

Workgroup: CFRG
Internet-Draft: draft-irtf-cfrg-vdaf-00
Published: 27 April 2022
Intended Status: Informational
Expires: 29 October 2022
Authors: R. L. Barnes C. Patton P. Schoppmann
 Cisco Cloudflare, Inc. Google
Verifiable Distributed Aggregation Functions

Abstract

This document describes Verifiable Distributed Aggregation Functions (VDAFs), a family of multi-party protocols for computing aggregate statistics over user measurements. These protocols are designed to ensure that, as long as at least one aggregation server executes the protocol honestly, individual measurements are never seen by any server in the clear. At the same time, VDAFs allow the servers to detect if a malicious or misconfigured client submitted an input that would result in an incorrect aggregate result.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Crypto Forum Research Group mailing list (cfrg@ietf.org), which is archived at https://mailarchive.ietf.org/arch/search/?email_list=cfrg.

Source for this draft and an issue tracker can be found at <https://github.com/cjpatton/vdaf>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 29 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Conventions and Definitions](#)
- [3. Overview](#)
- [4. Definition of VDAFs](#)
 - [4.1. Setup](#)
 - [4.2. Sharding](#)
 - [4.3. Preparation](#)
 - [4.4. Aggregation](#)
 - [4.5. Unsharding](#)
 - [4.6. Execution of a VDAF](#)
- [5. Preliminaries](#)
 - [5.1. Finite Fields](#)
 - [5.1.1. Auxiliary Functions](#)
 - [5.1.2. FFT-Friendly Fields](#)
 - [5.1.3. Parameters](#)
 - [5.2. Pseudorandom Generators](#)
 - [5.2.1. PrgAes128](#)
- [6. Prio3](#)
 - [6.1. Fully Linear Proof \(FLP\) Systems](#)
 - [6.1.1. Encoding the Input](#)
 - [6.2. Construction](#)
 - [6.2.1. Setup](#)
 - [6.2.2. Sharding](#)
 - [6.2.3. Preparation](#)
 - [6.2.4. Aggregation](#)
 - [6.2.5. Unsharding](#)
 - [6.2.6. Auxiliary Functions](#)
 - [6.3. A General-Purpose FLP](#)
 - [6.3.1. Overview](#)
 - [6.3.2. Validity Circuits](#)
 - [6.3.3. Construction](#)

- [6.4. Instantiations](#)
 - [6.4.1. Prio3Aes128Count](#)
 - [6.4.2. Prio3Aes128Sum](#)
 - [6.4.3. Prio3Aes128Histogram](#)
- [7. Poplar1](#)
 - [7.1. Incremental Distributed Point Functions \(IDPFs\)](#)
 - [7.2. Construction](#)
 - [7.2.1. Setup](#)
 - [7.2.2. Preparation](#)
 - [7.2.3. Aggregation](#)
 - [7.2.4. Unsharding](#)
 - [7.2.5. Helper Functions](#)
- [8. Security Considerations](#)
- [9. IANA Considerations](#)
- [10. References](#)
 - [10.1. Normative References](#)
 - [10.2. Informative References](#)
- [Acknowledgments](#)
- [Test Vectors](#)
 - [Prio3Aes128Count](#)
 - [Prio3Aes128Sum](#)
 - [Prio3Aes128Histogram](#)
- [Authors' Addresses](#)

1. Introduction

The ubiquity of the Internet makes it an ideal platform for measurement of large-scale phenomena, whether public health trends or the behavior of computer systems at scale. There is substantial overlap, however, between information that is valuable to measure and information that users consider private.

For example, consider an application that provides health information to users. The operator of an application might want to know which parts of their application are used most often, as a way to guide future development of the application. Specific users' patterns of usage, though, could reveal sensitive things about them, such as which users are researching a given health condition.

In many situations, the measurement collector is only interested in aggregate statistics, e.g., which portions of an application are most used or what fraction of people have experienced a given disease. Thus systems that provide aggregate statistics while protecting individual measurements can deliver the value of the measurements while protecting users' privacy.

Most prior approaches to this problem fall under the rubric of "differential privacy (DP)" [Dwo06]. Roughly speaking, a data aggregation system that is differentially private ensures that the

degree to which any individual measurement influences the value of the aggregate result can be precisely controlled. For example, in systems like RAPPOR [[EPK14](#)], each user samples noise from a well-known distribution and adds it to their input before submitting to the aggregation server. The aggregation server then adds up the noisy inputs, and because it knows the distribution from whence the noise was sampled, it can estimate the true sum with reasonable precision.

Differentially private systems like RAPPOR are easy to deploy and provide a useful guarantee. On its own, however, DP falls short of the strongest privacy property one could hope for. Specifically, depending on the "amount" of noise a client adds to its input, it may be possible for a curious aggregator to make a reasonable guess of the input's true value. Indeed, the more noise the clients add, the less reliable will be the server's estimate of the output. Thus systems employing DP techniques alone must strike a delicate balance between privacy and utility.

The ideal goal for a privacy-preserving measurement system is that of secure multi-party computation: No participant in the protocol should learn anything about an individual input beyond what it can deduce from the aggregate. In this document, we describe Verifiable Distributed Aggregation Functions (VDAFs) as a general class of protocols that achieve this goal.

VDAF schemes achieve their privacy goal by distributing the computation of the aggregate among a number of non-colluding aggregation servers. As long as a subset of the servers executes the protocol honestly, VDAFs guarantee that no input is ever accessible to any party besides the client that submitted it. At the same time, VDAFs are "verifiable" in the sense that malformed inputs that would otherwise garble the output of the computation can be detected and removed from the set of inputs.

The cost of achieving these security properties is the need for multiple servers to participate in the protocol, and the need to ensure they do not collude to undermine the VDAF's privacy guarantees. Recent implementation experience has shown that practical challenges of coordinating multiple servers can be overcome. The Prio system [[CGB17](#)] (essentially a VDAF) has been deployed in systems supporting hundreds of millions of users: The Mozilla Origin Telemetry project [[OriginTelemetry](#)] and the Exposure Notification Private Analytics collaboration among the Internet Security Research Group (ISRG), Google, Apple, and others [[ENPA](#)].

The VDAF abstraction laid out in [Section 4](#) represents a class of multi-party protocols for privacy-preserving measurement proposed in the literature. These protocols vary in their operational and

security considerations, sometimes in subtle but consequential ways. This document therefore has two important goals:

1. Providing applications like [[I-D.draft-gpew-priv-ppm](#)] with a simple, uniform interface for accessing privacy-preserving measurement schemes, and documenting relevant operational and security bounds for that interface:
 1. General patterns of communications among the various actors involved in the system (clients, aggregation servers, and measurement collectors);
 2. Capabilities of a malicious coalition of servers attempting to divulge information about client inputs; and
 3. Conditions that are necessary to ensure that malicious clients cannot corrupt the computation.
2. Providing cryptographers with design criteria that allow new constructions to be easily used by applications.

This document also specifies two concrete VDAF schemes, each based on a protocol from the literature.

*The aforementioned Prio system [[CGB17](#)] allows for the privacy-preserving computation of a variety aggregate statistics. The basic idea underlying Prio is fairly simple:

1. Each client shards its input into a sequence of additive shares and distributes the shares among the aggregation servers.
2. Next, each server adds up its shares locally, resulting in an additive share of the aggregate.
3. Finally, the aggregation servers combine their additive shares to obtain the final aggregate.

The difficult part of this system is ensuring that the servers hold shares of a valid input, e.g., the input is an integer in a specific range. Thus Prio specifies a multi-party protocol for accomplishing this task.

In [Section 6](#) we describe Prio3, a VDAF that follows the same overall framework as the original Prio protocol, but incorporates techniques introduced in [[BBCGGI19](#)] that result in significant performance gains.

*More recently, Boneh et al. [[BBCGGI21](#)] described a protocol called Poplar for solving the t-heavy-hitters problem in a

privacy-preserving manner. Here each client holds a bit-string of length n , and the goal of the aggregation servers is to compute the set of inputs that occur at least t times. The core primitive used in their protocol is a generalization of a Distributed Point Function (DPF) [GI14] that allows the servers to "query" their DPF shares on any bit-string of length shorter than or equal to n . As a result of this query, each of the servers has an additive share of a bit indicating whether the string is a prefix of the client's input. The protocol also specifies a multi-party computation for verifying that at most one string among a set of candidates is a prefix of the client's input.

In [Section 7](#) we describe a VDAF called Poplar1 that implements this functionality.

The remainder of this document is organized as follows: [Section 3](#) gives a brief overview of VDAFs; [Section 4](#) defines the syntax for VDAFs; [Section 5](#) defines various functionalities that are common to our constructions; [Section 6](#) describes the Prio3 construction; [Section 7](#) describes the Poplar1 construction; and [Section 8](#) enumerates the security considerations for VDAFs.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Algorithms in this document are written in Python 3. Type hints are used to define input and output types. A fatal error in a program (e.g., failure to parse one of the function parameters) is usually handled by raising an exception.

Some common functionalities:

*zeros(len: Unsigned) -> Bytes returns an array of zero bytes. The length of output **MUST** be len.

*gen_rand(len: Unsigned) -> Bytes returns an array of random bytes. The length of output **MUST** be len.


*byte(int: Unsigned) -> Bytes returns the representation of int as a byte string. The value of int **MUST** be in [0,256).

*xor(left: Bytes, right: Bytes) -> Bytes returns the bitwise XOR of left and right. An exception is raised if the inputs are not the same length.

*I2OSP and OS2IP from [RFC8017], which are used, respectively, to convert a non-negative integer to a byte string and convert a byte string to a non-negative integer.

*next_power_of_2(n: Unsigned) -> Unsigned returns the smallest integer greater than or equal to n that is also a power of two.

3. Overview



Input shares Aggregate shares

Figure 1: Overall data flow of a VDAF

In a VDAF-based private measurement system, we distinguish three types of actors: Clients, Aggregators, and Collectors. The overall flow of the measurement process is as follows:

*Clients are configured with public parameters for a set of Aggregators.

*To submit an individual measurement, the Client shards the measurement into "input shares" and sends one input share to each Aggregator.

*The Aggregators verify the validity of the input shares, producing a set of "output shares".

-Output shares are in one-to-one correspondence with the input shares.

- Just as each Aggregator receives one input share of each input, at the end of the validation process, each aggregator holds one output share.

- In most VDAFs, Aggregators will need to exchange information among themselves as part of the validation process.

- *Each Aggregator combines the output shares across inputs in the batch to compute the "aggregate share" for that batch, i.e., its share of the desired aggregate result.

- *The Aggregators submit their aggregate shares to the Collector, who combines them to obtain the aggregate result over the batch.

Aggregators are a new class of actor relative to traditional measurement systems where clients submit measurements to a single server. They are critical for both the privacy properties of the system and the correctness of the measurements obtained. The privacy properties of the system are assured by non-collusion among Aggregators, and Aggregators are the entities that perform validation of client inputs. Thus clients trust Aggregators not to collude (typically it is required that at least one Aggregator is honest), and Collectors trust Aggregators to properly verify Client inputs.

Within the bounds of the non-collusion requirements of a given VDAF instance, it is possible for the same entity to play more than one role. For example, the Collector could also act as an Aggregator, effectively using the other Aggregators to augment a basic client-server protocol.

In this document, we describe the computations performed by the actors in this system. It is up to applications to arrange for the required information to be delivered to the proper actors in the proper sequence. In general, we assume that all communications are confidential and mutually authenticated, with the exception that Clients submitting measurements may be anonymous.

4. Definition of VDAFs

A concrete VDAF specifies the algorithms involved in evaluating an aggregation function across a batch of inputs. This section specifies the interfaces of these algorithms as they would be exposed to applications.

