

Formalisation of Security for Federated Learning with DP and Attacker Advantage in IIIf for Satellite Swarms

Florian Kammüller^{1,2}

Abstract—In distributed applications, like swarms of satellites, machine learning can be efficiently applied even on small devices by using Federated Learning (FL). This allows to reduce the learning complexity by transmitting only updates to the general model in the server in the form of differences in stochastic gradient descent. FL naturally supports differential privacy but new attacks, so called Data Leakage from Gradient (DLG) have been discovered recently. There has been work on defenses against DLG but there is a lack of foundation and rigorous evaluation of their security. In the current work, we extend existing work on a formal notion of Differential Privacy for Federated Learning distributed dynamic systems and relate it to the notion of the attacker advantage. This formalisation is carried out within the Isabelle Insider and Infrastructure framework (III_f) allowing the machine supported verification of theory and applications within the proof assistant Isabelle. Satellite swarm systems are used as a motivating use case but also as a validation case study.

I. INTRODUCTION

In this paper, we review the formal characterization of Federated Learning (FL) and Differential Privacy (DP) as well as its formalisation in the Isabelle Insider and Infrastructure framework (III_f). This method allows the logical characterization of FL as an inductive protocol definition in the III_f providing a formal basis for expressing and verifying the security and privacy of applications. At the same time, it provides a formal definition of DP over dynamic executions of the FL protocol. As the formalisation is done in III_f, the formalisation can be used to prove “meta-level theory”, that is, theorems about the FL-algorithm can be proved within the interactive Isabelle theorem prover. As an example, a proof of a composition theorem is provided that allows to decompose the proof of DP of an FL architecture to the clients’ gradients. In addition to the meta-theory, it is possible to express application examples and reason about their security by applying the provided generic meta-theory theory.

In this paper, we extend the FL-DP formalisation in III_f by considering a more pragmatic way of characterizing security: Attacker Advantage ($\text{Adv}(\mathcal{A})$). To this end, we formalise the notion of $\text{Adv}(\mathcal{A})$ and relate it to the one of FL-DP. We further illustrate how the new meta-theory can be applied to a case study of satellite swarms.

II. BACKGROUND

A. Isabelle Infrastructure and Insider framework (III_f)

Isabelle is a generic Higher Order Logic (HOL) proof assistant. Its generic aspect allows the embedding of so-called object-logics as new theories on top of HOL. Object-logics, when added to Isabelle using constant and type definitions, constitute a so-called *conservative extension*. This means that no inconsistencies can be introduced. Isabelle supports modular reasoning, that is, we can prove theorems *in* an object logic but also *about* it. This allows the building of telescope-like structures in which a meta-theory at a lower level embeds a more concrete “application” at a higher level. Properties are proved at each level. Interactive proof is used to prove these properties but the meta-theory can be applied to immediately produce results. Moreover, Isabelle can be used to build frameworks by providing generic meta-theory for specific application domains.

The Isabelle Infrastructure and Insider framework (III_f) has been created initially for the modeling and analysis of Insider threats but is now a general framework for the state-based security analysis of infrastructures with policies and actors. Temporal logic and Kripke structure build the foundation. Meta-theoretical results have been established to show equivalence between attack trees and CTL statements. This foundation provides a generic notion of state transition on which attack trees and temporal logic can be used to express properties. In addition, proofs are facilitated because attacks can be derived using temporal logic analysis on Kripke structures by using the equivalence between attack trees and CTL statements. Thereby, the verification corresponds to a complete state exploration also known as model-checking.

B. Differential Privacy Federated Learning in III_f

The notion of Differential Privacy (DP) has become an ad hoc standard, most notably adopted for the US census [1] to obfuscate personal data in public data records. Federated Learning (FL) [14] is a widely used algorithm¹ for machine learning in heterogeneous distributed applications. FL facilitates machine learning for mobile distributed scenarios. Its natural relationship to DP has been investigated but the relationship to notions of noninterference in distributed systems is still largely open. Federated Learning and Differential Privacy are naturally related but formally their relationship is concerned with partitioning and distributing data securely over a network (the DP part) while safely only returning

¹Faculty of Science and Technology, Department of Computer Science, Middlesex University London, UK. <F.KammueLLer@mdx.ac.uk>

²Fakultät IV, Elektrotechnik und Informatik, TU Berlin, Germany. <info@tu-berlin.de>

