



Gábor Danner <gabor.danner@gmail.com>

IFIP DAIS 2015 notification for paper 24

IFIP DAIS 2015 <ifipdais2015@easychair.org>
Címzett: Gábor Danner <gabor.danner@gmail.com>

2015. március 23. 23:13

Dear Gábor Danner,

We are happy to inform you that your paper 24, entitled "Fully Distributed Privacy Preserving Mini-Batch Gradient Descent Learning", was accepted to DAIS 2015.

The reviewers worked very hard to give good feedback and select the best papers among the submitted. Please take into consideration their comments (for some few papers, we will contact the authors to highlight certain specific requests).

The instructions for preparing the camera-ready version will be sent to you in the next days. Be aware that we have a tight schedule and the final papers need to be provided by April 2nd.

Congratulations and see you in Grenoble,

Alysson and Sara

----- REVIEW 1 -----

PAPER: 24

TITLE: Fully Distributed Privacy Preserving Mini-Batch Gradient Descent Learning

AUTHORS: Gábor Danner and Mark Jelasity

OVERALL EVALUATION: 2 (accept)

REVIEWER'S CONFIDENCE: 2 (low)

Familiarity: Please indicate your familiarity with the paper's subject matter: 2 (Passing familiarity)

Novelty: How would you rate the novelty of the paper?: 4 (New contribution)

Impact: Will this paper be important over time?: 4 (Very likely)

----- REVIEW -----

This paper presents a distributed protocol to privately compute the sum of values contributed by a set of participants in the semi-honest adversarial model. The protocol deliberately undervalues precision over speed and robustness, as it is aimed to be used in a mini-batch gradient descent learning algorithm — an application domain where such a rough computation is acceptable.

The protocol assumes N participants to be connected by a random tree-shaped topology (set up by some other mean). In essence, the algorithm performs the usual bottom-up structural summation: each node aggregates the results from the descendants; sums its own local value and passes the result to the ancestor. However, in order to achieve privacy, computations are distributed among S consecutive nodes along the path to the root (S is a security parameter that sets the desired privacy level). Technically, this is performed by combining an additive secret sharing scheme with an additive homomorphic public-key encryption scheme. This way, the adversary needs to corrupt at least S consecutive nodes in order to gain access to the partial subtree summations. Apart from discussing the protocol security, the paper also discusses other aspects of the deployment of the protocol, such as the size/number of messages; robustness; and possible variations.

The protocol is ingenious and might well find other interesting applications. Overall, the paper is well structured and the presentation is clear. There is, however, an issue that has puzzled me: when the authors discuss the parameter choice for the underlying Paillier cryptosystem (Section 6.2), only constraints on the size of plaintext are taken into account. I would expect that security and efficiency of the Paillier cryptosystem itself to play a primary role (probably conditioning some of the remaining protocol parameters). Therefore, I would suggest the authors to clarify and to detail in the paper the choice for the parameters used in the experiments.

----- REVIEW 2 -----

PAPER: 24

TITLE: Fully Distributed Privacy Preserving Mini-Batch Gradient Descent Learning
AUTHORS: Gábor Danner and Mark Jelasity

OVERALL EVALUATION: 1 (weak accept)

REVIEWER'S CONFIDENCE: 2 (low)

Familiarity: Please indicate your familiarity with the paper's subject matter: 1 (Not my area)

Novelty: How would you rate the novelty of the paper?: 3 (Incremental improvement)

Impact: Will this paper be important over time?: 3 (Can't quite tell)

----- REVIEW -----

I do not have the necessary expertise to assess the novelty of the contributions. The paper is well written and I was able to follow the model, proposed solution, and the demonstration that the goals were achieved.

----- REVIEW 3 -----

PAPER: 24

TITLE: Fully Distributed Privacy Preserving Mini-Batch Gradient Descent Learning

AUTHORS: Gábor Danner and Mark Jelasity

OVERALL EVALUATION: 2 (accept)

REVIEWER'S CONFIDENCE: 4 (high)

Familiarity: Please indicate your familiarity with the paper's subject matter: 3 (Familiar with area)

Novelty: How would you rate the novelty of the paper?: 4 (New contribution)

Impact: Will this paper be important over time?: 4 (Very likely)

----- REVIEW -----

This paper makes a small but novel contribution in presenting an algorithm to efficiently calculate the sum of gradients as part of a fully distributed mini-batch stochastic gradient descent algorithm. It is a well-written paper and is mostly self-contained. The notion of a semi-honest adversary, however, needs further elaboration. The algorithm, its security properties and complexity are presented and analyzed, and the convergence speed of the algorithm on binary classification data is also evaluated.

Overall, this is a good paper and acceptable to DAIS

----- REVIEW 4 -----

PAPER: 24

TITLE: Fully Distributed Privacy Preserving Mini-Batch Gradient Descent Learning

AUTHORS: Gábor Danner and Mark Jelasity

OVERALL EVALUATION: 2 (accept)

REVIEWER'S CONFIDENCE: 4 (high)

Familiarity: Please indicate your familiarity with the paper's subject matter: 3 (Familiar with area)

Novelty: How would you rate the novelty of the paper?: 4 (New contribution)

Impact: Will this paper be important over time?: 4 (Very likely)

----- REVIEW -----

The authors propose an algorithm to ensure privacy and in particular avoid collusion attacks in gossip based learning.

The basic idea to use gossip-based learning is very interesting and appealing. Ensuring privacy in a large distributed system is therefore a relevant and important contribution.

Besides providing a thorough analysis of the algorithms properties, the authors also try to give characteristics of the behavior in real system environment. This part reads less convincing. For example, the overhead is measured in terms of communication cost. In the context of mobile devices (Figure 3) it would be much more important to understand the delays and overhead imposed by processing steps of the deployed paradigm. To my understanding the used cryptographic primitives come at a significant cost.

Nevertheless, overall the presented approach offers many interesting insights and introduces a promising approach.