The overall execution of a VDAF comprises the following steps:

- *Setup - Generating shared parameters for the Aggregators

- *Sharding - Computing input shares from an individual measurement

*Preparation - Conversion and verification of input shares to output shares compatible with the aggregation function being computed

*Aggregation - Combining a sequence of output shares into an aggregate share

*Unsharding - Combining a sequence of aggregate shares into an aggregate result

The setup algorithm is performed once for a given collection of Aggregators. Sharding and preparation are done once per measurement input. Aggregation and unsharding are done over a batch of inputs (more precisely, over the output shares recovered from those inputs).

Note that the preparation step performs two functions: Verification and conversion. Conversion translates input shares into output shares that are compatible with the aggregation function. Verification ensures that aggregating the recovered output shares will not lead to a garbled aggregate result.

The remainder of this section defines the VDAF interface in terms of an abstract base class `Vdaf`. This class defines the set of methods and attributes a concrete VDAF must provide. The attributes are listed in [Table 1](#); the methods are defined in the subsections that follow.

Parameter	Description
ROUNDS	Number of rounds of communication during the preparation phase (Section 4.3)
SHARES	Number of input shares into which each measurement is sharded (Section 4.2)
Measurement	Type of each measurement
PublicParam	Type of public parameter used by the Client during the sharding phase (Section 4.2)
VerifyParam	Type of verification parameter used by each Aggregator during the preparation phase (Section 4.3)
AggParam	Type of aggregation parameter
Prep	State of each Aggregator during the preparation phase (Section 4.3)
OutShare	Type of each output share
AggShare	Type of each aggregate share
AggResult	Type of the aggregate result

Table 1: Constants and types defined by each concrete VDAF.

4.1. Setup

Before execution of the VDAF can begin, it is necessary to distribute long-lived parameters to the Client and Aggregators. The long-lived parameters are generated by the following algorithm:

`*Vdaf.setup()` -> (PublicParam, Vec[VerifyParam]) is the randomized setup algorithm used to generate the public parameter used by the Clients and the verification parameters used by the Aggregators. The length of the latter **MUST** be equal to SHARES. In general, an Aggregator's verification parameter is considered secret and **MUST NOT** be revealed to the Clients, Collector or other Aggregators. The parameters **MAY** be reused across multiple VDAF evaluations. See [Section 8](#) for a discussion of the security implications this has depending on the threat model.

4.2. Sharding

In order to protect the privacy of its measurements, a VDAF Client splits its measurements into "input shares". The `measurement_to_input_shares` method is used for this purpose.

`*Vdaf.measurement_to_input_shares(public_param: PublicParam, input: Measurement)` -> Vec[Bytes] is the randomized input-distribution algorithm run by each Client. It consumes the public parameter and input measurement and produces a sequence of input shares, one for each Aggregator. The length of the output **MUST** be SHARES.

Client

=====

measurement

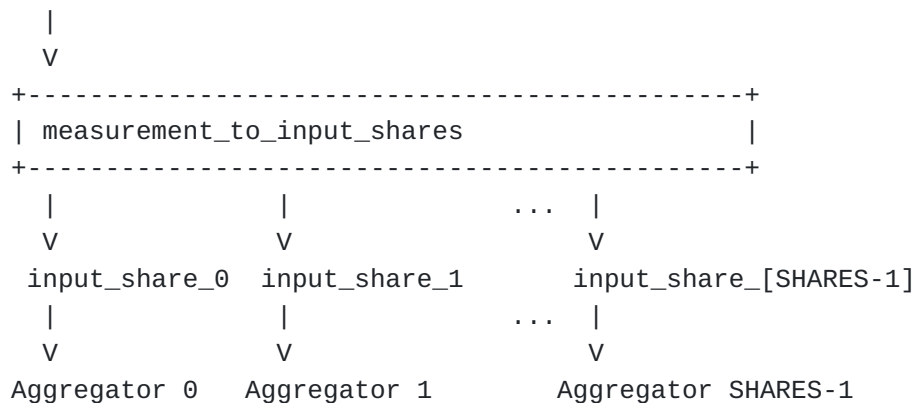


Figure 2: The Client divides its measurement input into input shares and distributes them to the Aggregators.

CP The `public_param` is intended to allow for protocols that require the Client to use a public key for sharding its measurement. When rotating the `verify_param` for such a scheme, it would be necessary to also update the `public_param` with which the clients are configured. For PPM it would be nice if we could rotate the `verify_param` without also having to update the clients. We should consider dropping this at some point. See <https://github.com/cjpatton/vdaf/issues/19>.

4.3. Preparation

To recover and verify output shares, the Aggregators interact with one another over `ROUNDS` rounds. Prior to each round, each Aggregator constructs an outbound message. Next, the sequence of outbound messages is combined into a single message, called a "preparation message". (Each of the outbound messages are called "preparation-message shares".) Finally, the preparation message is distributed to the Aggregators to begin the next round.

An Aggregator begins the first round with its input share and it begins each subsequent round with the previous preparation message. Its output in the last round is its output share and its output in each of the preceding rounds is a preparation-message share.

This process involves a value called the "aggregation parameter" used to map the input shares to output shares. The Aggregators need to agree on this parameter before they can begin preparing inputs for aggregation.




Figure 3: VDAF preparation process on the input shares for a single measurement. At the end of the computation, each Aggregator holds an output share or an error.

To facilitate the preparation process, a concrete VDAF implements the following class methods:

```
*Vdaf.prep_init(verify_param: VerifyParam, agg_param: AggParam,
nonce: Bytes, input_share: Bytes) -> Prep
```

is the deterministic preparation-state initialization algorithm run by each Aggregator to begin processing its input share into an output share. Its inputs are its verification parameter (`verify_param`), the aggregation parameter (`agg_param`), the nonce provided by the environment (`nonce`, see [Figure 6](#)), and one of the input shares generated by the client (`input_share`). Its output is the Aggregator's initial preparation state.

`*Vdaf.prep_next(prepare: Prep, inbound: Optional[Bytes]) -> Union[Tuple[Prep, Bytes], OutShare]` is the deterministic preparation-state update algorithm run by each Aggregator. It updates the Aggregator's preparation state (`prepare`) and returns either its next preparation state and its message share for the current round or, if this is the last round, its output share. An exception is raised if a valid output share could not be recovered. The input of this algorithm is the inbound preparation message or, if this is the first round, `None`.

`*Vdaf.prep_shares_to_prep(agg_param: AggParam, prep_shares: Vec[Bytes]) -> Bytes` is the deterministic preparation-message pre-processing algorithm. It combines the preparation-message shares generated by the Aggregators in the previous round into the preparation message consumed by each in the next round.

In effect, each Aggregator moves through a linear state machine with `ROUNDS+1` states. The Aggregator enters the first state on using the initialization algorithm, and the update algorithm advances the Aggregator to the next state. Thus, in addition to defining the number of rounds (`ROUNDS`), a VDAF instance defines the state of the Aggregator after each round.

TODO Consider how to bake this "linear state machine" condition into the syntax. Given that Python 3 is used as our pseudocode, it's easier to specify the preparation state using a class.

The preparation-state update accomplishes two tasks: recovery of output shares from the input shares and ensuring that the recovered output shares are valid. The abstraction boundary is drawn so that an Aggregator only recovers an output share if it is deemed valid (at least, based on the Aggregator's view of the protocol). Another way to draw this boundary would be to have the Aggregators recover output shares first, then verify that they are valid. However, this would allow the possibility of misusing the API by, say, aggregating an invalid output share. Moreover, in protocols like Prio+ [\[AGJOP21\]](#) based on oblivious transfer, it is necessary for the Aggregators to interact in order to recover aggregatable output shares at all.

Note that it is possible for a VDAF to specify `ROUNDS == 0`, in which case each Aggregator runs the preparation-state update algorithm once and immediately recovers its output share without interacting with the other Aggregators. However, most, if not all, constructions will require some amount of interaction in order to ensure validity of the output shares (while also maintaining privacy).

OPEN ISSUE Depending on what we do for issue#20, we may end up needing to revise the above paragraph.

4.4. Aggregation

Once an Aggregator holds output shares for a batch of measurements (where batches are defined by the application), it combines them into a share of the desired aggregate result. This algorithm is performed locally at each Aggregator, without communication with the other Aggregators.

*Vdaf.out_shares_to_agg_share(agg_param: AggParam, output_shares: Vec[OutShare]) -> agg_share: AggShare is the deterministic aggregation algorithm. It is run by each Aggregator over the output shares it has computed over a batch of measurement inputs.

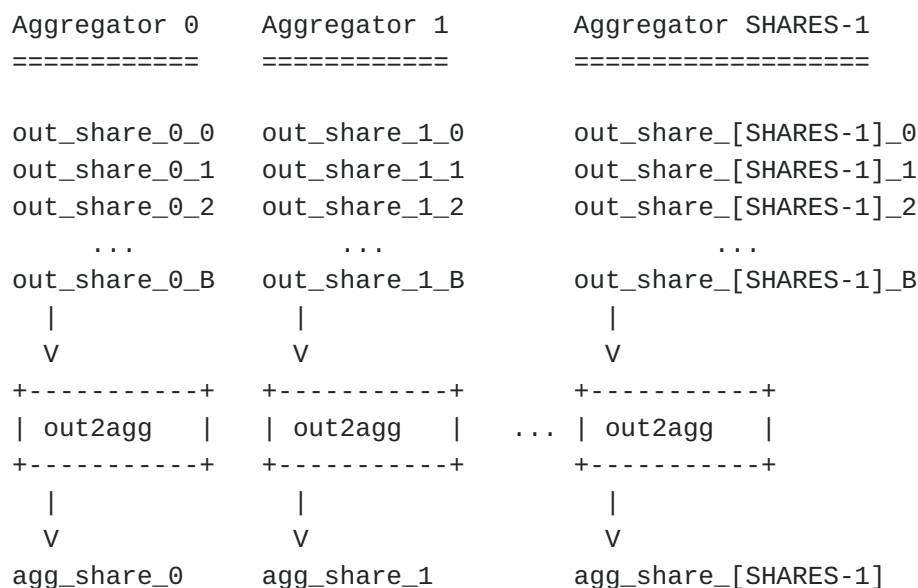


Figure 4: Aggregation of output shares. `B` indicates the number of measurements in the batch.

For simplicity, we have written this algorithm and the unsharding algorithm below in "one-shot" form, where all shares for a batch are provided at the same time. Some VDAFs may also support a "streaming" form, where shares are processed one at a time.

4.5. Unsharding

After the Aggregators have aggregated a sufficient number of output shares, each sends its aggregate share to the Collector, who runs the following algorithm to recover the following output:

*Vdaf.agg_shares_to_result(agg_param: AggParam, agg_shares: Vec[AggShare]) -> AggResult is run by the Collector in order to compute the aggregate result from the Aggregators' shares. The length of agg_shares **MUST** be SHARES. This algorithm is deterministic.

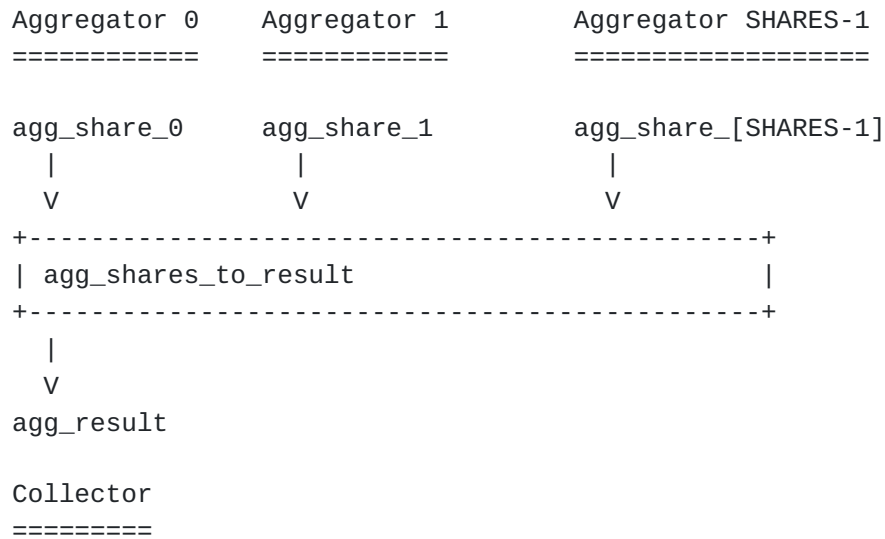


Figure 5: Computation of the final aggregate result from aggregate shares.