¹Google scholar counts more than 60K citations for the seminal paper [14].

the stochastic gradient descent (SGD) that ML has procured locally (the FL part). Both concepts have been studied excessively but not many (with a few notable exceptions, for example, [17], [13], [12]) have observed that DP resembles the classical notion of noninterference used to express secure information flows. In this paper, we use this observation and extend it even further addressing the distribution and partitioning of the data over a network. We use the expressive power of the Isabelle Insider and Infrastructure framework (IIIf) to support this key issue of Federated Learning and relate them to classical Security notions of computational security by exploring the relationship to the cryptographic concept of Attacker Advantage ($\text{Adv}(\mathcal{A})$).

C. DP and FL

The IIIf supports modeling infrastructures, human actors and policies. It also allows representing data and its distribution across a network as well as the exchange of the SGDs between the client nodes and the server. For modeling the notions of DP, an extension of the IIIf with probabilities [9], [12] supports quantified specification of security and privacy over infrastructures. Thus, DP can be expressed directly as a probabilistic relation over a given infrastructure system. The use of IIIf allows expressing DP in a dynamic way as noninterference while taking into account the nature of FL, that is, the (potentially unbalanced) distribution of data over a network. The novelty of the approach [12] lies in considering DP and FL in the context of dynamically changing (data) systems, that is, with an explicit representation of data distributed on clients and servers as well as representing the FL algorithm formally as a protocol. Compared to the vast amount on works exploring DP and FL on applications to data sets, the focus of [12] is on the dynamicity (of data and communication) and data distribution aspects of DP and FL. Also, by contrast, it formally models all concepts in IIIf thus introducing mathematical rigour and automated verification.

The current paper now extends the existing NI-FL-DP concept by establishing – within the formal framework IIIf – the conceptual relationship to classical complexity theory notions of computational security via the Attacker Advantage ($\text{Adv}(\mathcal{A})$).

D. Attacker Advantage

Many security notions are characterized by the *computational indistinguishability* between two distributions D_0 and D_1 . This computational indistinguishability is measured by the *advantage* an adversary $\mathcal{A}$ can have in distinguishing the two distributions. That is, we define the Advantage $\text{Adv}(\mathcal{A})$ as follows.

$$\text{Adv}(\mathcal{A}) = P_{x \in D_1}(\mathcal{A}(x) = 1) - P_{x \in D_0}(\mathcal{A}(x) = 1)$$

where $\mathcal{A}(x) = 1$ means that the attacker $\mathcal{A}$ guesses that the input x is in distribution D_1 (while $\mathcal{A}(x) = 0$ would mean $x \in D_0$).

For example, for the Diffie-Hellman key exchange protocol, with Diffie-Hellman tuples in $\mathbb{G} = \langle g \rangle$, a group of prime order p , spanned by a generator g , and denoted additively.

Definition 2.1 (Decisional Diffie-Hellman Problem (DDH)):

The DDH assumption in a group $\mathbb{G}$ of prime order p , with a generator g , states that the distributions D_0 and D_1 are computationally hard to distinguish, where

$$\begin{aligned} D_0 &= \{(a \cdot g, b \cdot g, ab \cdot g).a, b \xleftarrow{\$} Z_p\} \\ D_1 &= \{(a \cdot g, b \cdot g, c \cdot g).a, b, c \xleftarrow{\$} Z_p\} \end{aligned}$$

The notation $k \xleftarrow{\$} K$ is used in Cryptography to express “selecting a random string from [a uniformly distributed set of keys] K and naming it k ” [7][p. 63]. The set of keys in our case is Z_p which is another common mathematical notion for the prime number field $\{0, \dots, p-1\}$ with addition and multiplication modulo p . We then denote as $\text{Adv}_G^{\text{ddh}}(\mathcal{A})$ the advantage of an adversary $\mathcal{A}$ in distinguishing D_0 and D_1 . Looking back at the definition of Adversarial Advantage above, this means that if the DDH-problem is indeed computationally hard (as we believe it to be), then the attacker will not be able to distinguish whether c is the product ab of the secrets a and b . That is, they will output a 1 with about the same probability when exposed to an example that is in D_0 as for one that is in D_1 , so that the difference $\text{Adv}(\mathcal{A})$ (defined as $\text{Adv}(\mathcal{A}) = P_{x \in D_1}(\mathcal{A}(x) = 1) - P_{x \in D_0}(\mathcal{A}(x) = 1)$) is close to 0.

