Modern Dynamic Data Structures

Monika Henzinger
University of Vienna, Department of Computer Science, Vienna, Austria

Abstract

We give an overview of differentially private dynamic data structure, aka differentially private algorithms under continual release.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Differential privacy, data structures

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.2

Category Invited Talk

Funding This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101019564 "The Design of Modern Fully Dynamic Data Structures (MoDynStruct)" and from the Austrian Science Fund (FWF) project "Fast Algorithms for a Reactive Network Layer (ReactNet)", P 33775-N, with additional funding from the *netidee SCIENCE Stiftung*, 2020–2024

1 Extended Abstract

In the past the main goal when designing data structures was to achieve optimal time per operation and optimal space. However, in recent years new applications have lead to new requirements for data structures, such as differential privacy or fairness.

In this talk I will limit myself to the first topic, namely differential privacy. Since its invention in 2006 by Dwork, McSherry, Nissim, and Smith [3] differentially private algorithms and (static) data structures have been designed for many combinatorial problems (see e.g. [5] for a book on the topic). However, very little work has been done for *dynamic* data structures. A *dynamic data structure* is a data structure that supports not only *query* operations to the stored data, but also *update* operations, such as insertions and/or deletions. In the differentially privacy research community such data structures are frequently called data structures *in the continual release* (or *continual observation*) *model.*

The problem of *binary counting* the number of 1s in a binary sequence is equivalent to a dynamic data structure that supports the *AppendBit* operation and that outputs the number of 1s in the current sequence after each *AppendBit* operation. This problem has been well-studied in the differentially private setting [4, 1, 8, 2, 7], including a version that outputs a weighted average of the bits in the sequence so far [7]. Another extension of this problem leads to the *MaxSum* and the *SumSelect* problem: Assume the input is a sequence of d-dimensional binary vectors such that the t-th vector is denoted by b_t. The goal of the *SumSelect* problem is to output the value of the coordinate i such that $i = \text{argmax}_i \sum_t b_t[i]$, the goal of *MaxSum* is to output $\max_i \sum_t b_t[i]$. These problems were studied in the continual release model in [9]: Differentially private partially *dynamic graph algorithms* were also analyzed in the continual release setting [10, 6].

We will describe the algorithm of [7] in detail and explain why it is superior to the previous solutions for binary counting under continual observation.

© Monika Henzinger;
licensed under Creative Commons License CC-BY 4.0
47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).
Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 2; pp. 2:1–2:2
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

References

1 T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Trans. Inf. Syst. Secur.*, 14(3):26:1–26:24, 2011. doi:10.1145/2043621.2043626.

2 Sergey Denisov, Brendan McMahan, Keith Rush, Adam Smith, and Abhradeep Thakurta. Improved differential privacy for sgd via optimal private linear operators on adaptive streams. *arXiv preprint*, 2022. arXiv:2202.08312.

3 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*, pages 265–284, 2006.

4 Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *Proc. of the Forty-Second ACM Symp. on Theory of Computing (STOC'10)*, pages 715–724, 2010.

5 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. doi:10.1561/0400000042.

6 Hendrik Fichtenberger, Monika Henzinger, and Wolfgang Ost. Differentially private algorithms for graphs under continual observation. In *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, 2021.

7 Hendrik Fichtenberger, Monika Henzinger, and Jalaj Upadhyay. Constant matters: Fine-grained complexity of differentially private continual observation using completely bounded norms. *arXiv preprint*, 2022. arXiv:2202.11205.

8 James Honaker. Efficient use of differentially private binary trees. *Theory and Practice of Differential Privacy (TPDP 2015), London, UK*, 2015.

9 Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam Smith. The price of differential privacy under continual observation. *arXiv preprint*, 2021. arXiv:2112.00828.

10 Shuang Song, Susan Little, Sanjay Mehta, Staal Vinterbo, and Kamalika Chaudhuri. Differentially private continual release of graph statistics. *arXiv preprint*, 2018. arXiv:1809.02575.