QUESTION Maybe the aggregation algorithms should be randomized in order to allow the Aggregators (or the Collector) to add noise for differential privacy. (See the security considerations of [I-D.draft-gpew-priv-ppm](#).) Or is this out-of-scope of this document?

4.6. Execution of a VDAF

Executing a VDAF involves the concurrent evaluation of the VDAF on individual inputs and aggregation of the recovered output shares. This is captured by the following example algorithm:

```

def run_vdaf(Vdaf,
            agg_param: Vdaf.AggParam,
            nonces: Vec[Bytes],
            measurements: Vec[Vdaf.Measurement]):
    # Distribute long-lived parameters.
    (public_param, verify_params) = Vdaf.setup()

    out_shares = []
    for (nonce, measurement) in zip(nonces, measurements):

        # Each Client shards its input into shares.
        input_shares = Vdaf.measurement_to_input_shares(public_param,
                                                         measurement)

        # Each Aggregator initializes its preparation state.
        prep_states = []
        for j in range(Vdaf.SHARES):
            state = Vdaf.prep_init(verify_params[j],
                                   agg_param,
                                   nonce,
                                   input_shares[j])
            prep_states.append(state)

        # Aggregators recover their output shares.
        inbound = None
        for i in range(Vdaf.ROUNDS+1):
            outbound = []
            for j in range(Vdaf.SHARES):
                out = Vdaf.prep_next(prep_states[j], inbound)
                if i < Vdaf.ROUNDS:
                    (prep_states[j], out) = out
                outbound.append(out)
            # This is where we would send messages over the
            # network in a distributed VDAF computation.
            if i < Vdaf.ROUNDS:
                inbound = Vdaf.prep_shares_to_prep(agg_param,
                                                    outbound)

        # The final outputs of prepare phase are the output shares.
        out_shares.append(outbound)

    # Each Aggregator aggregates its output shares into an
    # aggregate share.
    agg_shares = []
    for j in range(Vdaf.SHARES):
        out_shares_j = [out[j] for out in out_shares]
        agg_share_j = Vdaf.out_shares_to_agg_share(agg_param,
                                                    out_shares_j)
        agg_shares.append(agg_share_j)

```



```
# Collector unshards the aggregate.  
agg_result = Vdaf.agg_shares_to_result(agg_param, agg_shares)  
return agg_result
```

Figure 6: Execution of a VDAF.

The inputs to this algorithm are the aggregation parameter `agg_param`, a list of nonces `nonces`, and a batch of Client inputs `input_batch`. The aggregation parameter is chosen by the Aggregators prior to executing the VDAF. This document does not specify how the nonces are chosen, but security requires that the nonces be unique for each measurement. See [Section 8](#) for details.

Another important question this document leaves out of scope is how a VDAF is to be executed by Aggregators distributed over a real network. Algorithm `run_vdaf` prescribes the protocol's execution in a "benign" environment in which there is no adversary and messages are passed among the protocol participants over secure point-to-point channels. In reality, these channels need to be instantiated by some "wrapper protocol", such as [\[I-D.draft-gpew-priv-ppm\]](#) that implements suitable cryptographic functionalities. Moreover, some fraction of the Aggregators (or Clients) may be malicious and diverge from their prescribed behaviors. [Section 8](#) describes the execution of the VDAF in various adversarial environments and what properties the wrapper protocol needs to provide in each.

5. Preliminaries

This section describes the primitives that are common to the VDAFs specified in this document.

5.1. Finite Fields

Both Prio3 and Poplar1 use finite fields of prime order. Finite field elements are represented by a class `Field` with the following associated parameters:

*`MODULUS`: Unsigned is the prime modulus that defines the field.

*`ENCODED_SIZE`: Unsigned is the number of bytes used to encode a field element as a byte string.

A concrete `Field` also implements the following class methods:

*`Field.zeros(length: Unsigned) -> output: Vec[Field]` returns a vector of zeros. The length of output **MUST** be `length`.

*`Field.rand_vec(length: Unsigned) -> output: Vec[Field]` returns a vector of random field elements. The length of output **MUST** be `length`.

A field element is an instance of a concrete `Field`. The concrete class defines the usual arithmetic operations on field elements. In

addition, it defines the following instance method for converting a field element to an unsigned integer:

```
*elem.as_unsigned() -> Unsigned returns the integer representation
of field element elem.
```

Likewise, each concrete Field implements a constructor for converting an unsigned integer into a field element:

```
*Field(integer: Unsigned) returns integer represented as a field
element. If integer >= Field.MODULUS, then integer is first
reduced modulo Field.MODULUS.
```

Finally, each concrete Field has two derived class methods, one for encoding a vector of field elements as a byte string and another for decoding a vector of field elements.

```
def encode_vec(Field, data: Vec[Field]) -> Bytes:
    encoded = Bytes()
    for x in data:
        encoded += I2OSP(x.as_unsigned(), Field.ENCODED_SIZE)
    return encoded

def decode_vec(Field, encoded: Bytes) -> Vec[Field]:
    L = Field.ENCODED_SIZE
    if len(encoded) % L != 0:
        raise ERR_DECODE

    vec = []
    for i in range(0, len(encoded), L):
        encoded_x = encoded[i:i+L]
        x = Field(OS2IP(encoded_x))
        vec.append(x)
    return vec
```

Figure 7: Derived class methods for finite fields.

5.1.1. Auxiliary Functions

The following auxiliary functions on vectors of field elements are used in the remainder of this document. Note that an exception is raised by each function if the operands are not the same length.

```

# Compute the inner product of the operands.
def inner_product(left: Vec[Field], right: Vec[Field]) -> Field:
    return sum(map(lambda x: x[0] * x[1], zip(left, right)))

# Subtract the right operand from the left and return the result.
def vec_sub(left: Vec[Field], right: Vec[Field]):
    return list(map(lambda x: x[0] - x[1], zip(left, right)))

# Add the right operand to the left and return the result.
def vec_add(left: Vec[Field], right: Vec[Field]):
    return list(map(lambda x: x[0] + x[1], zip(left, right)))

```

Figure 8: Common functions for finite fields.

5.1.2. FFT-Friendly Fields

Some VDAFs require fields that are suitable for efficient computation of the discrete Fourier transform. (One example is Prio3 ([Section 6](#)) when instantiated with the generic FLP of [Section 6.3.3](#).) Specifically, a field is said to be "FFT-friendly" if, in addition to satisfying the interface described in [Section 5.1](#), it implements the following method:

*Field.gen() -> Field returns the generator of a large subgroup of the multiplicative group.

FFT-friendly fields also define the following parameter:

*GEN_ORDER: Unsigned is the order of a multiplicative subgroup generated by Field.gen(). This value **MUST** be a power of 2.

5.1.3. Parameters

The tables below define finite fields used in the remainder of this document.

Parameter	Value
MODULUS	$2^{32} * 4294967295 + 1$
ENCODED_SIZE	8
Generator	$7^{4294967295}$
GEN_ORDER	2^{32}

Table 2: Field64, an FFT-friendly field.

Parameter	Value
MODULUS	$2^{66} * 4611686018427387897 + 1$
ENCODED_SIZE	16
Generator	$7^{4611686018427387897}$

Parameter	Value
GEN_ORDER	2 ⁶⁶

Table 3: Field128, an FFT-friendly field.

5.2. Pseudorandom Generators

A pseudorandom generator (PRG) is used to expand a short, (pseudo)random seed into a long string of pseudorandom bits. A PRG suitable for this document implements the interface specified in this section. Concrete constructions are described in the subsections that follow.

PRGs are defined by a class `Prg` with the following associated parameter:

*SEED_SIZE: Unsigned is the size (in bytes) of a seed.

A concrete `Prg` implements the following class method:

*`Prg(seed: Bytes, info: Bytes)` constructs an instance of `Prg` from the given seed and info string. The seed **MUST** be of length SEED_SIZE and **MUST** be generated securely (i.e., it is either the output of `gen_rand` or a previous invocation of the PRG). The info string is used for domain separation.

*`prg.next(length: Unsigned)` returns the next length bytes of output of PRG. If the seed was securely generated, the output can be treated as pseudorandom.

Each `Prg` has two derived class methods. The first is used to derive a fresh seed from an existing one. The second is used to compute a sequence of pseudorandom field elements. For each method, the seed **MUST** be of length SEED_SIZE and **MUST** be generated securely (i.e., it is either the output of `gen_rand` or a previous invocation of the PRG).

```

# Derive a new seed.
def derive_seed(Prg, seed: Bytes, info: Bytes) -> bytes:
    prg = Prg(seed, info)
    return prg.next(Prg.SEED_SIZE)

# Expand a seed into a vector of Field elements.
def expand_into_vec(Prg,
                    Field,
                    seed: Bytes,
                    info: Bytes,
                    length: Unsigned):
    m = next_power_of_2(Field.MODULUS) - 1
    prg = Prg(seed, info)
    vec = []
    while len(vec) < length:
        x = OS2IP(prg.next(Field.ENCODDED_SIZE))
        x &= m
        if x < Field.MODULUS:
            vec.append(Field(x))
    return vec

```

Figure 9: Derived class methods for PRGs.

5.2.1. PrgAes128

OPEN ISSUE Phillipp points out that a fixed-key mode of AES may be more performant (<https://eprint.iacr.org/2019/074.pdf>). See <https://github.com/cjpatton/vdaf/issues/32> for details.

Our first construction, PrgAes128, converts a blockcipher, namely AES-128, into a PRG. Seed expansion involves two steps. In the first step, CMAC [[RFC4493](#)] is applied to the seed and info string to get a fresh key. In the second step, the fresh key is used in CTR-mode to produce a key stream for generating the output. A fixed initialization vector (IV) is used.

```

class PrgAes128:

    SEED_SIZE: Unsigned = 16

    def __init__(self, seed, info):
        self.length_consumed = 0

        # Use CMAC as a pseudorandom function to derive a key.
        self.key = AES128-CMAC(seed, info)

    def next(self, length):
        self.length_consumed += length

        # CTR-mode encryption of the all-zero string of the desired
        # length and using a fixed, all-zero IV.
        stream = AES128-CTR(key, zeros(16), zeros(self.length_consumed))
        return stream[-length:]

```

Figure 10: Definition of PRG PrgAes128.

6. Prio3

NOTE This construction has not undergone significant security analysis.

This section describes "Prio3", a VDAF for Prio [CGB17]. Prio is suitable for a wide variety of aggregation functions, including (but not limited to) sum, mean, standard deviation, estimation of quantiles (e.g., median), and linear regression. In fact, the scheme described in this section is compatible with any aggregation function that has the following structure:

*Each measurement is encoded as a vector over some finite field.

*Input validity is determined by an arithmetic circuit evaluated over the encoded input. (An "arithmetic circuit" is a function comprised of arithmetic operations in the field.) The circuit's output is a single field element: if zero, then the input is said to be "valid"; otherwise, if the output is non-zero, then the input is said to "invalid".

*The aggregate result is obtained by summing up the encoded input vectors and computing some function of the sum.

At a high level, Prio3 distributes this computation as follows. Each Client first shards its measurement by first encoding it, then splitting the vector into secret shares and sending a share to each Aggregator. Next, in the preparation phase, the Aggregators carry out a multi-party computation to determine if their shares correspond to a valid input (as determined by the arithmetic

circuit). This computation involves a "proof" of validity generated by the Client. Next, each Aggregator sums up its input shares locally. Finally, the Collector sums up the aggregate shares and computes the aggregate result.

