


A Versatile NDP Server Architecture and Its Impact on Data Center Applications



Existing near-data processing (NDP) techniques have demonstrated their strength for some specific data-intensive applications. However, they may be inadequate for a data center server that needs to perform a diverse range of applications from data-intensive to compute-intensive. How to develop a

versatile NDP-powered server to support various data center applications remains an open question. Further, a good understanding of the impact of NDP on data center applications is still missing. For example, can a compute-intensive application also benefit from NDP? Which type of NDP engine is a better choice, an FPGA-based engine or an ARM-based engine? To address these issues, we first propose a new NDP server architecture that is equipped with a parallel NDP engine array. Based on the architecture, two NDP servers ANS (ARM-based NDP Server) and FNS (FPGA-based NDP Server) are introduced. Next, we implement a single-engine prototype for each of them. Finally, we measure performance, energy efficiency, and cost/performance ratio of six typical data center applications running on the two prototypes. Some new findings have been observed.

Xiaojia Song, Tao Xie, and Stephen Fischer

This research is supported by the National Science Foundation (CNS-1320738) grant, and the Computational Science Research Center (CSRC) at San Diego State University

