

VDAF Analysis

Christopher Patton

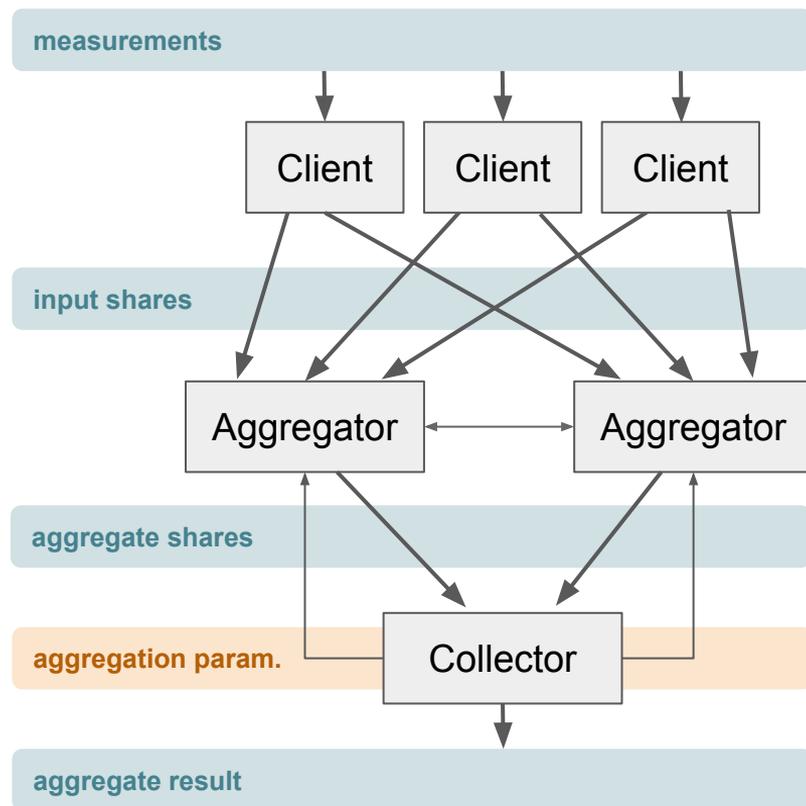
IETF 116 – UFMRG

Background

- [PPM \("Privacy Preserving Measurement"\) WG](#) will standardize methods of securely aggregating privacy-sensitive measurements generated client-side:
 - Probable exposure to illness (COVID-19 exposure notifications et al.)
 - Performance or threat-detection telemetry (operating system, browser, app, ...)
 - Ad conversion (how many sales are attributable to a given ad campaign)
- [draft-irtf-cfrg-vdaf](#): VDAFs ("Verifiable Distributed Aggregation Functions") are a class of lightweight **multi-party computation (MPC)** protocols that help address these use cases.
- [draft-ietf-ppm-dap](#): DAP ("Distributed Aggregation Protocol") specifies execution of VDAFs over HTTP.