This VDAF does not have an aggregation parameter. Instead, the output share is derived from the input share by applying a fixed map. See [Section 7](#) for an example of a VDAF that makes meaningful use of the aggregation parameter.

As the name implies, "Prio3" is a descendant of the original Prio construction. A second iteration was deployed in the [\[ENPA\]](#) system, and like the VDAF described here, the ENPA system was built from techniques introduced in [\[BBCGGI19\]](#) that significantly improve communication cost. That system was specialized for a particular aggregation function; the goal of Prio3 is to provide the same level of generality as the original construction.

The core component of Prio3 is a "Fully Linear Proof (FLP)" system. Introduced by [\[BBCGGI19\]](#), the FLP encapsulates the functionality required for encoding and validating inputs. Prio3 can be thought of as a transformation of a particular class of FLPs into a VDAF.

The remainder of this section is structured as follows. The syntax for FLPs is described in [Section 6.1](#). The generic transformation of an FLP into Prio3 is specified in [Section 6.2](#). Next, a concrete FLP suitable for any validity circuit is specified in [Section 6.3](#). Finally, instantiations of Prio3 for various types of measurements are specified in [Section 6.4](#). Test vectors can be found in [Appendix "Test Vectors"](#).

6.1. Fully Linear Proof (FLP) Systems

Conceptually, an FLP is a two-party protocol executed by a prover and a verifier. In actual use, however, the prover's computation is carried out by the Client, and the verifier's computation is distributed among the Aggregators. The Client generates a "proof" of its input's validity and distributes shares of the proof to the Aggregators. Each Aggregator then performs some a computation on its input share and proof share locally and sends the result to the other Aggregators. Combining the exchanged messages allows each Aggregator to decide if it holds a share of a valid input. (See [Section 6.2](#) for details.)

As usual, we will describe the interface implemented by a concrete FLP in terms of an abstract base class `Flp` that specifies the set of methods and parameters a concrete FLP must provide.

The parameters provided by a concrete FLP are listed in [Table 4](#).

Parameter	Description
PROVE_RAND_LEN	Length of the prover randomness, the number of random field elements consumed by the prover when generating a proof
QUERY_RAND_LEN	Length of the query randomness, the number of random field elements consumed by the verifier
JOINT_RAND_LEN	Length of the joint randomness, the number of random field elements consumed by both the prover and verifier
INPUT_LEN	Length of the encoded measurement (Section 6.1.1)
OUTPUT_LEN	Length of the aggregatable output (Section 6.1.1)
PROOF_LEN	Length of the proof
VERIFIER_LEN	Length of the verifier message generated by querying the input and proof
Measurement	Type of the measurement
Field	As defined in (Section 5.1)

Table 4: Constants and types defined by a concrete FLP.

An FLP specifies the following algorithms for generating and verifying proofs of validity (encoding is described below in [Section 6.1.1](#)):

`*Flp.prove(input: Vec[Field], prove_rand: Vec[Field], joint_rand: Vec[Field]) -> Vec[Field]` is the deterministic proof-generation algorithm run by the prover. Its inputs are the encoded input, the "prover randomness" `prove_rand`, and the "joint randomness" `joint_rand`. The proof randomness is used only by the prover, but the joint randomness is shared by both the prover and verifier.

`*Flp.query(input: Vec[Field], proof: Vec[Field], query_rand: Vec[Field], joint_rand: Vec[Field]) -> Vec[Field]` is the query-generation algorithm run by the verifier. This is used to "query" the input and proof. The result of the query (i.e., the output of this function) is called the "verifier message". In addition to the input and proof, this algorithm takes as input the query randomness `query_rand` and the joint randomness `joint_rand`. The former is used only by the verifier.

`*Flp.decide(verifier: Vec[Field]) -> Bool` is the deterministic decision algorithm run by the verifier. It takes as input the verifier message and outputs a boolean indicating if the input from which it was generated is valid.

Our application requires that the FLP is "fully linear" in the sense defined in [\[BBCGGI19\]](#) As a practical matter, what this property implies is that, when run on a share of the input and proof, the query-generation algorithm outputs a share of the verifier message. Furthermore, the "zero-knowledge" property of the FLP system ensures that the verifier message reveals nothing about the input's

validity. Therefore, to decide if an input is valid, the Aggregators will run the query-generation algorithm locally, exchange verifier shares, combine them to recover the verifier message, and run the decision algorithm.

An FLP is executed by the prover and verifier as follows:

```
def run_flp(Flp, inp: Vec[Flp.Field], num_shares: Unsigned):
    joint_rand = Flp.Field.rand_vec(Flp.JOINT_RAND_LEN)
    prove_rand = Flp.Field.rand_vec(Flp.PROVE_RAND_LEN)
    query_rand = Flp.Field.rand_vec(Flp.QUERY_RAND_LEN)

    # Prover generates the proof.
    proof = Flp.prove(inp, prove_rand, joint_rand)

    # Verifier queries the input and proof.
    verifier = Flp.query(inp, proof, query_rand, joint_rand, num_shares)

    # Verifier decides if the input is valid.
    return Flp.decide(verifier)
```

Figure 11: Execution of an FLP.

The proof system is constructed so that, if input is a valid input, then `run_flp(Flp, input)` always returns `True`. On the other hand, if input is invalid, then as long as `joint_rand` and `query_rand` are generated uniform randomly, the output is `False` with overwhelming probability.

We remark that [BBCGG19] defines a much larger class of fully linear proof systems than we consider here. In particular, what is called an "FLP" here is called a 1.5-round, public-coin, interactive oracle proof system in their paper.

6.1.1. Encoding the Input

The type of measurement being aggregated is defined by the FLP. Hence, the FLP also specifies a method of encoding raw measurements as a vector of field elements:

```
*Flp.encode(measurement: Measurement) -> Vec[Field] encodes a raw
measurement as a vector of field elements. The return value MUST
be of length INPUT_LEN.
```

For some FLPs, the encoded input also includes redundant field elements that are useful for checking the proof, but which are not needed after the proof has been checked. An example is the "integer sum" data type from [CGB17] in which an integer in range $[0, 2^k)$ is encoded as a vector of k field elements (this type is also defined

in [Section 6.4](#)). After consuming this vector, all that is needed is the integer it represents. Thus the FLP defines an algorithm for truncating the input to the length of the aggregated output:

```
*Flp.truncate(input: Vec[Field]) -> Vec[Field] maps an encoded
input to an aggregatable output. The length of the input MUST be
INPUT_LEN and the length of the output MUST be OUTPUT_LEN.
```

We remark that, taken together, these two functionalities correspond roughly to the notion of "Affine-aggregatable encodings (AFEs)" from [\[CGB17\]](#).

6.2. Construction

This section specifies Prio3, an implementation of the Vdaf interface ([Section 4](#)). It has two generic parameters: an Flp ([Section 6.1](#)) and a Prg ([Section 5.2](#)). The associated constants and types required by the Vdaf interface are defined in [Table 5](#). The methods required for sharding, preparation, aggregation, and unsharding are described in the remaining subsections.

Parameter	Value
ROUNDS	1
SHARES	in [2, 255)
Measurement	Flp.Measurement
PublicParam	None
VerifyParam	Tuple[Unsigned, Bytes]
AggParam	None
Prep	Tuple[Vec[Flp.Field], Optional[Bytes], Bytes]
OutShare	Vec[Flp.Field]
AggShare	Vec[Flp.Field]
AggResult	Vec[Unsigned]

Table 5: Associated parameters for the Prio3 VDAF.

6.2.1. Setup

The setup algorithm generates a symmetric key shared by all of the Aggregators. The key is used to derive query randomness for the FLP query-generation algorithm run by the Aggregators during preparation. An Aggregator's verification parameter also includes its "ID", a unique integer in [0, SHARES).

```
def setup(Prio3):
    k_query_init = gen_rand(Prio3.Prg.SEED_SIZE)
    verify_param = [(j, k_query_init) for j in range(Prio3.SHARES)]
    return (None, verify_param)
```

Figure 12: The setup algorithm for Prio3.

6.2.2. Sharding

Recall from [Section 6.1](#) that the FLP syntax calls for "joint randomness" shared by the prover (i.e., the Client) and the verifier (i.e., the Aggregators). VDAFs have no such notion. Instead, the Client derives the joint randomness from its input in a way that allows the Aggregators to reconstruct it from their input shares. (This idea is based on the Fiat-Shamir heuristic and is described in Section 6.2.3 of [\[BBCGG19\]](#).)

The input-distribution algorithm involves the following steps:

1. Encode the Client's raw measurement as an input for the FLP
2. Shard the input into a sequence of input shares
3. Derive the joint randomness from the input shares
4. Run the FLP proof-generation algorithm using the derived joint randomness
5. Shard the proof into a sequence of proof shares

The algorithm is specified below. Notice that only one set input and proof shares (called the "leader" shares below) are vectors of field elements. The other shares (called the "helper" shares) are represented instead by PRG seeds, which are expanded into vectors of field elements.

The code refers to a pair of auxiliary functions for encoding the shares. These are called `encode_leader_share` and `encode_helper_share` respectively and they are described in [Section 6.2.6](#).

```

def measurement_to_input_shares(Prio3, _public_param, measurement):
    dst = b"vdaf-00 prio3"
    inp = Prio3.Flp.encode(measurement)
    k_joint_rand = zeros(Prio3.Prg.SEED_SIZE)

    # Generate input shares.
    leader_input_share = inp
    k_helper_input_shares = []
    k_helper_blinds = []
    k_helper_hints = []
    for j in range(Prio3.SHARES-1):
        k_blind = gen_rand(Prio3.Prg.SEED_SIZE)
        k_share = gen_rand(Prio3.Prg.SEED_SIZE)
        helper_input_share = Prio3.Prg.expand_into_vec(
            Prio3.Flp.Field,
            k_share,
            dst + byte(j+1),
            Prio3.Flp.INPUT_LEN
        )
        leader_input_share = vec_sub(leader_input_share,
                                     helper_input_share)
        encoded = Prio3.Flp.Field.encode_vec(helper_input_share)
        k_hint = Prio3.Prg.derive_seed(k_blind,
                                     byte(j+1) + encoded)
        k_joint_rand = xor(k_joint_rand, k_hint)
        k_helper_input_shares.append(k_share)
        k_helper_blinds.append(k_blind)
        k_helper_hints.append(k_hint)
    k_leader_blind = gen_rand(Prio3.Prg.SEED_SIZE)
    encoded = Prio3.Flp.Field.encode_vec(leader_input_share)
    k_leader_hint = Prio3.Prg.derive_seed(k_leader_blind,
                                     byte(0) + encoded)
    k_joint_rand = xor(k_joint_rand, k_leader_hint)

    # Finish joint randomness hints.
    for j in range(Prio3.SHARES-1):
        k_helper_hints[j] = xor(k_helper_hints[j], k_joint_rand)
    k_leader_hint = xor(k_leader_hint, k_joint_rand)

    # Generate the proof shares.
    prove_rand = Prio3.Prg.expand_into_vec(
        Prio3.Flp.Field,
        gen_rand(Prio3.Prg.SEED_SIZE),
        dst,
        Prio3.Flp.PROVE_RAND_LEN
    )
    joint_rand = Prio3.Prg.expand_into_vec(
        Prio3.Flp.Field,
        k_joint_rand,

```

```

        dst,
        Prio3.Flp.JOINT_RAND_LEN
    )
    proof = Prio3.Flp.prove(inp, prove_rand, joint_rand)
    leader_proof_share = proof
    k_helper_proof_shares = []
    for j in range(Prio3.SHARES-1):
        k_share = gen_rand(Prio3.Prg.SEED_SIZE)
        k_helper_proof_shares.append(k_share)
        helper_proof_share = Prio3.Prg.expand_into_vec(
            Prio3.Flp.Field,
            k_share,
            dst + byte(j+1),
            Prio3.Flp.PROOF_LEN
        )
        leader_proof_share = vec_sub(leader_proof_share,
                                      helper_proof_share)

    input_shares = []
    input_shares.append(Prio3.encode_leader_share(
        leader_input_share,
        leader_proof_share,
        k_leader_blind,
        k_leader_hint,
    ))
    for j in range(Prio3.SHARES-1):
        input_shares.append(Prio3.encode_helper_share(
            k_helper_input_shares[j],
            k_helper_proof_shares[j],
            k_helper_blinds[j],
            k_helper_hints[j],
        ))
    return input_shares

```

Figure 13: Input-distribution algorithm for Prio3.

