mathbf{j}$-$\mathbf{k}$ basis is not convenient to work in, however, because we would like to choose a basis for the encounter plane where the basis vectors are aligned with the principle axes of the ellipse. Thus we need to diagonalize the (symmetric) matrix $\mathbf{C}$. See Appendix B for further discussion.

- Let $\mathbf{U}$ be the 2×2 matrix whose columns are orthonormal eigenvectors of $\mathbf{C}$, thus

$$\mathbf{U^T C U} = \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix}$$

  The eigenvalues of $\mathbf{C}$ are the scalars σ_x^2, σ_y^2. Because $\mathbf{C}$ is a covariance matrix it will be positive-definite, so the square roots of its eigenvalues will be real. See Appendix C for notes on calculating eigenvectors and eigenvalues of a symmetric 2×2 matrix.

- We are only interested in the relative positions of the objects, translating, we find that the center of the "hardbody" is at $\mathbf{p}_b - \mathbf{p}_a$ (in $\mathbb{R}^3$). Projecting this onto the encounter plane (multiplying by $\mathbf{Q}^T$), and then putting it in the eigenvector basis (multiplying by $\mathbf{U}^T$), we have

$$\begin{bmatrix} x_m \\ y_m \end{bmatrix} = \mathbf{U}^T\mathbf{Q}^T(\mathbf{p}_b - \mathbf{p}_a)$$

B   Mathematics Ellipses and Ellipsoids

In this section we review some of the basic mathematical facts concerning ellipses.

B.1   Definition of an Ellipse

An ellipsoid whose principle axes are parallel to the coordinate axes and centered at the origin is given by an equation

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1 \tag{3}$$

Where a, b, c are the lengths of the semi-principle axes of the ellipse.[3] We can write equation 3 concisely as

$$\mathbf{x}^T \mathbf{A} \mathbf{x} = 1 \tag{4}$$

where

$$\mathbf{x} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} a^2 & 0 & 0 \\ 0 & b^2 & 0 \\ 0 & 0 & c^2 \end{pmatrix}$$

Since $\mathbf{x}^T \mathbf{A} \mathbf{x} = \mathbf{x}^T \mathbf{A}^T \mathbf{x}$, then $\mathbf{x}^T \mathbf{A} \mathbf{x} = \mathbf{x}^T ((\mathbf{A} + \mathbf{A}^T)/2)\mathbf{x}$. Since $(\mathbf{A} + \mathbf{A}^T)/2$ is symmetric, when dealing with quadratic forms, we can always assume $\mathbf{A}$ is symmetric.

This leads to a simple formulation of the formula for an ellipse when the principle axes of the ellipse are not aligned with the coordinate axes. If $\mathbf{A}$ is a symmetric matrix, then the equation $\mathbf{x}^T \mathbf{A} \mathbf{x} = 1$ is the equation of an ellipse where the eigenvectors of $\mathbf{A}$ represent the principle axes of the ellipse and the inverses of the eigenvalues are the squares of the lengths of the corresponding semi-principle axes. Note that because $\mathbf{A}$ is symmetric, the spectral theorem tells us that the eigenvalues of $\mathbf{A}$ are orthogonal (which we would expect since they are the principle axes of the ellipsoid).

B.2   Slicing an Ellipsoid by a Plane

If $\mathbf{u}$ and $\mathbf{v}$ are orthonormal vectors spanning a plane in $\mathbb{R}^3$, then the matrix

$$\mathbf{B} = \begin{pmatrix} | & | \\ \mathbf{u} & \mathbf{v} \\ | & | \end{pmatrix}$$

transforms vectors in the $(\mathbf{u}, \mathbf{v})$-plane to vectors in $\mathbb{R}^3$.

If the (3×3) matrix $\mathbf{A}$ defines an ellipsoid, then vectors in the $(\mathbf{u}, \mathbf{v})$ plane are in the ellipsoid exactly when they satisfy the equation

$$\mathbf{z}^T \mathbf{B}^T \mathbf{A} \mathbf{B} \mathbf{z} = 1 \tag{5}$$

where $\mathbf{z}$ is a vector in $\mathbb{R}^2$ representing the point in $(\mathbf{u}, \mathbf{v})$-plane, written in the $\{\mathbf{u}, \mathbf{v}\}$ basis.