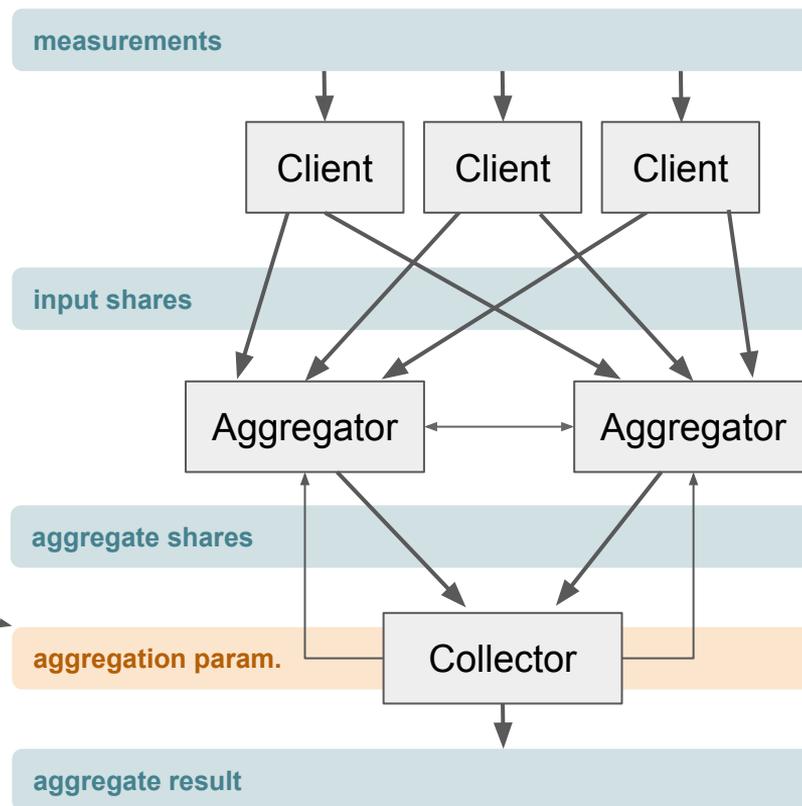
VDAF overview

- VDAF execution is distributed amongst a (large) set of **Clients**, a (small) set of **Aggregators** and a **Collector**:
 - Each client shards its **measurement** into **input shares**
 - Aggregators interact in order to prepare each set of input shares for aggregation
 - Refine them into an aggregatable form
 - Verify the refined shares correspond to a well-formed value
 - Each aggregator locally aggregates its refined shares into an **aggregate share**
 - Collector unshards the aggregate shares into the **aggregate result**.



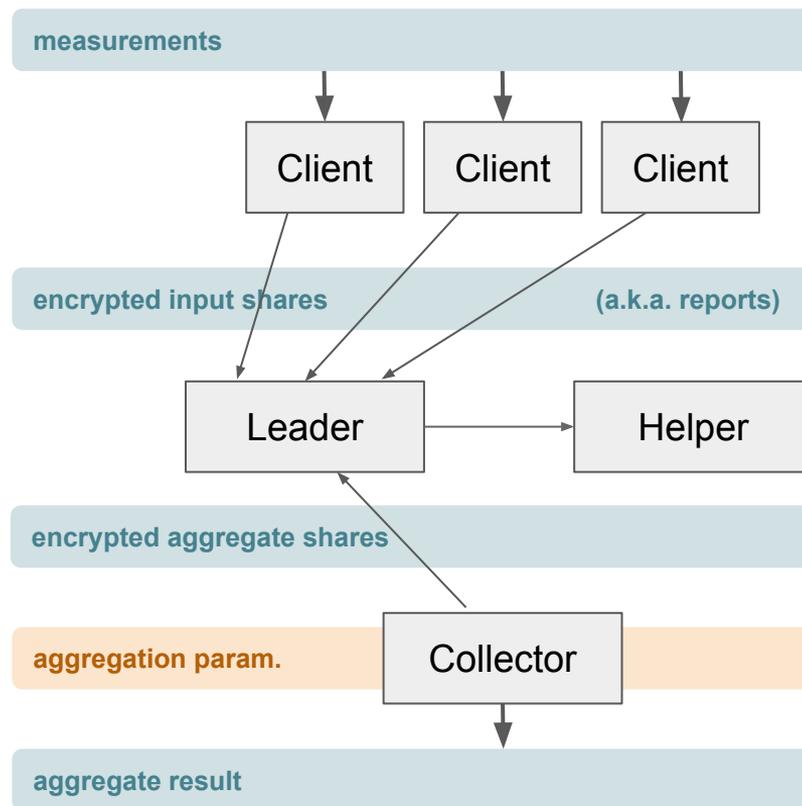
VDAF overview

- Notable features of this architecture
 - Bulk of the computation (sharding and preparation) is fully parallelizable across all measurements being aggregated
 - **Aggregation functions** $F(m_1, \dots, m_N)$ computable by VDAFs can be decomposed into $f(g(m_1), \dots, g(m_N))$, where f is linear and g is determined by the **aggregation parameter**
 - Not all of MPC, but a useful subset