6.2.3. Preparation

This section describes the process of recovering output shares from the input shares. The high-level idea is that each Aggregator first queries its input and proof share locally, then exchanges its verifier share with the other Aggregators. The verifier shares are then combined into the verifier message, which is used to decide whether to accept.

In addition, the Aggregators must ensure that they have all used the same joint randomness for the query-generation algorithm. The joint randomness is generated by a PRG seed. Each Aggregator derives an XOR secret share of this seed from its input share and the "blind" generated by the client. Thus, before running the query-generation algorithm, it must first gather the XOR secret shares derived by the other Aggregators.

In order to avoid extra round of communication, the Client sends each Aggregator a "hint" equal to the XOR of the other Aggregators' shares of the joint randomness seed. This leaves open the possibility that the Client cheated by, say, forcing the Aggregators to use joint randomness that biases the proof check procedure some way in its favor. To mitigate this, the Aggregators also check that they have all computed the same joint randomness seed before accepting their output shares. To do so, they exchange their XOR shares of the PRG seed along with their verifier shares.

NOTE This optimization somewhat diverges from Section 6.2.3 of [\[BBCGGI19\]](#). Security analysis is needed.

The algorithms required for preparation are defined as follows. These algorithms make use of encoding and decoding methods defined in [Section 6.2.6](#).

```

def prep_init(Prio3, verify_param, _agg_param, nonce, input_share):
    dst = b"vdaf-00 prio3"
    (j, k_query_init) = verify_param

    (input_share, proof_share, k_blind, k_hint) = \
        Prio3.decode_leader_share(input_share) if j == 0 else \
        Prio3.decode_helper_share(dst, j, input_share)

    out_share = Prio3.Flp.truncate(input_share)

    k_query_rand = Prio3.Prg.derive_seed(k_query_init,
                                         byte(255) + nonce)
    query_rand = Prio3.Prg.expand_into_vec(
        Prio3.Flp.Field,
        k_query_rand,
        dst,
        Prio3.Flp.QUERY_RAND_LEN
    )
    joint_rand, k_joint_rand, k_joint_rand_share = [], None, None
    if Prio3.Flp.JOINT_RAND_LEN > 0:
        encoded = Prio3.Flp.Field.encode_vec(input_share)
        k_joint_rand_share = Prio3.Prg.derive_seed(k_blind,
                                                  byte(j) + encoded)
        k_joint_rand = xor(k_hint, k_joint_rand_share)
        joint_rand = Prio3.Prg.expand_into_vec(
            Prio3.Flp.Field,
            k_joint_rand,
            dst,
            Prio3.Flp.JOINT_RAND_LEN
        )
    verifier_share = Prio3.Flp.query(input_share,
                                     proof_share,
                                     query_rand,
                                     joint_rand,
                                     Prio3.SHARES)

    prep_msg = Prio3.encode_prepare_message(verifier_share,
                                             k_joint_rand_share)
    return (out_share, k_joint_rand, prep_msg)

def prep_next(Prio3, prep, inbound):
    (out_share, k_joint_rand, prep_msg) = prep

    if inbound is None:
        return (prep, prep_msg)

    (verifier, k_joint_rand_check) = \
        Prio3.decode_prepare_message(inbound)

    if k_joint_rand_check != k_joint_rand or \

```


Figure 14: Preparation state for Prio3.

6.2.4. Aggregation

Aggregating a set of output shares is simply a matter of adding up the vectors element-wise.

```
def out_shares_to_agg_share(Prio3, _agg_param, out_shares):
    agg_share = Prio3.Flp.Field.zeros(Prio3.Flp.OUTPUT_LEN)
    for out_share in out_shares:
        agg_share = vec_add(agg_share, out_share)
    return agg_share
```

Figure 15: Aggregation algorithm for Prio3.

6.2.5. Unsharding

To unshard a set of aggregate shares, the Collector first adds up the vectors element-wise. It then converts each element of the vector into an integer.

```
def agg_shares_to_result(Prio3, _agg_param, agg_shares):
    agg = Prio3.Flp.Field.zeros(Prio3.Flp.OUTPUT_LEN)
    for agg_share in agg_shares:
        agg = vec_add(agg, agg_share)
    return list(map(lambda x: x.as_unsigned(), agg))
```

Figure 16: Computation of the aggregate result for Prio3.

6.2.6. Auxiliary Functions

[illegible]

```

k_proof_share, encoded = encoded[:1], encoded[1:]
proof_share = Prio3.Prg.expand_into_vec(Prio3.Flp.Field,
                                         k_proof_share,
                                         dst + byte(j),
                                         Prio3.Flp.PROOF_LEN)

k_blind, k_hint = None, None
if Prio3.Flp.JOINT_RAND_LEN > 0:
    k_blind, encoded = encoded[:1], encoded[1:]
    k_hint, encoded = encoded[:1], encoded[1:]
if len(encoded) != 0:
    raise ERR_DECODE
return (input_share, proof_share, k_blind, k_hint)

def encode_prepare_message(Prio3, verifier, k_joint_rand):
    encoded = Bytes()
    encoded += Prio3.Flp.Field.encode_vec(verifier)
    if Prio3.Flp.JOINT_RAND_LEN > 0:
        encoded += k_joint_rand
    return encoded

def decode_prepare_message(Prio3, encoded):
    l = Prio3.Flp.Field.ENCODED_SIZE * Prio3.Flp.VERIFIER_LEN
    encoded_verifier, encoded = encoded[:l], encoded[l:]
    verifier = Prio3.Flp.Field.decode_vec(encoded_verifier)
    k_joint_rand = None
    if Prio3.Flp.JOINT_RAND_LEN > 0:
        l = Prio3.Prg.SEED_SIZE
        k_joint_rand, encoded = encoded[:l], encoded[l:]
    if len(encoded) != 0:
        raise ERR_DECODE
    return (verifier, k_joint_rand)

```

Figure 17: Helper functions required for Prio3.

6.3. A General-Purpose FLP

This section describes an FLP based on the construction from in [BBCGGI19], Section 4.2. We begin in [Section 6.3.1](#) with an overview of their proof system and the extensions to their proof system made here. The construction is specified in [Section 6.3.3](#).

OPEN ISSUE We're not yet sure if specifying this general-purpose FLP is desirable. It might be preferable to specify specialized FLPs for each data type that we want to standardize, for two reasons. First, clear and concise specifications are likely easier to write for specialized FLPs rather than the general one. Second, we may end up tailoring each FLP to the measurement type in a way that improves performance, but breaks compatibility with the general-purpose FLP.

In any case, we can't make this decision until we know which data types to standardize, so for now, we'll stick with the general-purpose construction. The reference implementation can be found at <https://github.com/cjpatton/vdaf/tree/main/poc>.

OPEN ISSUE Chris Wood points out that the this section reads more like a paper than a standard. Eventually we'll want to work this into something that is readily consumable by the CFRG.

6.3.1. Overview

In the proof system of [BBCGGI19], validity is defined via an arithmetic circuit evaluated over the input: If the circuit output is zero, then the input is deemed valid; otherwise, if the circuit output is non-zero, then the input is deemed invalid. Thus the goal of the proof system is merely to allow the verifier to evaluate the validity circuit over the input. For our application ([Section 6](#)), this computation is distributed among multiple Aggregators, each of which has only a share of the input.

Suppose for a moment that the validity circuit C is affine, meaning its only operations are addition and multiplication-by-constant. In particular, suppose the circuit does not contain a multiplication gate whose operands are both non-constant. Then to decide if an input x is valid, each Aggregator could evaluate C on its share of x locally, broadcast the output share to its peers, then combine the output shares locally to recover $C(x)$. This is true because for any SHARES-way secret sharing of x it holds that

$$\begin{aligned} C(x_shares[0] + \dots + x_shares[SHARES-1]) = \\ C(x_shares[0]) + \dots + C(x_shares[SHARES-1]) \end{aligned}$$

(Note that, for this equality to hold, it may be necessary to scale any constants in the circuit by SHARES.) However this is not the case if C is not-affine (i.e., it contains at least one multiplication gate whose operands are non-constant). In the proof system of [BBCGGI19], the proof is designed to allow the (distributed) verifier to compute the non-affine operations using only linear operations on (its share of) the input and proof.

To make this work, the proof system is restricted to validity circuits that exhibit a special structure. Specifically, an arithmetic circuit with "G-gates" (see [BBCGGI19], Definition 5.2) is composed of affine gates and any number of instances of a distinguished gate G , which may be non-affine. We will refer to this class of circuits as "gadget circuits" and to G as the "gadget".

As an illustrative example, consider a validity circuit C that recognizes the set $L = \text{set}([0], [1])$. That is, C takes as input a length-1 vector x and returns 0 if $x[0]$ is in $[0, 2)$ and outputs something else otherwise. This circuit can be expressed as the following degree-2 polynomial:

$$C(x) = (x[0] - 1) * x[0] = x[0]^2 - x[0]$$

This polynomial recognizes L because $x[0]^2 = x[0]$ is only true if $x[0] == 0$ or $x[0] == 1$. Notice that the polynomial involves a non-affine operation, $x[0]^2$. In order to apply [BBCGGI19], Theorem 4.3, the circuit needs to be rewritten in terms of a gadget that subsumes this non-affine operation. For example, the gadget might be multiplication:

$$\text{Mul}(\text{left}, \text{right}) = \text{left} * \text{right}$$

The validity circuit can then be rewritten in terms of Mul like so:

$$C(x[0]) = \text{Mul}(x[0], x[0]) - x[0]$$

The proof system of [BBCGGI19] allows the verifier to evaluate each instance of the gadget (i.e., $\text{Mul}(x[0], x[0])$ in our example) using a linear function of the input and proof. The proof is constructed roughly as follows. Let C be the validity circuit and suppose the gadget is arity- L (i.e., it has L input wires.). Let $\text{wire}[j-1, k-1]$ denote the value of the j th wire of the k th call to the gadget during the evaluation of $C(x)$. Suppose there are M such calls and fix distinct field elements $\alpha[0], \dots, \alpha[M-1]$. (We will require these points to have a special property, as we'll discuss in [Section 6.3.1.1](#); but for the moment it is only important that they are distinct.)

The prover constructs from wire and α a polynomial that, when evaluated at $\alpha[k-1]$, produces the output of the k th call to the

gadget. Let us call this the "gadget polynomial". Polynomial evaluation is linear, which means that, in the distributed setting, the Client can disseminate additive shares of the gadget polynomial that the Aggregators then use to compute additive shares of each gadget output, allowing each Aggregator to compute its share of $C(x)$ locally.

There is one more wrinkle, however: It is still possible for a malicious prover to produce a gadget polynomial that would result in $C(x)$ being computed incorrectly, potentially resulting in an invalid input being accepted. To prevent this, the verifier performs a probabilistic test to check that the gadget polynomial is well-formed. This test, and the procedure for constructing the gadget polynomial, are described in detail in [Section 6.3.3](#).

6.3.1.1. Extensions

The FLP described in the next section extends the proof system [\[BBCGGI19\]](#), Section 4.2 in three ways.