This example illustrates the notion of $\text{Adv}(\mathcal{A})$ on the well-known DH problem; we will adapt it to the setting of FL and DP.

III. FORMALISING THE ATTACKER ADVANTAGE FOR FL-DP

A. FL Algorithm

FL combines the requirements of training on distributed clients where the data may be unevenly distributed (“unbalanced” [14]) and local data sets may not be representative of the global one (“non-iid” [14]), massively distributed and with limited communication.

Algorithm 1 FL

```

1: Server executes:
2: initialize  $w_0$ ;  $K := \text{set of clients}$ 
3:   for each round  $t = 1, 2, \dots$ 
4:     for each client  $k \in K$  in parallel do
5:        $w_{t+1}^k := \text{ClientUpdate}(k, w_t)$ 
6:      $w_{t+1} := \sum_{k=1}^K \frac{n_k}{n} w_{t+1}^k$ 
7:   ClientUpdate( $k, w$ ):
8:      $w := w - \eta \nabla \mathcal{L}(w; P_k)$ 
9:   return  $w$  to Server

```

The FL algorithm can be applied to any finite-sum objective as follows [14]

$$\min_{w \in \mathbb{R}^d} f(w) \quad \text{where} \quad f(w) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(w).$$

where the f_i are loss functions, that is, $f_i(w) = \mathcal{L}(x_i, y_i; w)$ where (x_i, y_i) is an example data and w is the model parameter, that is, the current value for the parameters of the ML algorithm to get the best prediction. So, the FL algorithm

tries to approximate a prediction function by minimizing a loss function (sometimes called cost function) on data example points (the training data). What is more is that the FL algorithm considers a set of clients K over which the data is partitioned. Let P_k be the partition for client $k \in K$ and $n_k = |P_k|$ the number of data points in that partition. This adds a layer of detail to highlight the unbalanced scenario for the FL algorithm. We can re-write the above description by detailing the averaging in each client.

$$f(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w) \quad \text{where} \quad F_k(w) = \frac{1}{n_k} \sum_{i \in P_k} f_i w.$$

B. Formalisation of FL-DP protocol

The formalization of FL in IIIf implements the FL algorithm introduced in the previous sections as protocol. That is, the execution of the FL algorithm is a sequence of exchanges between the server and a number of distributed clients. The framework IIIf supports well this formalization of protocols. For example, Quantum Key Distribution [9], [11] also models protocols built on top of the basic Kripke structures with CTL and Attack Trees thus providing concepts and techniques for formalizing FL in IIIf.

In general, infrastructures in the IIIf are graphs with actors and local policies attached to the nodes. For the FL algorithm, the actors are the server and its clients attached each to their nodes (location). The data can be explicitly represented in infrastructures and is also attached to the nodes corresponding to the server and its clients. The constructor `Lgraph` of the data type `igraph` assembles the components `gra`, `aloc`, `server`, `clients`, `ready`, `gradient`, `curmodpar`, `partition`, and `dataset` which represent: the actual topology of locations, the allocation of actors (that is, also devices) to the locations, the server and clients, the clients finished (“ready”) with learning in a round of the FL algorithm, the local gradients currently held in a location, the global current model parameter used for the distributed learning, the partition of the global data set onto the distributed system, and the global data set. The name introduced for each component can be used as a projection functions on the `igraph` to select each of the components. The parameter type α in the components `gradient` and `curmodpar` is given as a type variable. This permits representing arbitrary vectors of model parameters (coefficients).

```
datatype igraph = Lgraph gra: location set
                    aloc: actor  $\Rightarrow$  location
                    server: actor
                    clients: actor set
                    ready: actor set
                    gradient: actor  $\Rightarrow \alpha :: \text{num}$ 
                    curmodpar:  $\alpha$ 
                    partition: actor  $\Rightarrow$  data set
                    dataset: data set
```

The type `igraph` represents infrastructures that we use for modeling distributed scenario. Protocol execution traces are sets of lists of events: each list represents a finite execution

trace of the protocol. The set of such traces is a snapshot of execution traces of the system. The possibility of consideration of parallel traces is crucial for security; many attacks come from exploits between simultaneous protocol runs. For FL the events in a protocol trace correspond to the communication actions of the FL algorithm: as an initial step *putting* the partition of the data on the clients, *getting* back the local gradient descents from the clients to the server, and *evaluation* of the combined local gradient descents by the server once all clients have been harvested to compute and redistribute the new current model parameter for the next learning round.

