


Communication-optimal algorithms for contracting distributed tensors

29 July 2014



Iteration Space and Data Space Mapping for Matrix-Matrix Multiplication on a 2D Torus Network

Tensor contractions, generalized matrix multiplications that are time-consuming to calculate, make them among the most compute-intensive operations in several ab initio computational quantum chemistry methods. In this work, scientists from Pacific Northwest National Laboratory and The Ohio State University developed a systematic framework that uses three fundamental communication operators—recursive broadcast, rotation, and reduction, or RRR,—to derive communication-efficient algorithms for distributed contraction of arbitrary dimensional tensors on the IBM Blue Gene/Q Mira supercomputer. The framework automatically models potential space-performance trade-offs to optimize the communication costs incurred in executing tensor contractions on supercomputers. The paper documenting this work, "Communication-optimal Framework for Contracting Distributed Tensors," is a SC14 Best Paper award finalist.

In computational physics and chemistry, [tensor algebra](#) is important because it provides a mathematical framework for formulating and solving problems related to areas such as fluid mechanics. By offering a comprehensive framework that automatically generates communication-optimal algorithms for contracting distributed tensors, redundancy is avoided and the total computation load is balanced, improving the

overall communication costs. By deconstructing these distributed tensor contractions, the work also afforded insights into the fundamental building blocks of these widely studied computations.

The researchers characterized distributed tensor contraction algorithms on tori (mesh circles with wraparound connected in more than one dimension) networks, defining tensor indices, iteration space, and their mappings. By mapping the iteration space, they could precisely define where each computation of a tensor contraction occurs, as well as define the data that needs to be present in each processor. For each tensor contraction, the researchers sought an iteration space mapping, a data space mapping, and an algorithm to minimize the communication cost (per contraction) for a given amount of memory per processor.

Then, for a given iteration space mapping, their RRR framework identified the fundamental data movement directions required by a distributed algorithm, which also are elemental to the tensor contraction, called "reuse dimensions." With these reuse dimensions, the framework can compute compatible input and output tensor distributions and systematically generate a contraction algorithm for them using communication operators. In their work, the researchers also showed a cost model that predicted the communication cost for a given iteration space mapping, compatible input and output distribution, and the generated contraction algorithm. The cost model then was used to identify iteration and data space mapping that minimized the overall communication cost.

In their experiments, the researchers showed their framework was scalable up to 16,384 nodes (262,144 cores) on Blue Gene/Q supercomputers. They also demonstrated how their framework improves commutation optimality—even exceeding the Cyclops Tensor Framework, which stands as the current state of the art.

In addition to their distributed and symmetric nature, tensors also might exhibit various forms of sparsity. The team is working on combining this work with the approach published in "A Framework for Load Balancing of Tensor Contraction Expressions via Dynamic Task Partitioning," presented last year at SC13, to dynamically load balance tensor contractions. The outcome would be a hybrid approach that exploits the [communication](#) efficiency of this work while dynamically adapting to load imbalances introduced by sparsity.

More information: Rajbhandari S, A Nikam, P Lai, K Stock, S Krishnamoorthy, and P Sadayappan. 2014. "Communication-optimal framework for contracting distributed tensors." Presented at: International Conference for High Performance Computing, Networking, Storage and Analysis (SC14). November 16-21, 2014, New Orleans, Louisiana (Best Paper Finalist).

Lai P, K Stock, S Rajbhandari, S Krishnamoorthy, and P Sadayappan. 2013. "A framework for load balancing of tensor contraction expressions via dynamic task partitioning." Presented at: International Conference for High Performance Computing, Networking, Storage and Analysis (SC13). November 17-22, 2013, Denver, Colorado.

Provided by Pacific Northwest National Laboratory

APA citation: Communication-optimal algorithms for contracting distributed tensors (2014, July 29) retrieved 28 October 2020 from <https://phys.org/news/2014-07-communication-optimal-algorithms-tensors.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.