

1145-94-553

Ankur Mallick*, 4720 Forbes Avenue, CIC 4th floor, Pittsburgh, PA 15213, and **Malhar Chaudhari** and **Gauri D Joshi**, 4720 Forbes Avenue, CIC 4105, Pittsburgh, PA 15213. *Fast and Efficient Distributed Matrix-Vector Multiplication Using Rateless Fountain Codes.*

Large-scale machine learning and data mining applications require computer systems to perform massive computations that need to be parallelized across multiple nodes, for example, massive matrix-vector and matrix-matrix multiplication. The presence of straggling nodes – computing nodes that unpredictably slowdown or fail – is a major bottleneck in such distributed computations. We propose a *rateless fountain coding* strategy to alleviate the problem of stragglers in distributed matrix-vector multiplication. Our algorithm generates linear combinations of the m rows of the matrix and assigns them to different worker nodes, which then perform row-vector products with the encoded rows. The original matrix-vector product can be decoded as soon as slightly more than m row-vector products are collectively completed by the nodes. This strategy enables fast nodes to steal work from slow nodes, without requiring the knowledge of node speeds. Compared to recently proposed fixed-rate erasure coding strategies which ignore partial work done by straggling nodes, rateless coding achieves significantly lower overall delay, as well as small computational overhead. (Received September 09, 2018)