The theory of Kripke structures and temporal logics CTL in the IIIf supports a generic state transition relation that can be instantiated as $\rightarrow_{FL}$ to implement the FL algorithm. This relation can then be defined as an inductive relation over the infrastructure type. Consequently, it has three cases:

- **put_part** is the initial deployment of the partitioned global data set to all the clients by inserting the data into disjoint partitions.
- **get_grad** is the case where a local gradient results from the local learning in a client the position `c`. The result is loaded into the graphs’ component `gradient G` by updating this function for client `c`. Also, the current client `c` is added to the set `ready G` signifying that `c` has updated its gradient in this round of FL.
- **eval_server** finally, pulls together the new model parameter by averaging all local gradients. The event **eval_server** is only enabled if the set `ready` contains all clients. After the `eval` step, `ready` is initialized to the empty set for a new round.

For the inductive definition of the state transition system of $\rightarrow_{FL}$ that defines the above informally described semantics of the FL algorithm in IIIf see [12], [10]. We do not need that level of detail for the considerations in this paper.

When formalizing the security properties of FL systems we need to express probabilities of state transition execution traces. To this end, the IIIf offers a probabilistic path quantifier [9] that allows to assign probabilities to paths that end in a specified set of states. Thus, specific properties can be expressed and quantified by probabilities, for example, that an FL system leads to states with certain model parameters or in which specific gradients were transmitted. This path quantifier is relative to a Kripke structure M because “gradient transmission” paths must start from initial states and be paths along states within M . The IIIf provides the following specialized path selector $M \vdash F s$ for Kripke model M and state s . The list operator `hd` yields the head (first element) of a list – here ensuring the trace `l` starts in an initial state – and the operator `F` selects a set of finite paths leading into a state s .

definition $M \vdash F s \equiv$
 $\{ l. \text{ set } l \subseteq \text{states } M \wedge \text{hd } l \in \text{init } M \} \cap F s$

This is now combined with the probability notions provided in the IIIf framework. For example, to introduce the probability for a set of finite paths given an assignment `ops` of

probabilities to paths representing outcomes of protocol runs, the operation pmap ops produces a probability distribution for this protocol. This can then be applied to the previous path selection: $\text{pmap ops } (M \vdash Fs)$ computes the probability $P(M \vdash Fs)$ of the event that paths end in state s .

Given these notions, probabilistic noninterference for FL can be defined based on DP.

C. Definition of NI-FL-DP and Compositionality Theorem

Noninterference (NI) [5] is a formal definition of secure information flows. It is an equivalence property over sequences of state transitions (or sequences of events – depending on the system view). Intuitively, it characterizes that from a certain viewpoint various executions of a system do not reveal any secret information. For programming languages, noninterference is defined as an indistinguishability relation of program executions. The indistinguishability is respective to visible system states or visible input/output variables. It thus requires some observer viewpoint level which is commonly added as $H(igh)$ and $L(ow)$ labels that are variables. It then declares system execution (states) to be equivalent if they only differ in values of H -variables. Noninterference is often defined as a Low-bisimulation: in any program run the indistinguishability of pairs of system states is preserved. This implies that the computation of L -variables is independent of that of H -variables. No unintended “illicit” information flows from H to L can appear in program execution.

The notion of NI of FL by DP (NI-FL-DP) [12] is inspired by Volpano and Smith’s probabilistic NI [18]: “under two memories that may differ [only] on H -variables, we want to know that the probability that a concurrent program can reach some global configuration under one of the memories is the same as the probability that it [the concurrent program] reaches an equivalent configuration under the other” [18]. For “concurrent program”, we read “distributed infrastructure”, for “differs on H -variables” we have “the datasets differ only by one” and for “equivalent configuration” we have differential privacy, that is, the probability of reaching equivalent configurations differs only by the DP-factor. Thus, this notion gives rise to a notion of NI-FL-DP: noninterference of federated learning as differential privacy,

