EDITOR **RON VETTER**
University of North Carolina Wilmington;
vetterr@uncw.edu

Dapper: An Adaptive Manager for Large-Capacity Persistent Memory

Dongliang Xue, Linpeng Huang, Chao Li, and Chentao Wu, Shanghai Jiao Tong University

This installment of Computer's *series highlighting the work published in IEEE Computer Society journals comes from* IEEE Transactions on Computers.

Smart applications are becoming more data centric, and the need for businesses to quickly manage vast amounts of data continues to challenge today's computing infrastructure. Nevertheless, main memory systems consisting entirely of dynamic random-access memory (DRAM) are already hitting their power and capacity limits. It would be beneficial to unleash significant system performance combined emerging memory modules without being worried about design complexities. Design innovations in highly scalable hybrid memory systems hold great promise to realize this vision.

A recent paper[1] proposes an adaptive management mechanism called *Dapper* (see Figure 1) for large-capacity memory integration. From the start of its design, Dapper has given high priority to lower maintenance costs, better user experience, and higher scalability. It enables the user to build larger-capacity memory systems while remaining within the system's overall cost and resource limits.

This article focuses on persistent memory (PM), which is a recent groundbreaking technological innovation. PM delivers a unique combination of data persistence with an unprecedented large capacity. For example, Intel has issued its new Optane dc PM module (DCPMM), available in 128-, 256-, and 512-GB modules. However, integrating PM into current commercial servers can be a great undertaking. Many designers and developers may decide against changing to PM solutions because of challenging tasks such as metadata management, power consumption issue,

Digital Object Identifier 10.1109/MC.2020.2964916
Date of current version: 12 March 2020