First, the validity circuit in our construction includes an additional, random input (this is the "joint randomness" derived from the input shares in Prio3; see [Section 6.2](#)). This allows for circuit optimizations that trade a small soundness error for a shorter proof. For example, consider a circuit that recognizes the set of length- N vectors for which each element is either one or zero. A deterministic circuit could be constructed for this language, but it would involve a large number of multiplications that would result in a large proof. (See the discussion in [\[BBCGGI19\]](#), Section 5.2 for details). A much shorter proof can be constructed for the following randomized circuit:

$$C(\text{inp}, r) = r * \text{Range2}(\text{inp}[0]) + \dots + r^N * \text{Range2}(\text{inp}[N-1])$$

(Note that this is a special case of [\[BBCGGI19\]](#), Theorem 5.2.) Here inp is the length- N input and r is a random field element. The gadget circuit Range2 is the "range-check" polynomial described above, i.e., $\text{Range2}(x) = x^2 - x$. The idea is that, if inp is valid (i.e., each $\text{inp}[j]$ is in $[0, 2)$), then the circuit will evaluate to 0 regardless of the value of r ; but if $\text{inp}[j]$ is not in $[0, 2)$ for some j , the output will be non-zero with high probability.

The second extension implemented by our FLP allows the validity circuit to contain multiple gadget types. (This generalization was suggested in [\[BBCGGI19\]](#), Remark 4.5.) For example, the following circuit is allowed, where Mul and Range2 are the gadgets defined above (the input has length $N+1$):

$$C(\text{inp}, r) = r * \text{Range2}(\text{inp}[0]) + \dots + r^N * \text{Range2}(\text{inp}[N-1]) + \backslash \\ 2^{\wedge}0 * \text{inp}[0] + \dots + 2^{\wedge}(N-1) * \text{inp}[N-1] - \backslash \\ \text{Mul}(\text{inp}[N], \text{inp}[N])$$

Finally, [BBCGGI19], Theorem 4.3 makes no restrictions on the choice of the fixed points $\alpha[0], \dots, \alpha[M-1]$, other than to require that the points are distinct. In this document, the fixed points are chosen so that the gadget polynomial can be constructed efficiently using the Cooley-Tukey FFT ("Fast Fourier Transform") algorithm. Note that this requires the field to be "FFT-friendly" as defined in [Section 5.1.2](#).

6.3.2. Validity Circuits

The FLP described in [Section 6.3.3](#) is defined in terms of a validity circuit `Valid` that implements the interface described here.

A concrete `Valid` defines the following parameters:

Parameter	Description
GADGETS	A list of gadgets
GADGET_CALLS	Number of times each gadget is called
INPUT_LEN	Length of the input
OUTPUT_LEN	Length of the aggregatable output
JOINT_RAND_LEN	Length of the random input
Measurement	The type of measurement
Field	An FFT-friendly finite field as defined in Section 5.1.2

Table 6: Validity circuit parameters.

Each gadget G in `GADGETS` defines a constant `DEGREE` that specifies the circuit's "arithmetic degree". This is defined to be the degree of the polynomial that computes it. For example, the `Mul` circuit in [Section 6.3.1](#) is defined by the polynomial $\text{Mul}(x) = x * x$, which has degree 2. Hence, the arithmetic degree of this gadget is 2.

Each gadget also defines a parameter `ARITY` that specifies the circuit's arity (i.e., the number of input wires).

A concrete `Valid` provides the following methods for encoding a measurement as an input vector and truncating an input vector to the length of an aggregatable output:

`*Valid.encode(measurement: Measurement) -> Vec[Field]` returns a vector of length `INPUT_LEN` representing a measurement.

`*Valid.truncate(input: Vec[Field]) -> Vec[Field]` returns a vector of length `OUTPUT_LEN` representing an aggregatable output.

Finally, the following class methods are derived for each concrete Valid:

```
# Length of the prover randomness.
def prove_rand_len(Valid):
    return sum(map(lambda g: g.ARITY, Valid.GADGETS))

# Length of the query randomness.
def query_rand_len(Valid):
    return len(Valid.GADGETS)

# Length of the proof.
def proof_len(Valid):
    length = 0
    for (g, g_calls) in zip(Valid.GADGETS, Valid.GADGET_CALLS):
        P = next_power_of_2(1 + g_calls)
        length += g.ARITY + g.DEGREE * (P - 1) + 1
    return length

# Length of the verifier message.
def verifier_len(Valid):
    length = 1
    for g in Valid.GADGETS:
        length += g.ARITY + 1
    return length
```

Figure 18: Derived methods for validity circuits.

6.3.3. Construction

This section specifies FlpGeneric, an implementation of the Flp interface ([Section 6.1](#)). It has as a generic parameter a validity circuit Valid implementing the interface defined in [Section 6.3.2](#).

NOTE A reference implementation can be found in https://github.com/cjpatton/vdaf/blob/main/poc/flp_generic.sage.

The FLP parameters for FlpGeneric are defined in [Table 7](#). The required methods for generating the proof, generating the verifier, and deciding validity are specified in the remaining subsections.

In the remainder, we let $[n]$ denote the set $\{1, \dots, n\}$ for positive integer n . We also define the following constants:

```
*Let  $H = \text{len}(\text{Valid.GADGETS})$ 

*For each  $i$  in  $[H]$ :

    -Let  $G_i = \text{Valid.GADGETS}[i]$ 
```

```

-Let L_i = Valid.GADGETS[i].ARITY

-Let M_i = Valid.GADGET_CALLS[i]

-Let P_i = next_power_of_2(M_i+1)

-Let alpha_i = Field.gen()^(Field.GEN_ORDER / P_i)

```

Parameter	Value
PROVE_RAND_LEN	Valid.prove_rand_len() (see Section 6.3.2)
QUERY_RAND_LEN	Valid.query_rand_len() (see Section 6.3.2)
JOINT_RAND_LEN	Valid.JOINT_RAND_LEN
INPUT_LEN	Valid.INPUT_LEN
OUTPUT_LEN	Valid.OUTPUT_LEN
PROOF_LEN	Valid.proof_len() (see Section 6.3.2)
VERIFIER_LEN	Valid.verifier_len() (see Section 6.3.2)
Measurement	Valid.Measurement
Field	Valid.Field

Table 7: FLP Parameters of FlpGeneric.

6.3.3.1. Proof Generation

On input `inp`, `prove_rand`, and `joint_rand`, the proof is computed as follows:

1. For each i in $[H]$ create an empty table `wire_i`.
2. Partition the prover randomness `prove_rand` into subvectors `seed_1, ..., seed_H` where $\text{len}(\text{seed}_i) == L_i$ for all i in $[H]$. Let us call these the "wire seeds" of each gadget.
3. Evaluate Valid on input of `inp` and `joint_rand`, recording the inputs of each gadget in the corresponding table. Specifically, for every i in $[H]$, set `wire_i[j-1,k-1]` to the value on the j th wire into the k th call to gadget G_i .
4. Compute the "wire polynomials". That is, for every i in $[H]$ and j in $[L_i]$, construct `poly_wire_i[j-1]`, the j th wire polynomial for the i th gadget, as follows:

```

*Let w = [seed_i[j-1], wire_i[j-1,0], ...,
          wire_i[j-1,M_i-1]].

```

```

*Let padded_w = w + Field.zeros(P_i - len(w)).

```

NOTE We pad w to the nearest power of 2 so that we can use FFT for interpolating the wire polynomials. Perhaps there is

some clever math for picking wire_inp in a way that avoids having to pad.

*Let $\text{poly_wire_i}[j-1]$ be the lowest degree polynomial for which $\text{poly_wire_i}[j-1](\alpha_i^k) == \text{padded_w}[k]$ for all k in $[P_i]$.

5. Compute the "gadget polynomials". That is, for every i in $[H]$:

*Let $\text{poly_gadget_i} = G_i(\text{poly_wire_i}[0], \dots, \text{poly_wire_i}[L_i-1])$. That is, evaluate the circuit G_i on the wire polynomials for the i th gadget. (Arithmetic is in the ring of polynomials over Field .)

The proof is the vector $\text{proof} = \text{seed}_1 + \text{coeff}_1 + \dots + \text{seed}_H + \text{coeff}_H$, where coeff_i is the vector of coefficients of poly_gadget_i for each i in $[H]$.

6.3.3.2. Query Generation

On input of inp , proof , query_rand , and joint_rand , the verifier message is generated as follows:

1. For every i in $[H]$ create an empty table wire_i .
2. Partition proof into the subvectors $\text{seed}_1, \text{coeff}_1, \dots, \text{seed}_H, \text{coeff}_H$ defined in [Section 6.3.3.1](#).
3. Evaluate Valid on input of inp and joint_rand , recording the inputs of each gadget in the corresponding table. This step is similar to the prover's step (3.) except the verifier does not evaluate the gadgets. Instead, it computes the output of the k th call to G_i by evaluating $\text{poly_gadget_i}(\alpha_i^k)$. Let v denote the output of the circuit evaluation.
4. Compute the wire polynomials just as in the prover's step (4.).
5. Compute the tests for well-formedness of the gadget polynomials. That is, for every i in $[H]$:

*Let $t = \text{query_rand}[i]$. Check if $t^{(P_i)} == 1$: If so, then raise ERR_ABORT and halt. (This prevents the verifier from inadvertently leaking a gadget output in the verifier message.)

*Let $y_i = \text{poly_gadget_i}(t)$.

*For each j in $[0, L_i)$ let $x_i[j-1] = \text{poly_wire_i}[j-1](t)$.

The verifier message is the vector $\text{verifier} = [v] + x_1 + [y_1] + \dots + x_H + [y_H]$.

6.3.3.3. Decision

On input of vector verifier , the verifier decides if the input is valid as follows:

1. Parse verifier into $v, x_1, y_1, \dots, x_H, y_H$ as defined in [Section 6.3.3.2](#).
2. Check for well-formedness of the gadget polynomials. For every i in $[H]$:
 - *Let $z = G_i(x_i)$. That is, evaluate the circuit G_i on x_i and set z to the output.
 - *If $z \neq y_i$, then return False and halt.
3. Return True if $v == 0$ and False otherwise.

6.3.3.4. Encoding

The FLP encoding and truncation methods invoke `Valid.encode` and `Valid.truncate` in the natural way.

6.4. Instantiations

This section specifies instantiations of Prio3 for various measurement types. Each uses `FlpGeneric` as the FLP ([Section 6.3](#)) and is determined by a validity circuit ([Section 6.3.2](#)) and a PRG ([Section 5.2](#)). Test vectors for each can be found in [Appendix "Test Vectors"](#).

NOTE Reference implementations of each of these VDAFs can be found in https://github.com/cjpatton/vdaf/blob/main/poc/vdaf_prio3.sage.

6.4.1. Prio3Aes128Count

Our first instance of Prio3 is for a simple counter: Each measurement is either one or zero and the aggregate result is the sum of the measurements.

This instance uses `PrgAes128` ([Section 5.2.1](#)) as its PRG. Its validity circuit, denoted `Count`, uses `Field64` ([Table 2](#)) as its finite field. Its gadget, denoted `Mul`, is the degree-2, arity-2 gadget defined as

```
def Mul(x, y):
    return x * y
```

The validity circuit is defined as

```
def Count(inp: Vec[Field64]):
    return Mul(inp[0], inp[0]) - inp[0]
```

The measurement is encoded as a singleton vector in the natural way. The parameters for this circuit are summarized below.

Parameter	Value
GADGETS	[Mul]
GADGET_CALLS	[1]
INPUT_LEN	1
OUTPUT_LEN	1
JOINT_RAND_LEN	0
Measurement	Unsigned, in range $[0, 2)$
Field	Field64 (Table 2)

Table 8: Parameters of validity circuit Count.

6.4.2. Prio3Aes128Sum

The next instance of Prio3 supports summing of integers in a pre-determined range. Each measurement is an integer in range $[0, 2^{\text{bits}})$, where bits is an associated parameter.