The definition of NI-FL-DP in IIIf states that for two infrastructure states s_0 and s_1 that both originate in one initial infrastructure I and whose data sets differ only by one element (expressed using the $\overset{1}{\Diamond}$ operator), the probabilities of finite traces leading into such states differ only by the DP-factor e^ϵ .

definition $NI_FL_DP \ I \ s_0 \ s_1 \ \epsilon \equiv$
 $I \xrightarrow{*}_{FL} s_0 \implies I \xrightarrow{*}_{FL} s_1 \implies$
 $M = \{\text{Kripke } \{t \ . \ I \xrightarrow{*}_{FL} t\} \{I\}\} \implies$
 $G_0 = \text{igra } s_0 \implies G_1 = \text{igra } s_1 \implies$
 $\text{dataset } G_0 \overset{1}{\Diamond} \text{dataset } G_1 \implies$
 $P(M \vdash F \{s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_0\}) \leq$
 $e^\epsilon * P(M \vdash F \{s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_1\})$

Since this notion is inspired by probabilistic NI for distributed systems, it is not surprising that the methodologies

and techniques applied in the NI domain can be adapted to provide useful tools for the verification and analysis also for NI-FL-DP. Many works on NI (and already Goguen and Meseguer in their seminal paper) introduced so-called “unwinding relations” to simplify proving noninterference [6]. Unwinding theorems allow stepwise (sequential) decomposition of NI. Since the NI-FL-DP property is a security notion for *distributed* systems, an unwinding theorem for this notion consequently needs to decompose the global property to its (distributed) components. The following compositionality theorem is provided in the IIIf framework to break down the proof of security of an FL system to that of its client components.

theorem $NI_Decomposition:$

$$\begin{aligned} I \xrightarrow{*}_{FL} s_0 &\implies I \xrightarrow{*}_{FL} s_1 \implies \\ M = \{\text{Kripke } \{t \ . \ I \xrightarrow{*}_{FL} t\} \{I\}\} &\implies \\ G_0 = \text{igra } s_0 \implies G_1 = \text{igra } s_1 &\implies \\ \text{clients } G_0 = \text{clients } G_1 &\implies \\ \forall c \in \text{clients } G_0. \text{partition } G_0 \ c \overset{1}{\Diamond} &\text{partition } G_1 \ c \\ \implies & \\ P(M \vdash F \{s. \text{gradient}(\text{igra } s) \ c = \text{gradient } G_0 \ c\}) &\leq \\ e^\epsilon * P(M \vdash F \{s. \text{gradient}(\text{igra } s) \ c = \text{gradient } G_1 \ c\}) & \\ \implies NI_FL_DP \ I \ s_0 \ s_1 \ \epsilon & \end{aligned}$$

The precondition for this Decomposition Theorem assumes that for two state graphs that differ in their local dataset by only one item of the data set, each client’s gradient is differentially private in the local learning of the FL-algorithm. Given this precondition, the theorem allows to infer that the global algorithm computing the average of the local gradients, is also differentially private. That is, the NI_FL_DP property holds.

With the notion of NI-FL-DP, we have a good method for expressing information flows – and thus leakage – for FL. In fact, it is very desirable to have quantifiable notions of information flow because in reality we cannot hope to have 100% security but at least need a way to quantify it. The additional Decomposition Theorem is furthermore a very practical tool since it allows to break down the global NI-FL-DP property to that of the local components of the FL system. Practically, the ϵ can be computed as the maximum value of the local components’ ϵ . Ultimately, the ϵ values are, however, measured values providing quantification of an upper bound based on empirical data. To counter this rather empirical way of estimation a security and privacy boundary, it is desirable to have a more classical metric. This motivates the following extension of relating the NI-FL-DP property to $\text{Adv}(\mathcal{A})$ (see Section II-D).

D. Relationship to Attacker Advantage

Decisional NI-FL-DP problem: The assumption NI-FL-DP can be expressed as a decision problem following the outline of Definition 2.1.

