

## PicoDriver: Fast-path Device Drivers for Multi-kernel Operating Systems

Balazs Gerofi  
RIKEN R-CCS  
bgerofi@riken.jp

Aram Santogidis  
CERN, Switzerland  
aram.santogidis@cern.ch

Dominique Martinet  
CEA, France  
dominique.martinet@cea.fr

Yutaka Ishikawa  
RIKEN R-CCS  
yutaka.ishikawa@riken.jp

出典 : HPDC '18: International Symposium on High-Performance Parallel and Distributed Computing

Lightweight kernel (LWK) operating systems (OS) in high-end supercomputing have a proven track record of excellent scalability. However, the lack of full Linux compatibility and limited availability of device drivers in LWKs have prohibited their widespread deployment. Multi-kernels, where an LWK is run side-by-side with Linux on many-core CPUs, have been proposed to address these shortcomings. In a multi-kernel system the LWK implements only performance critical kernel services and the rest of the OS functionality is offloaded to Linux. Access to device drivers is usually attained via offloading. Although high-performance interconnects are commonly driven from user-space, there are networks (e.g., Intel's OmniPath or Cray's Gemini) that require device driver interaction for a number of performance sensitive operations, which in turn can be adversely impacted by system call offloading.

In this paper, we propose PicoDriver, a novel device driver architecture for multi-kernels, where only a small part of the driver (i.e., the performance critical piece) is ported to the LWK and access to the rest remains transparent via Linux. Our solution requires no modifications to the original Linux driver, yet it enables optimization opportunities in the lightweight kernel. We implemented this system in the IHK/McKernel multi-kernel OS and demonstrate that on 16,384 Intel Knight's Landing (KNL) CPU cores (interconnected by OmniPath network) we can outperform Linux by up to 30% on various mini-applications.