This instance of Prio3 uses PrgAes128 ([Section 5.2.1](#)) as its PRG. Its validity circuit, denoted Sum, uses Field128 ([Table 3](#)) as its finite field. The measurement is encoded as a length-bits vector of field elements, where the l th element of the vector represents the l th bit of the summand:

```

def encode(Sum, measurement: Integer):
    if 0 > measurement or measurement >= 2^Sum.INPUT_LEN:
        raise ERR_INPUT

    encoded = []
    for l in range(Sum.INPUT_LEN):
        encoded.append(Sum.Field((measurement >> l) & 1))
    return encoded

def truncate(Sum, inp):
    decoded = Sum.Field(0)
    for (l, b) in enumerate(inp):
        w = Sum.Field(1 << l)
        decoded += w * b
    return [decoded]

```

The validity circuit checks that the input comprised of ones and zeros. Its gadget, denoted `Range2`, is the degree-2, arity-1 gadget defined as

```

def Range2(x):
    return x^2 - x

```

The validity circuit is defined as

```

def Sum(inp: Vec[Field128], joint_rand: Vec[Field128]):
    out = Field128(0)
    r = joint_rand[0]
    for x in inp:
        out += r * Range2(x)
        r *= joint_rand[0]
    return out

```

Parameter	Value
GADGETS	[Range2]
GADGET_CALLS	[bits]
INPUT_LEN	bits
OUTPUT_LEN	1
JOINT_RAND_LEN	1
Measurement	Unsigned, in range $[0, 2^{\text{bits}})$
Field	Field128 (Table 3)

Table 9: Parameters of validity circuit Sum.

6.4.3. Prio3Aes128Histogram

This instance of Prio3 allows for estimating the distribution of the measurements by computing a simple histogram. Each measurement is an arbitrary integer and the aggregate result counts the number of measurements that fall in a set of fixed buckets.

This instance of Prio3 uses PrgAes128 ([Section 5.2.1](#)) as its PRG. Its validity circuit, denoted Histogram, uses Field128 ([Table 3](#)) as its finite field. The measurement is encoded as a one-hot vector representing the bucket into which the measurement falls (let bucket denote a sequence of monotonically increasing integers):

```
def encode(Histogram, measurement: Integer):
    boundaries = buckets + [Infinity]
    encoded = [Field128(0) for _ in range(len(boundaries))]
    for i in range(len(boundaries)):
        if measurement <= boundaries[i]:
            encoded[i] = Field128(1)
    return encoded

def truncate(Histogram, inp: Vec[Field128]):
    return inp
```

The validity circuit uses Range2 (see [Section 6.4.2](#)) as its single gadget. It checks for one-hotness in two steps, as follows:

```
def Histogram(inp: Vec[Field128],
              joint_rand: Vec[Field128],
              num_shares: Unsigned):
    # Check that each bucket is one or zero.
    range_check = Field128(0)
    r = joint_rand[0]
    for x in inp:
        range_check += r * Range2(x)
        r *= joint_rand[0]

    # Check that the buckets sum to 1.
    sum_check = -Field128(1) * Field128(num_shares).inv()
    for b in inp:
        sum_check += b

    out = joint_rand[1] * range_check + \
          joint_rand[1]^2 * sum_check
    return out
```

Note that this circuit depends on the number of shares into which the input is sharded. This is provided to the FLP by Prio3.

Parameter	Value
GADGETS	[Range2]
GADGET_CALLS	[buckets + 1]
INPUT_LEN	buckets + 1
OUTPUT_LEN	buckets + 1
JOINT_RAND_LEN	2
Measurement	Integer
Field	Field128 (Table 3)

Table 10: Parameters of validity circuit Histogram.

7. Poplar1

NOTE The spec for Poplar1 is still a work-in-progress. A partial implementation can be found at <https://github.com/abetterinternet/libprio-rs/blob/main/src/vdaf/poplar1.rs>. The verification logic is nearly complete, however as of this draft the code is missing the IDPF. An implementation of the IDPF can be found at https://github.com/google/distributed_point_functions/.

This section specifies Poplar1, a VDAF for the following task. Each Client holds a BITS-bit string and the Aggregators hold a set of 1-bit strings, where $1 \leq \text{BITS}$. We will refer to the latter as the set of "candidate prefixes". The Aggregators' goal is to count how many inputs are prefixed by each candidate prefix.

This functionality is the core component of Poplar [[BBCGGI21](#)]. At a high level, the protocol works as follows.

1. Each Clients runs the input-distribution algorithm on its n-bit string and sends an input share to each Aggregator.
2. The Aggregators agree on an initial set of candidate prefixes, say 0 and 1.
3. The Aggregators evaluate the VDAF on each set of input shares and aggregate the recovered output shares. The aggregation parameter is the set of candidate prefixes.
4. The Aggregators send their aggregate shares to the Collector, who combines them to recover the counts of each candidate prefix.
5. Let H denote the set of prefixes that occurred at least t times. If the prefixes all have length BITS, then H is the set of t -heavy-hitters. Otherwise compute the next set of candidate prefixes as follows. For each p in H , add $p \parallel 0$ and $p \parallel 1$

to the set. Repeat step 3 with the new set of candidate prefixes.

Poplar1 is constructed from an "Incremental Distributed Point Function (IDPF)", a primitive described by [BBCGGI21] that generalizes the notion of a Distributed Point Function (DPF) [GI14]. Briefly, a DPF is used to distribute the computation of a "point function", a function that evaluates to zero on every input except at a programmable "point". The computation is distributed in such a way that no one party knows either the point or what it evaluates to.

An IDPF generalizes this "point" to a path on a full binary tree from the root to one of the leaves. It is evaluated on an "index" representing a unique node of the tree. If the node is on the path, then function evaluates to a non-zero value; otherwise it evaluates to zero. This structure allows an IDPF to provide the functionality required for the above protocol, while at the same time ensuring the same degree of privacy as a DPF.

Our VDAF composes an IDPF with the "secure sketching" protocol of [BBCGGI21]. This protocol ensures that evaluating a set of input shares on a unique set of candidate prefixes results in shares of a "one-hot" vector, i.e., a vector that is zero everywhere except for one element, which is equal to one.

7.1. Incremental Distributed Point Functions (IDPFs)

An IDPF is defined over a domain of size 2^{BITS} , where BITS is constant defined by the IDPF. The Client specifies an index α and a pair of values β , one for each "level" $1 \leq l \leq \text{BITS}$. The key generation generates two IDPF keys, one for each Aggregator. When evaluated at index $0 \leq x < 2^{\text{BITS}}$, each IDPF share returns an additive share of $\beta[l]$ if x is the l -bit prefix of α and shares of zero otherwise.

CP What does it mean for x to be the l -bit prefix of α ? We need to be a bit more precise here.

CP Why isn't the domain size actually $2^{(\text{BITS}+1)}$, i.e., the number of nodes in a binary tree of height BITS (excluding the root)?

Each $\beta[l]$ is a pair of elements of a finite field. Each level **MAY** have different field parameters. Thus a concrete IDPF specifies associated types $\text{Field}[1]$, $\text{Field}[2]$, ..., and $\text{Field}[\text{BITS}]$ defining, respectively, the field parameters at level 1, level 2, ..., and level BITS.

An IDPF is comprised of the following algorithms (let type `Value[l]` denote `(Field[l], Field[l])` for each level `l`):

`*idpf_gen(alpha: Unsigned, beta: (Value[1], ..., Value[BITS])) -> key: (IDPFKey, IDPFKey)` is the randomized key-generation algorithm run by the client. Its inputs are the index `alpha` and the values `beta`. The value of `alpha` **MUST** be in range `[0, 2^BITS)`.

`*IDPFKey.eval(l: Unsigned, x: Unsigned) -> value: Value[l]` is deterministic, stateless key-evaluation algorithm run by each Aggregator. It returns the value corresponding to index `x`. The value of `l` **MUST** be in `[1, BITS]` and the value of `x` **MUST** be in range `[2^(l-1), 2^l)`.

A concrete IDPF specifies a single associated constant:

`*BITS: Unsigned` is the length of each Client input.

A concrete IDPF also specifies the following associated types:

`*Field[l]` for each level `1 <= l <= BITS`. Each defines the same methods and associated constants as `Field` in [Section 6](#).

Note that IDPF construction of [BBCGGI21](#) uses one field for the inner nodes of the tree and a different, larger field for the leaf nodes. See [BBCGGI21](#), Section 4.3.

Finally, an implementation note. The interface for IDPFs specified here is stateless, in the sense that there is no state carried between IDPF evaluations. This is to align the IDPF syntax with the VDAF abstraction boundary, which does not include shared state across VDAF evaluations. In practice, of course, it will often be beneficial to expose a stateful API for IDPFs and carry the state across evaluations.

7.2. Construction

The VDAF involves two rounds of communication (`ROUNDS == 2`) and is defined for two Aggregators (`SHARES == 2`).

7.2.1. Setup

The verification parameter is a symmetric key shared by both Aggregators. This VDAF has no public parameter.

```
def vdaf_setup():
    k_verify_init = gen_rand(SEED_SIZE)
    return (None, [(0, k_verify_init), (1, k_verify_init)])
```

Figure 19: The setup algorithm for poplar1.

7.2.1.1. Client

The client's input is an IDPF index, denoted α . The values are pairs of field elements $(1, k)$ where each k is chosen at random. This random value is used as part of the secure sketching protocol of [BBCGGI21]. After evaluating their IDPF key shares on the set of candidate prefixes, the sketching protocol is used by the Aggregators to verify that they hold shares of a one-hot vector. In addition, for each level of the tree, the prover generates random elements a , b , and c and computes

$$A = -2*a + k$$

$$B = a*a + b - k*a + c$$

and sends additive shares of a , b , c , A and B to the Aggregators. Putting everything together, the input-distribution algorithm is defined as follows. Function `encode_input_share` is defined in [Section 7.2.5](#).

```

def measurement_to_input_shares(_, alpha):
    if alpha < 2**BITS: raise ERR_INVALID_INPUT

    # Prepare IDPF values.
    beta = []
    correlation_shares_0, correlation_shares_1 = [], []
    for l in range(1,BITS+1):
        (k, a, b, c) = Field[l].rand_vec(4)

        # Construct values of the form (1, k), where k
        # is a random field element.
        beta += [(1, k)]

        # Create secret shares of correlations to aid
        # the Aggregators' computation.
        A = -2*a+k
        B = a*a + b - a * k + c
        correlation_share = Field[l].rand_vec(5)
        correlation_shares_1.append(correlation_share)
        correlation_shares_0.append(
            [a, b, c, A, B] - correlation_share)

    # Generate IDPF shares.
    (key_0, key_1) = idpf_gen(alpha, beta)

    input_shares = [
        encode_input_share(key_0, correlation_shares_0),
        encode_input_share(key_1, correlation_shares_1),
    ]

    return input_shares

```

Figure 20: The input-distribution algorithm for poplar1.

TODO It would be more efficient to represent the shares of a , b , and c using PRG seeds as suggested in [\[BBCGGI21\]](#).

7.2.2. Preparation

The aggregation parameter encodes a sequence of candidate prefixes. When an Aggregator receives an input share from the Client, it begins by evaluating its IDPF share on each candidate prefix, recovering a pair of vectors of field elements `data_share` and `auth_share`. The Aggregators use `auth_share` and the correlation shares provided by the Client to verify that their `data_share` vectors are additive shares of a one-hot vector.