Definition 3.1 (Decisional NI-FL-DP Problem): Given that

$$\begin{aligned} I \xrightarrow{*}_{FL} s_0 \wedge I \xrightarrow{*}_{FL} s_1 \wedge \\ M = \{\text{Kripke } \{t \ . \ I \xrightarrow{*}_{FL} t\} \{I\}\} \wedge \\ G_0 = \text{igra } s_0 \wedge G_1 = \text{igra } s_1 \wedge \\ \text{dataset } G_0 \overset{1}{\Diamond} \text{dataset } G_1 \end{aligned}$$

the NI-FL-DP assumption states that the distributions D_0 and D_1 are computationally hard to distinguish, where

$$D_0 = M \vdash F \{ s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_0 \}$$

$$D_1 = M \vdash F \{ s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_1 \}$$

We denote as $\text{Adv}^{\text{dnifldp}}(\mathcal{A})$ the advantage of an adversary $\mathcal{A}$ in distinguishing D_0 and D_1 when seeing (public) current model parameters $\text{curmodpar } G_i, i \in \{0, 1\}$.

Note, that the above distributions are sets of traces of the state transition relation of the FL system leading into the state sets designated after the $\vdash F$ -operator. Thus the distributions D_0 and D_1 are effectively the same probability distributions as used in the NI-FL-DP definition. Also it is important to note at this point that the NI-FL-DP problem is defined as an assumption. Clearly, the computational hardness of the indistinguishability depends on the problem of inverting the prediction function based on the gradients. However, Definition 3.1 provides new metrics for security and privacy of FL systems based on the concept of DP.

The definition of the assumptions can be applied to categorize the security of techniques for defending FL architectures using the DP techniques. For example, Chang and Zhu [2] report on the concept of “Deep Leakage from Gradient (DLG)” [19] which refers generally to clients inadvertently revealing sensitive properties of their local training data. In [2], the authors report on two techniques, *gradient sparsification* and *pseudo-gradients*, to protect from DLG. The contribution of our paper is to make the security and privacy of such protection mechanisms quantifiable and amenable to automated reasoning and verification in IIIf. To this end the definitions presented in this section can be applied.

We can use here the notion of $\text{Adv}(\mathcal{A})$ (see Section II-D) to effectively relate the computational security characterization to DP. The following equivalence holds between $\text{Adv}(\mathcal{A})$ and the DP-factor for FL.

Theorem 1 (Adv-FL-DP-Eq): Let Kripke structure M be defined as $\{\text{Kripke } \{t. I \rightarrow_{FL}^* t\} \{I\}\}$. Furthermore, let $I \rightarrow_{FL}^* s_0$ and $I \rightarrow_{FL}^* s_1$ where $G_0 = \text{igra } s_0$ and $G_1 = \text{igra } s_1$ such that that is, their datasets differ only in x . Given two distributions

$$D_0 = M \vdash F \{ s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_0 \}$$

$$D_1 = M \vdash F \{ s. \text{curmodpar}(\text{igra } s) = \text{curmodpar } G_1 \}$$

Then $\text{Adv}^{\text{dnifldp}}(\mathcal{A}) = 1 - \frac{1}{e^\epsilon}$ where $e^\epsilon = \frac{P(D_0)}{P(D_1)}$.

Proof: Intuitively, the distributions D_0, D_1 are computationally hard to distinguish iff their NI-FL-DP factor ϵ is close to 0. By Definition 3.1, the attacker’s advantage and NI-FL-DP s_0, s_1, ϵ have the following equivalences

$$\begin{aligned} \text{Adv}^{\text{dnifldp}}(\mathcal{A}) = 0 &\quad \equiv \quad \frac{1}{e^\epsilon} = 1, \text{ if } \epsilon = 0 \\ 0 < \text{Adv}^{\text{dnifldp}}(\mathcal{A}) \leq 1 &\quad \equiv \quad \frac{1}{e^\epsilon} \rightarrow 0, \text{ if } \epsilon > 0 \end{aligned}$$

For the case $\epsilon = 0$, we have perfect indistinguishability, i.e. $\text{NI_FL_DP } s_0, s_1, 0$, and the attacker’s advantage is also 0. For the non-zero case, the advantage grows towards 1 as ϵ goes to infinity. Unfolding the definition of advantage (see Section II-D), we need to show $P_{x \in \text{dataset } G_1}(\mathcal{A}(x) = 1) - P_{x \in \text{dataset } G_0}(\mathcal{A}(x) = 1) = 1 - \frac{1}{e^\epsilon}$. This is straightforward calculus (for details, see [10]). ■