DAP overview

- VDAF execution coordinated by a designated Aggregator (the **Leader**)
 - Upload: Client push encrypted input shares (a.k.a. **reports**) to the Leader
 - Aggregate: Leader works with **Helper** to prepare and aggregate the reports
 - Collect: The Collector pulls encrypted aggregate shares from the Leader
- HTTP used for transport
- [HPKE \(RFC 9180\)](#) for share encryption
- "Business logic": picking the VDAF, partitioning reports into batches, additional privacy parameters, etc.



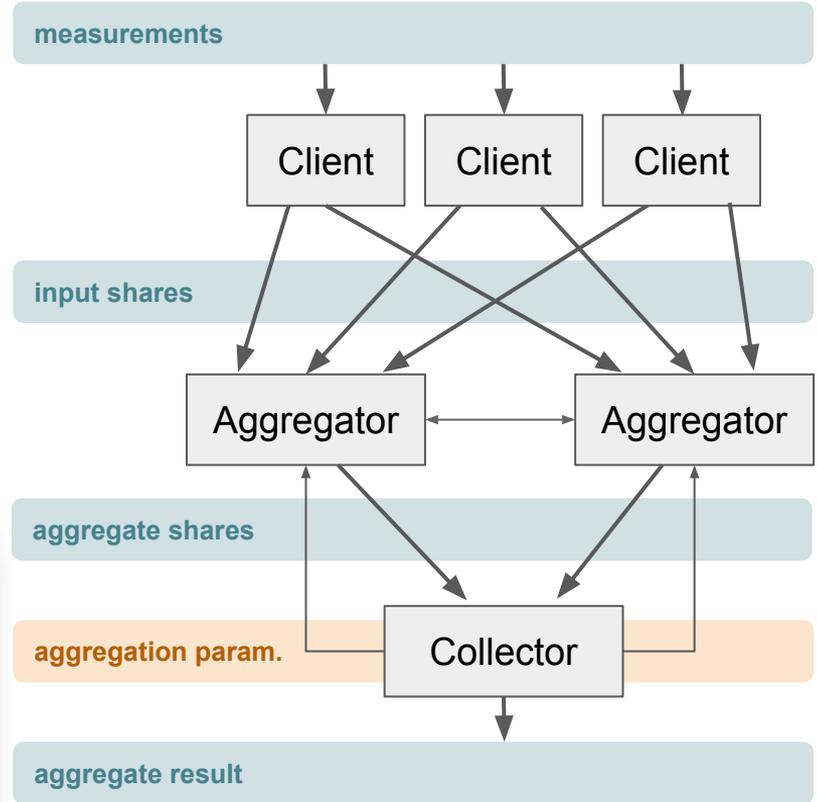
VDAF security considerations

- (Informal) security goals: **privacy** of measurements and **robustness** of the computation
 - Different trust models for privacy (one Aggregator is honest) and robustness (all Aggregators are honest)

9. Security Considerations

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.



VDAF security analysis of [DPRS23]

- *Starting goal*: Formalize security considerations of [draft-irtfg-cfrg-vdaf-03](#) (current draft at time of writing) and prove that the candidates (Prio3 and Poplar1) meet these goals
 - Model needs to fully account for how DAP uses VDAFs ([issue #161](#)):
 - Aggregators share a secret key ("VDAF verification key") for verifying each report.
 - DAP does not specify how the key is distributed. For maximum flexibility, it should be safe for a (malicious) Aggregator to pick the key unilaterally and distribute it to the Aggregators out-of-band.
- Our results:
 - Security proofs for a tweaked version of Prio3
 - Security proofs for a round-reduced variant of Poplar1 (called Doplar)