If $\mathbf{A}$ is a symmetric matrix, then $\mathbf{B}^T \mathbf{A} \mathbf{B}$ is also symmetric, so equation 5 represents an ellipse (in the two dimensional plane defined by $\mathbf{u}, \mathbf{v}$).

The eigenvectors of $\mathbf{B}^T \mathbf{A} \mathbf{B}$ represent the directions of the principle axes of the ellipse (in the $(\mathbf{u}, \mathbf{v})$-basis) and the eigenvalues of $\mathbf{B}^T \mathbf{A} \mathbf{B}$ represent the inverses of the squares of the lengths of the semi-principle axes.

---

[3]If the center of the ellipsoid is at the point $p = (p_x, p_y, p_z)$, then the equation becomes

$$\frac{(x - p_x)^2}{a^2} + \frac{(y - p_y)^2}{b^2} + \frac{(z - p_z)^2}{c^2} = 1,$$

but we will focus on ellipsoids centered at 0.

B.3 Ellipsoids of Multivariate Normals

The density function for a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{C}$, is given by

$$\frac{1}{\sqrt{(2\pi)^k |\mathbf{C}|}} \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{C}^{-1}\mathbf{x}\right)$$

The argument to exp is actually in the form of an ellipse $(\mathbf{x}^T \mathbf{C}^{-1}\mathbf{x})$. In fact, the equation $\mathbf{x}^T \mathbf{C}^{-1}\mathbf{x} = 1$ is an equation of an ellipse centered at 0, extending one variance in each direction.

Thus the square roots of the eigenvalues of $\mathbf{C}$ are the lengths of the semi-principle axes of the ellipsoid that extends one standard deviation in each deviation.

C   Eigenvectors and Eigenvalues and of 2×2 Symmetric Matrices

Calculating the semi-principle axes of the projected ellipsoid requires calculating the eigenvectors and eigenvalues of the symmetrix 2×2 matrix $\mathbf{Q}^T(\mathbf{C}_b + \mathbf{C}_a)\mathbf{Q}$. To facilitate calculations, the equations for the eigenvectors of a 2×2 symmetric matrix are given below.

If $\mathbf{A} = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$, then the eigenvectors of $\mathbf{A}$ are

$$\begin{bmatrix} \frac{2b}{d-a\pm\sqrt{(d-a)^2+(2b)^2}} \\ 1 \end{bmatrix}$$

with eigenvalues $\frac{1}{2}\left[d + a \pm \sqrt{(d-a)^2 + (2b)^2}\right]$.

Normalizing the eigenvectors, we find the two normalized eigenvectors are

$$\begin{bmatrix} \frac{2b}{\left(d-a\pm\sqrt{(d-a)^2+(2b)^2}\right)^2+(2b)^2} \\ \frac{d-a\pm\sqrt{(d-a)^2+(2b)^2}}{\left(d-a\pm\sqrt{(d-a)^2+(2b)^2}\right)^2+(2b)^2} \end{bmatrix}$$

D   Taylor Expansions of erf

A natural choice for representing erf as a circuit is to use the Taylor expansion. We can Taylor expand $\text{erf}(\cdot)$ as

$$\text{erf}\,x = \frac{2}{\sqrt{\pi}}\sum_{n=0}^{\infty}\frac{(-1)^n x^{2n+1}}{n!(2n+1)},$$

which, by [CR08] has absolute error

$$\left|\text{erf}\,x - \frac{2}{\sqrt{\pi}}\sum_{n=0}^{N}\frac{(-1)^n x^{2n+1}}{n!(2n+1)}\right| \leq \frac{2}{\sqrt{\pi}}\frac{x^{2N+1}}{N!(2N+1)}.$$

Unfortunately, the error function has horizontal asymptotes (see Figure 7, a property that is not shared by any polynomial. Thus Taylor expansions for erf fare poorly for large inputs (see Figure 8).

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2}\, dt$$

Figure 7: The error function

$$\left| \mathrm{erf}(x) - \sum_{n=0}^{50} \frac{(-1)^n x^{2n+1}}{n!(2n+1)} \right|$$

Figure 8: The absolute error of the 50 term Taylor expansion for erf.