Practically, the theorem is useful when considering applications as discussed in the next section. It allows to compute concrete bounds on NI-FL-DP as we illustrate now by a general application example. From the literature on complexity theory, we can assume that (stochastic) gradient descent is computationally hard [4]. Furthermore, its inversion is also computationally hard [3]. In terms of security of the SGD algorithm, this means that the attacker is bound with polynomial time complexity. That is, the attacker’s advantage decreases faster than any inverse polynomial $p(n)$ for sufficiently large n : $\text{Adv}(\mathcal{A}) < \frac{1}{p(x)}$. With Theorem 1, we have thus $\text{Adv}(\mathcal{A}) = 1 - \frac{1}{e^\epsilon} < \frac{1}{p(x)}$ which is equivalent to $e^\epsilon < \frac{p(n)}{p(n)-1}$. Illustrative examples for smaller polynomial grades are contained in the following table. It shows that the NI-FL-DP factor ϵ corresponds to the inverse grade of the bounding polynomial. Since the attacker’s advantage decreases faster than any polynomial, ϵ gets arbitrarily close to 0.

$p(n)$	n	n^2	n^3	n^{10}
$n = 100$	100	10^4	10^6	10^{20}
$e^\epsilon < \dots$	1.01	1.0001	1.000001	$10^{20}/10^{20} - 1$
$\epsilon < \dots$	10^{-3}	10^{-4}	10^{-6}	10^{-20}

Theorem 1 can also be decomposed by applying Theorem NI-Decomposition to break down the advantage to the client components. As seen in the above application example, even if we restrict the batch size for SGD to $n = 100$, we still get near perfect security because ϵ grows arbitrarily close to 0.

IV. SATELLITE SWARM APPLICATION

Satellite swarms are collaborations of satellites that function as distributed systems in space. Their spatial distribution allows them to operate from different locations – often for observational purposes – and share tasks to reduce the operational load. An example is the European Space Agency’s (ESA) project Swarm (see Figure 1). Scientific projects like Swarm usually have a humanitarian and ethical motivation and do not seek to hide their data. However, it can be easily imagined that security and privacy can become an issue, for example, to protect private data of individuals or security criticality, like locations of secret military assets. As an illustrative case study, let *Moniteo* be a Swarm-like application to monitor temperature data on the Earth’s surface using FL. That is, each *Moniteo*’s agent satellites σ is geo-stationary and equipped with an infrared camera. The mission is to find a hypothesis (model) for the temperature in given locations using FL. Initially, each σ is equipped with a dataset Part_σ of locations and their temperatures. Certain locations must be omitted in the data set because they convey private information. That is, for geographic data sets D_0, D_1 we have that $D_0 \dot{\Delta} D_1$ for the partition of data sets. Even if σ learns no data, it will only serve to improve the local hypothesis: the new data does not need to be transported to the server. Now, *Moniteo*’s security depends on low (or no) information flows from the local data to the SGDs transmitted to the server. Here, Theorem 1 and its corollary




Fig. 1. Swarm is ESA's first constellation of Earth observation satellites designed to measure the magnetic signals from Earth's core, mantle, crust, oceans, ionosphere and magnetosphere, providing data that will allow scientists to study the complexities of our protective magnetic field, ©ESA.

for the decomposition of the Advantage yield that local deviations of data, e.g., by severe temperature increase in a new data point, cannot be distinguished by an attacker. That is, if a client satellite σ would have two data partitions that differ in one data point because of a different measurement of that one data point, then the difference in the gradients would result in a difference below e^ϵ . Consequently, if the satellite σ 's infrared cameras measure a significant difference of temperature in a location that is marked as secret, it may safely omit the data point's value from the gradient computation without risking significant information flows above the threshold of computational security.

V. CONCLUSIONS AND RELATED WORK

In this paper, we have revisited the notion of Information Flow Security (NI) for Federated Learning (FL) by Differential Privacy (DP) (NI-FL-DP) and related it to Computational Security given by the cryptographic notion of Attacker Advantage $\text{Adv}(\mathcal{A})$. We illustrated its importance by a satellite swarm use case.