Security model of [DPRS23]: Robustness

Game $\text{Exp}_{II}^{\text{robust}}(A)$:

```
1  $sk \leftarrow \mathcal{SK}$ ;  $w \leftarrow \text{false}$ ;  $A^{\text{Prep}}()$ ; ret  $w$ 
Prep( $n \in \mathcal{N}$ ,  $\vec{x} \in \mathcal{X}^s$ ,  $msg_{\text{Init}} \in \mathcal{M}$ ,  $st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$ ):
2 if not  $II.\text{validSt}(\text{Used}[n], st_{\text{Init}})$ : ret  $\perp$ 
3  $\text{Used}[n] \leftarrow \text{Used}[n] \parallel (st_{\text{Init}},)$ 
4  $\text{Msg}[0, 1] \leftarrow msg_{\text{Init}}$ 
5  $y \leftarrow II.\text{refineFromShares}(st_{\text{Init}}, msg_{\text{Init}}, \vec{x})$ 
6 for  $\hat{j} \in [s]$ :  $\text{St}[\hat{j}] \leftarrow st_{\text{Init}}$ 
7 for  $\hat{\ell} \in [r + 1]$ :
8   for  $\hat{j} \in [s]$ :
9      $(sts, out) \leftarrow II.\text{Prep}(\hat{j}, sk, \text{St}[\hat{j}],$ 
10        $n, \text{Msg}[\hat{\ell}-1, \cdot], \vec{x}[\hat{j}])$ 
11     if  $sts = \text{running}$ :
12        $(\text{St}[\hat{j}], msg) \leftarrow out$ 
13        $\text{Msg}[\hat{\ell}, \hat{j}] \leftarrow msg$ 
14     else if  $sts = \text{finished}$ :
15        $y_{\hat{j}} \leftarrow out$ ;  $\tilde{w} \leftarrow [y \notin \mathcal{V}_{st_{\text{Init}}}]$ 
16     else if  $sts = \text{failed}$ : pass
17 if not  $\tilde{w}$ :
18    $\tilde{w} \leftarrow [y \neq II.\text{Unshard}(1, (II.\text{Agg}(y_{\hat{j}}))_{\hat{j} \in [s]})]$ 
19  $w \leftarrow w \vee \tilde{w}$ ; ret  $(w, \text{Msg})$ 
```

- Defined concretely in terms of a game played by a set of corrupt Clients interacting with a set of honest Aggregators.
- Adversary wins if:
 - (1) any Aggregator accepts an invalid report
 - (2) all Aggregators accept, but compute shares of the wrong value

Security model of [DPRS23]: Privacy

Game $\text{Exp}_{II,t}^{\text{Priv}}(A)$:

```

1   $(st_A, \mathcal{V}, (sk_j)_{j \in \mathcal{V}}) \leftarrow s A()$ 
2  if  $|\mathcal{V}| + t \neq s$  return  $\perp$ 
3   $b \leftarrow s \{0, 1\}$ 
4   $b' \leftarrow s A^{\text{Shard, Setup, Prep, Agg}}(st_A)$ 
5  ret  $b = b'$ 

```

Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):

```

6  if  $\text{Used}[\hat{k}] \neq \perp$ : ret  $\perp$ 
7   $n \leftarrow s \mathcal{N}$ 
8   $(\text{Pub}[\hat{k}], \text{In}[\hat{k}, \cdot]) \leftarrow s II.\text{Shard}(m_b, n)$ 
9   $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 
10 ret  $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, j])_{j \in \mathcal{T}})$ 

```

Setup($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$):

```

11 if  $\text{Status}[\hat{i}, \hat{j}] \neq \perp$ 
12   or not  $II.\text{validSt}(\text{Setup}[\cdot, \hat{j}], st_{\text{Init}})$ :
13   ret  $\perp$ 
14  $\text{Setup}[\hat{i}, \hat{j}] \leftarrow st_{\text{Init}}$ 
15  $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{running}$ 