CP Consider adding aggregation parameter as input to `k_verify_rand` derivation.

```

class PrepState:
    def __init__(verify_param, agg_param, nonce, input_share):
        (self.l, self.candidate_prefixes) = decode_indexes(agg_param)
        (self.idpf_key,
         self.correlation_shares) = decode_input_share(input_share)
        (self.party_id, k_verify_init) = verify_param
        self.k_verify_rand = get_key(k_verify_init, nonce)
        self.step = "ready"

    def next(self, inbound: Optional[Bytes]):
        l = self.l
        (a_share, b_share, c_share,
         A_share, B_share) = correlation_shares[l-1]

        if self.step == "ready" and inbound == None:
            # Evaluate IDPF on candidate prefixes.
            data_share, auth_share = [], []
            for x in self.candidate_prefixes:
                value = self.idpf_key.eval(l, x)
                data_share.append(value[0])
                auth_share.append(value[1])

            # Prepare first sketch verification message.
            r = Prg.expand_into_vec(Field[l], self.k_verify_rand, len(data_share))
            verifier_share_1 = [
                a_share + inner_product(data_share, r),
                b_share + inner_product(data_share, r * r),
                c_share + inner_product(auth_share, r),
            ]

            self.output_share = data_share
            self.step = "sketch round 1"
            return verifier_share_1

        elif self.step == "sketch round 1" and inbound != None:
            verifier_1 = Field[l].decode_vec(inbound)
            verifier_share_2 = [
                (verifier_1[0] * verifier_1[0] \
                 - verifier_1[1] \
                 - verifier_1[2]) * self.party_id \
                + A_share * verifier_1[0] \
                + B_share
            ]

            self.step = "sketch round 2"
            return Field[l].encode_vec(verifier_share_2)

        elif self.step == "sketch round 2" and inbound != None:
            verifier_2 = Field[l].decode_vec(inbound)

```

```

        if verifier_2 != 0: raise ERR_INVALID
        return Field[1].encode_vec(self.output_share)

    else: raise ERR_INVALID_STATE

def prep_shares_to_prep(agg_param, inbound: Vec[Bytes]):
    if len(inbound) != 2:
        raise ERR_INVALID_INPUT

    (l, _) = decode_indexes(agg_param)
    verifier = Field[1].decode_vec(inbound[0]) + \
        Field[1].decode_vec(inbound[1])

    return Field[1].encode_vec(verifier)

```

Figure 21: Preparation state for poplar1.

7.2.3. Aggregation

```
def out_shares_to_agg_share(agg_param, output_shares: Vec[Bytes]):
    (l, candidate_prefixes) = decode_indexes(agg_param)
    if len(output_shares) != len(candidate_prefixes):
        raise ERR_INVALID_INPUT

    agg_share = Field[l].zeros(len(candidate_prefixes))
    for output_share in output_shares:
        agg_share += Field[l].decode_vec(output_share)

    return Field[l].encode_vec(agg_share)
```

Figure 22: Aggregation algorithm for poplar1.

7.2.4. Unsharding

```
def agg_shares_to_result(agg_param, agg_shares: Vec[Bytes]):
    (l, _) = decode_indexes(agg_param)
    if len(agg_shares) != 2:
        raise ERR_INVALID_INPUT

    agg = Field[l].decode_vec(agg_shares[0]) + \
        Field[l].decode_vec(agg_shares[1])

    return Field[l].encode_vec(agg)
```

Figure 23: Computation of the aggregate result for poplar1.

7.2.5. Helper Functions

TODO Specify the following functionalities:

*encode_input_share is used to encode an input share, consisting of an IDPF key share and correlation shares.

*decode_input_share is used to decode an input share.

*decode_indexes(encoded: Bytes) -> (l: Unsigned, indexes: Vec[Unsigned]) decodes a sequence of indexes, i.e., candidate indexes for IDPF evaluation. The value of l **MUST** be in range [1, BITS] and indexes[i] **MUST** be in range $[2^{(l-1)}, 2^l]$ for all i. An error is raised if encoded cannot be decoded.

8. Security Considerations

NOTE: This is a brief outline of the security considerations. This section will be filled out more as the draft matures and security analyses are completed.

VDAFs have two essential security goals:

1. Privacy: An attacker that controls the network, the Collector, and a subset of Clients and Aggregators learns nothing about the measurements of honest Clients beyond what it can deduce from the aggregate result.
2. Robustness: An attacker that controls the network and a subset of Clients cannot cause the Collector to compute anything other than the aggregate of the measurements of honest Clients.

Note that, to achieve robustness, it is important to ensure that the verification parameters distributed to the Aggregators (verify_params, see [Section 6.2.1](#)) is never revealed to the Clients.

It is also possible to consider a stronger form of robustness, where the attacker also controls a subset of Aggregators (see [[BBCGGI19](#)], Section 6.3). To satisfy this stronger notion of robustness, it is necessary to prevent the attacker from sharing the verification parameters with the Clients. It is therefore **RECOMMENDED** that the Aggregators generate verify_params only after a set of Client inputs has been collected for verification, and re-generate them for each such set of inputs.

In order to achieve robustness, the Aggregators **MUST** ensure that the nonces used to process the measurements in a batch are all unique.

A VDAF is the core cryptographic primitive of a protocol that achieves the above privacy and robustness goals. It is not sufficient on its own, however. The application will need to assure a few security properties, for example:

- *Securely distributing the long-lived parameters.

- *Establishing secure channels:

 - Confidential and authentic channels among Aggregators, and between the Aggregators and the Collector; and

 - Confidential and Aggregator-authenticated channels between Clients and Aggregators.

- *Enforcing the non-collusion properties required of the specific VDAF in use.

In such an environment, a VDAF provides the high-level privacy property described above: The Collector learns only the aggregate measurement, and nothing about individual measurements aside from what can be inferred from the aggregate result. The Aggregators learn neither individual measurements nor the aggregate result. The Collector is assured that the aggregate statistic accurately reflects the inputs as long as the Aggregators correctly executed their role in the VDAF.

On their own, VDAFs do not mitigate Sybil attacks [Dou02]. In this attack, the adversary observes a subset of input shares transmitted by a Client it is interested in. It allows the input shares to be processed, but corrupts and picks bogus inputs for the remaining Clients. Applications can guard against these risks by adding additional controls on measurement submission, such as client authentication and rate limits.

VDAFs do not inherently provide differential privacy [Dwo06]. The VDAF approach to private measurement can be viewed as complementary to differential privacy, relying on non-collusion instead of statistical noise to protect the privacy of the inputs. It is possible that a future VDAF could incorporate differential privacy features, e.g., by injecting noise before the sharding stage and removing it after unsharding.

9. IANA Considerations

This document makes no request of IANA.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC4493] Song, JH., Poovendran, R., Lee, J., and T. Iwata, "The AES-CMAC Algorithm", RFC 4493, DOI 10.17487/RFC4493, June 2006, <<https://www.rfc-editor.org/rfc/rfc4493>>.
- [RFC8017] Moriarty, K., Ed., Kaliski, B., Jonsson, J., and A. Rusch, "PKCS #1: RSA Cryptography Specifications Version 2.2", RFC 8017, DOI 10.17487/RFC8017, November 2016, <<https://www.rfc-editor.org/rfc/rfc8017>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

10.2. Informative References

- [AGJOP21] Addanki, S., Garbe, K., Jaffe, E., Ostrovsky, R., and A. Polychroniadou, "Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares", 2021, <<https://ia.cr/2021/576>>.
- [BBCGGI19] Boneh, D., Boyle, E., Corrigan-Gibbs, H., Gilboa, N., and Y. Ishai, "Zero-Knowledge Proofs on Secret-Shared Data via Fully Linear PCPs", CRYPTO 2019 , 2019, <<https://ia.cr/2019/188>>.
- [BBCGGI21] Boneh, D., Boyle, E., Corrigan-Gibbs, H., Gilboa, N., and Y. Ishai, "Lightweight Techniques for Private Heavy Hitters", IEEE S&P 2021 , 2021, <<https://ia.cr/2021/017>>.
- [CGB17] Corrigan-Gibbs, H. and D. Boneh, "Prio: Private, Robust, and Scalable Computation of Aggregate Statistics", NSDI 2017 , 2017, <<https://dl.acm.org/doi/10.5555/3154630.3154652>>.
- [Dou02] Douceur, J., "The Sybil Attack", IPTPS 2002 , 2002, <https://doi.org/10.1007/3-540-45748-8_24>.
- [Dwo06] Dwork, C., "Differential Privacy", ICALP 2006 , 2006, <https://link.springer.com/chapter/10.1007/11787006_1>.
- [ENPA] "Exposure Notification Privacy-preserving Analytics (ENPA) White Paper", 2021, <https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf>.
- [EPK14] Erlingsson, Ú., Pihur, V., and A. Korolova, "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response", CCS 2014 , 2014, <<https://dl.acm.org/doi/10.1145/2660267.2660348>>.
- [GI14] Gilboa, N. and Y. Ishai, "Distributed Point Functions and Their Applications", EUROCRYPT 2014 , 2014, <https://link.springer.com/chapter/10.1007/978-3-642-55220-5_35>.
- [I-D.draft-gpew-priv-ppm] Geoghegan, T., Patton, C., Rescorla, E., and C. A. Wood, "Privacy Preserving Measurement", Work in Progress, Internet-Draft, draft-gpew-priv-ppm-01, 7 March 2022, <<https://datatracker.ietf.org/doc/html/draft-gpew-priv-ppm-01>>.
- [OriginTelemetry] "Origin Telemetry", 2020, <<https://firefox-source-docs.mozilla.org/toolkit/components/telemetry/collection/origin.html>>.

[Vad16]

Vadhan, S., "The Complexity of Differential Privacy", 2016, <https://link.springer.com/chapter/10.1007/978-3-319-57048-8_7>.

Acknowledgments

Thanks to David Cook, Henry Corrigan-Gibbs, Armando Faz-Hernandez, Mariana Raykova, and Christopher Wood for useful feedback on and contributions to the spec.

Test Vectors

Test vectors cover the generation of input shares and the conversion of input shares into output shares. Vectors specify the public and verification parameters, the measurement, the aggregation parameter, the expected input shares, the prepare messages, and the expected output shares.

Test vectors are encoded in JSON. Input shares and prepare messages are represented as hexadecimal streams. To make the tests deterministic, `gen_rand()` was replaced with a function that returns the requested number of 0x01 octets.

Prio3Aes128Count

For this test, the value of SHARES is 2.

```

{
  "public_param": null,
  "verify_params": [
    [
      0,
      "01010101010101010101010101010101"
    ],
    [
      1,
      "01010101010101010101010101010101"
    ]
  ],
  "agg_param": null,
  "prep": [
    {
      "measurement": 1,
      "nonce": "01010101010101010101010101010101",
      "input_shares": [
        "05ac22db75e9b262e9642de9ec1ec37990625f92bf426c52e12c88d",
        "0101010101010101010101010101010101010101010101010101010101010101"
      ],
      "prep_shares": [
        [
          "48771012eeda70a056cf2fd53022cf7b2edf45090eaa765c2b6",
          "b788efec11258f61fa53dd238a164da076bfd3f2e4bd966634b"
        ]
      ],
      "out_shares": [
        [
          408739992155304546
        ],
        [
          18038004077259279776
        ]
      ]
    }
  ]
}

```

Prio3Aes128Sum

For this test:

*The value of SHARES is 2.

*The value of bits is 8.

```

{
  "public_param": null,
  "verify_params": [
    [
      0,
      "01010101010101010101010101010101"
    ],
    [
      1,
      "01010101010101010101010101010101"
    ]
  ],
  "agg_param": null,
  "prep": [
    {
      "measurement": 100,
      "nonce": "01010101010101010101010101010101",
      "input_shares": [
        "05ac22dc75e9b243140aab68977b81a3958c5e79d27edaa00c6b338",
        "0101010101010101010101010101010101010101010101010101010101010101"
      ],
      "prep_shares": [
        [
          "2bcc0144c56fcf120a7aab22d57cde99fd2ee2301be4c59983d",
          "d433febb3a9030d1f58554dd2a8321682eff019a7a26e88471e"
        ]
      ],
      "out_shares": [
        [
          178602842237398423407215704739732627917
        ],
        [
          161679524683540039539650068628168138392
        ]
      ]
    }
  ]
}

```

Prio3Aes128Histogram

For this test:

*The value of SHARES is 2.

*The value of buckets is [1, 10, 100].

Email: rlb@ipv.sx

Christopher Patton
Cloudflare, Inc.

Email: chrispatton+ietf@gmail.com

Phillipp Schoppmann
Google

Email: schoppmann@google.com