Information theoretic aspects of DP are studied intensively at the moment as is documented in the survey paper [15]. While the foundations appear to become clarified, the relationship to practical implementations of DP, that is, inference or verification systems and thus notions like noninterference are still fairly open. In particular, works that support the formalization and automated reasoning are very scarce with a few notable exceptions. To our knowledge, amongst this there are none in relation to federated learning and modeling distribution of data, communication and processing as we do using the infrastructures, actors, and actions of the IIIf. Tschantz et al [17] already observed the similarity between noninterference and differential privacy with respect to verification. More recently, Liu et al [13] took these concepts

further and presented a dedicated model checking technique. However, both works do not consider the conceptual relation of FL and DP to fundamental notions of Cryptography like computational hardness and $\text{Adv}(\mathcal{A})$.

The notion of DP naturally corresponds to the NI property which gives rise to a natural way of defining a quantified notion of NI-FL-DP (see Section III-C). The quantification is useful to relate practically motivated attack techniques, like the DLG, and their defense techniques, that is, gradient sparsification and pseudo-gradients or fundamental notions of cryptography and to rigorous formalisation and automated reasoning and verification in IIIf.

REFERENCES

- [1] U. C. Bureau. Disclosure avoidance for the 2020 census: An introduction, 2021. accessed 2.10.2025.
- [2] W. Chang and T. Zhu. Gradient-based defense methods for data leakage in vertical federated learning. *Computers & Security* **139**, 103744, Elsevier, 2024.
- [3] F. Duam Keyes, C. Hegde. On the Fine-Grained Hardness of Inverting Generative Models. arXiv 23.09.05795v1, <https://arxiv.org/pdf/2309.05795>, 2023.
- [4] J. Fearnley, P. W. Goldberg, A. Hollender, R. Savani. The Complexity of Gradient Descent: $\text{CLS} = \text{PPAD} \cap \text{PLS}$. *Journal of the ACM* **70**:1, <https://doi.org/10.1145/3568163>, 2022
- [5] J. Goguen and J. Meseguer. Security policies and security models. *Symposium on Security and Privacy, SOSP'82*, pages 11–22. IEEE Computer Society Press, 1982.
- [6] J. A. Goguen and J. Meseguer. Unwinding and inference control. *Symposium on Security and Privacy, SOSP'84*, pages 75–87. IEEE Computer Society, 1984.
- [7] S. Goldwasser and M. Bellare. Lecture Notes on Cryptography, 2008. Available at <https://cseweb.ucsd.edu/~mihir/papers/gb.pdf>. Accessed 2. September 2025.
- [8] H. Gong, L. Jiang, X. Liu, Y. Wang, L. Wang, K. Zhang. Recover user's private training image data by gradient in federated learning. *Sensors*: **22**, 7157. 2022.
- [9] F. Kammüller. Attack trees in isabelle extended with probabilities for quantum cryptography. *Computers & Security*, 87, Elsevier, 2019.
- [10] F. Kammüller. Formalisation of Security for Federated Learning with DP and Attacker Advantage in IIIf for Satellite Swarms, extended version. arxiv, 2025.
- [11] F. Kammüller, R. Nagarajan, M. C. Parker, C. White. Formalisation and Analysis of Decoy QKD in the IIIf using Refinement and Attack Trees. *Systems, Men, and Cybernetics, SMC 2025*, IEEE, 2025
- [12] F. Kammüller, L. Piras, B. Fields, R. Nagarajan. Formalizing Federated Learning and Differential Privacy for GIS systems in IIIf. *SecAssure 2024, co-located with ESORICS'24*. LNCS, volume 15264, Springer 2024.
- [13] D. Liu, B. Wang, C. Fu, and L. Zhang. Model checking differentially private properties. *Theor. Comput. Sci.*, 943:153–170, 2023.
- [14] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. *International Conference on Artificial Intelligence and Statistics, AISTATS PMLR***54**(1273–1282), 2017.
- [15] A. Ünsal and M. Önen. Information-theoretic approaches to differential privacy. *ACM Comput. Surv.*, 56(3), oct 2023.
- [16] European Space Agency. Swarm reveals growing weak spot in Earth's magnetic field. https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Swarm/Swarm_reveals_growing_weak_spot_in_Earth_s_magnetic_field. Accessed 7. 11. 2025.
- [17] M. C. Tschantz, D. K. Kaynar, and A. Datta. Formal verification of differential privacy for interactive systems MFPS, ENTCS **276**(61–79), Elsevier, 2011.
- [18] D. M. Volpano and G. Smith. Probabilistic noninterference in a concurrent language. *J. Comput. Secur.*, 7(1), 1999.
- [19] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.