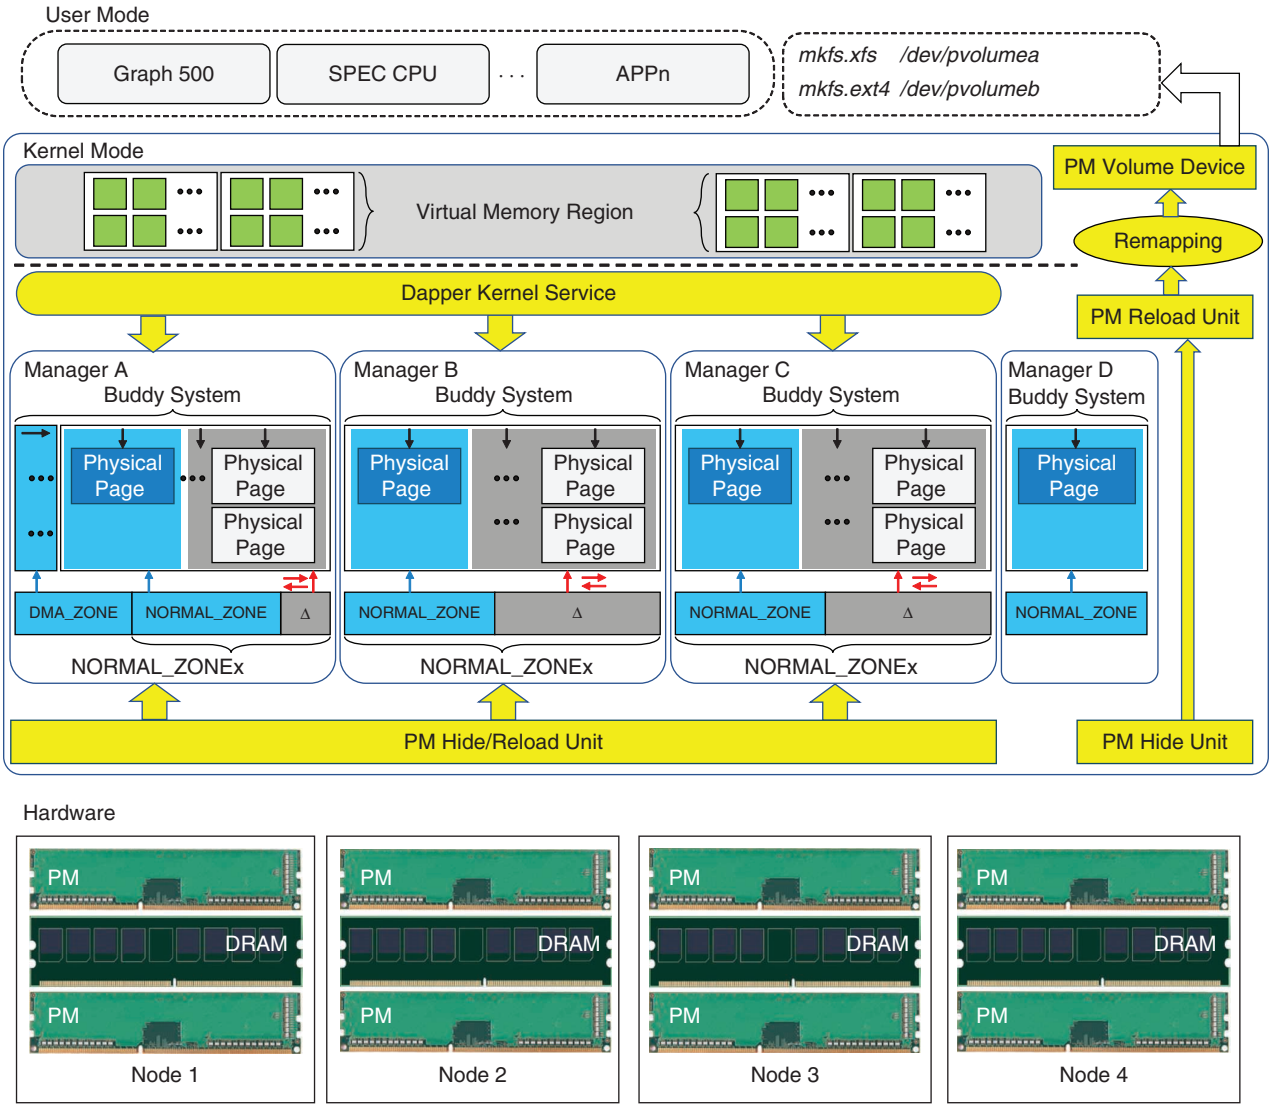

FIGURE 1. The Dapper architecture overview.

and remote memory access latency. Dapper's main advantage is that it gracefully addresses these memory-scalability challenges by balancing performance and overhead.

Dapper provides a convenient approach for an enterprise to introduce large-capacity PM into nonuniform memory access-based commodity servers. Building with scalability and simplicity in mind, Dapper allows servers to hold increasing amounts of data using a multifaceted optimization strategy. For example, Dapper manages PM using a DRAM-like manner, thereby bypassing the vast modifications to the memory management subsystem of the OS. In addition, considering the expense of managing large-capacity PM, Dapper chooses to dynamically integrate a moderate amount of memory to gradually match the memory demand. Confronted with the bottleneck of remote memory access, Dapper further reduces its reliance on remote memory access by expanding local memory with a moderate amount of PM space.

Aside from enabling more memory and larger data volumes per server, the authors also infuse the attribute of persistency into the storage subsystem of the OS by converting part of the PM into a volume device. Experimental results show that compared to typical memory management approaches, Dapper can achieve

average performance improvements of 13 and 34% on Graph500 BFS_SSSP benchmarks and SPEC CPU2006 floating point workloads, respectively.

It is increasingly important to build scalable, efficient, and high-performance computers. Systems such as Dapper prepare enterprise systems for future growth while creating a much leaner product that suits current in-memory computing needs without extra design complexity. For more than 65 years, *IEEE Transactions on Computers* has served the computing community with top-quality research contributions. It will continually pay attention to the state-of-the-art computing system designs from both academia and industry.

DONGLIANG XUE is an assistant research fellow in the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. Contact him at xuedongliang010@sjtu.edu.cn.

LINPENG HUANG is a professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. Contact him at lphuang@sjtu.edu.cn.

CHAO LI is an associate professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. Contact him at chaol@sjtu.edu.cn.

CHENTAO WU is a professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. Contact him at wuct@cs.sjtu.edu.cn.

REFERENCE

1. D. Xue, L. Huang, C. Li, and C. Wu, "Dapper: An adaptive manager for large-capacity persistent memory," *IEEE Trans. Comput.*, vol. 68, no. 7, pp. 1019–1034, July 2019. doi: 10.1109/TC.2019.2899836.

Digital Object Identifier 10.1109/MC.2020.2974641