ABSTRACT

Efficient computation and communication management for all-pairs interactions

Cory James Kleinheksel

Arun K. Somani, Major Professor
Iowa State University

Big data continues to grow in size for all sciences. New methods like those proposed are needed to further reduce memory footprints and distribute work equally across compute nodes both in local HPC systems and rented cluster resources in the cloud. Modern infrastructures have evolved to support these big data computations and that includes key pieces like our internet backbones and data center networks. Many optical networks face heterogeneous communication requests requiring topologies to be efficient and fault tolerant. The all-pairs problem requires all elements (computation datasets or communication nodes) to be paired with all other elements. These all-pairs problems occur in many research fields and have significant impacts, which has led to their continued interest.

We proposed using cyclic quorum sets to efficiently manage all-pairs computations. We proved these sets have an "all-pairs" property that allows for minimal data replication and for distributed, load balanced, and communication-less computation management. The quorums are $O\left(\frac{N}{\sqrt{P}}\right)$ in size, up to 50% smaller than dual $\frac{N}{\sqrt{P}}$ array implementations, and significantly smaller than solutions requiring all data. Scaling from 16 to 512 cores (1 to 32 compute nodes) and using real dataset inputs, application experiments demonstrated scalability with greater than 150x (super-linear) speedup and less than 1/4th the memory usage per process.

Cyclic quorum sets can provided benefits to more than just computations. The sets can provide a guarantee that all pairs of optical nodes in a network can communicate also. Our evaluation analyzed the fault tolerance of routing optical cycles based on cyclic quorum sets. With this method of topology construction, unicast and multicast communication requests do not need to be known or even modeled a priori. In the presence of network single-link faults, our simulated cycle routing had greater than 99% average fault coverage. Hence, even in the presence of a network fault, the optical networks could continue operation of nearly all node pair communications.

Lastly, we proposed a generalized R redundant cyclic quorum set. These sets guarantee all pairs of nodes occur at least R times. When applied to routing cycles in optical networks, this technique provided almost fault-tolerant communications. More importantly, when applied using only single cycles rather than the standard paired cycles, the generalized R redundancy technique almost halved the necessary light-trail resources while maintaining the fault tolerance and dependability expected from cycle-based routing.