LA-UR-24-28102

Author(s): Jackson Wesley

Mentors: David DeBonis

Title: Communication Performance Assessment of Sapphire Rapids Architecture

Abstract:

The Sapphire Rapids architecture from Intel is an approach to microarchitecture design using chiplets, discrete and modular chips assembled into a package. This new chiplet design offers powerful compute as well as cost-effective customization options and is the foundation of the Cray Shasta architecture used in modern supercomputer designs such as Crossroads, Tycho, and Rocinante at LANL. Because communication is the dominating cost in parallel applications, it is valuable to understand the on node and off node bandwidth and latency performance of these microarchitectures. In this study we present the bandwidth and latency trends over a range of message sizes for on node chiplets in on-socket and off-socket communication using Fabtests, a Libfabric testing package. Additionally, we measure on node and off node communication and compare to MPI communication measurements using the OSU Benchmarks library. We observe differences in bandwidth performance at certain packet sizes that were unanticipated, occurring at small and very large message sizes. This study provides an assessment of the communication performance of Sapphire Rapids architectures allowing for communication improvements in HPC codes of interest.