```

Prep($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, \hat{k} \in \mathbb{N}, m\text{Msg} \in \mathcal{M}^*$):

```

16 if  $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$  or  $\text{In}[\hat{k}, \hat{j}] = \perp$ : ret  $\perp$ 
17 if  $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \perp$ :
18    $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \text{Setup}[\hat{i}, \hat{j}]; m\text{Msg} \leftarrow (\text{Pub}[\hat{k}], )$ 
19    $(n, m_0, m_1) \leftarrow \text{Used}[\hat{k}]$ 
20    $(sts, out) \leftarrow$ 
21      $II.\text{Prep}(\hat{j}, sk_{\hat{j}}, \text{St}[\hat{i}, \hat{j}, \hat{k}], n, m\text{Msg}, \text{In}[\hat{k}, \hat{j}])$ 
22   if  $sts = \text{running}$ :
23      $(st, msg) \leftarrow out; \text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow st$ 
24   else if  $sts = \text{finished}$ :
25      $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp; \text{Out}[\hat{i}, \hat{j}, \hat{k}] \leftarrow out$ 
26      $\text{Batch}_0[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_0; \text{Batch}_1[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_1$ 
27   else if  $sts = \text{failed}$ :  $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp$ 
28   ret  $(sts, msg)$ 

```

Agg($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}$):

```

29 if  $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$ : ret  $\perp$ 
30  $(st_1, \dots, st_s) \leftarrow \text{Setup}[\hat{i}, \cdot]$ 
31 if  $F(st_{\hat{j}}, \text{Batch}_0[\hat{i}, \hat{j}, \cdot]) \neq F(st_{\hat{j}}, \text{Batch}_1[\hat{i}, \hat{j}, \cdot])$ 
32   and  $(\forall j, j' \in \mathcal{V}) st_j = st_{j'} \wedge sk_j = sk_{j'}$ :
33   ret  $\perp$ 
34  $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{finished}$ 
35 ret  $II.\text{Agg}(\text{Out}[\hat{i}, \hat{j}, \cdot])$ 

```

- Defined concretely in terms of a game played by a corrupt Collector and subset of Aggregators
- Adversary mounts **chosen-batch attack**:
 - Via the **Shard** oracle, it chooses a pair of measurements for each honest Client
 - Game processes one of these, depending on the outcome of a coin flip
 - Adversary wins if it correctly guesses the coin flip

Impact of this work

- Revisions to Prio3
- Revised security considerations for DAP/VDAF documents (summarized in [this thread](#) on the CFRG mailing list) based on analysis of Prio3 and Doplar (Poplar1 variant)
 - VDAF verification key: Aggregators MUST commit to the key prior to reports being generated.
 - Nonce (accompanying each report): Clients MUST choose this at random.
 - Aggregation parameter: VDAFs MAY specify constraints on sequences of aggregation parameters used for the same report
- (Not a new observation, but bears repeating) Papers regularly leave gaps that need to be filled by corresponding specs
 - In our case: Instantiate "ideal coin-flipping functionality" presumed in prior work
 - Question for RG: **Can anyone think of examples?**

Aside: why games?

- Simulation paradigm is more conventional for MPC.
 - Universal Composability (UC): Define the ideal functionality for computing an aggregation function such that any VDAF that UC-realizes that functionality implies privacy and robustness.
- Consideration:
 - To guide parameter selection, it is helpful to obtain concrete security bounds for privacy and robustness *separately*.
 - Loose robustness bound may be tolerable in some applications, but a loose privacy bound is never tolerable.
 - Ideal functionality is complicated by the fact that privacy and robustness have different corruption models.
 - (Personal view of speaker): Protocols designed in the UC paradigm push details to spec authors that are not always fully addressed.

Future work

- Machine-checked proofs for VDAF candidates
 - Proof for Prio3 is pen-and-paper
 - No proof for Poplar1
- Symbolic or computational analysis of DAP. Two ways to frame this:
 - Prove that DAP preserves the security of the VDAF being executed (i.e., its usage in DAP does not result in an attack vector not accounted for in the existing security model)
 - Prove that DAP meets its own security goals, which are likely to go beyond privacy and robustness (i.e., differential privacy)