Yitian YANG, Youyou LU, 2023. NICFS: a file system based on persistent memory and SmartNIC. *Frontiers of Information Technology & Electronic Engineering*, 24(5): 675-687. https://doi.org/10.1631/FITEE.2200469

NICFS: a file system based on persistent memory and SmartNIC

Key words: Non-volatile memory; Persistent memory; Data processing unit; Smart network interface card (SmartNIC); File system

Corresponding author: Youyou LU
E-mail: luyouyou@tsinghua.edu.cn
ORCID: https://orcid.org/0000-0002-6214-5390

Motivation

Emergence of new hardware, including persistent memory and smart network interface card (SmartNIC), has brought new opportunities to file system design:

- **Persistent memory** offers high performance and byte addressability, making it ideal for storing logs and metadata.
- **SmartNIC** has the ability to offload host CPU tasks, thus saving host CPU cycles to perform computational tasks and gaining greater economic benefits.

With these considerations, we designed NICFS, a file system based on persistent memory and SmartNIC.

Main idea

1. NICFS uses persistent memory to store logs and metadata to take advantage of its low latency and granularity. For simplicity, the data are also stored on the persistent memory.

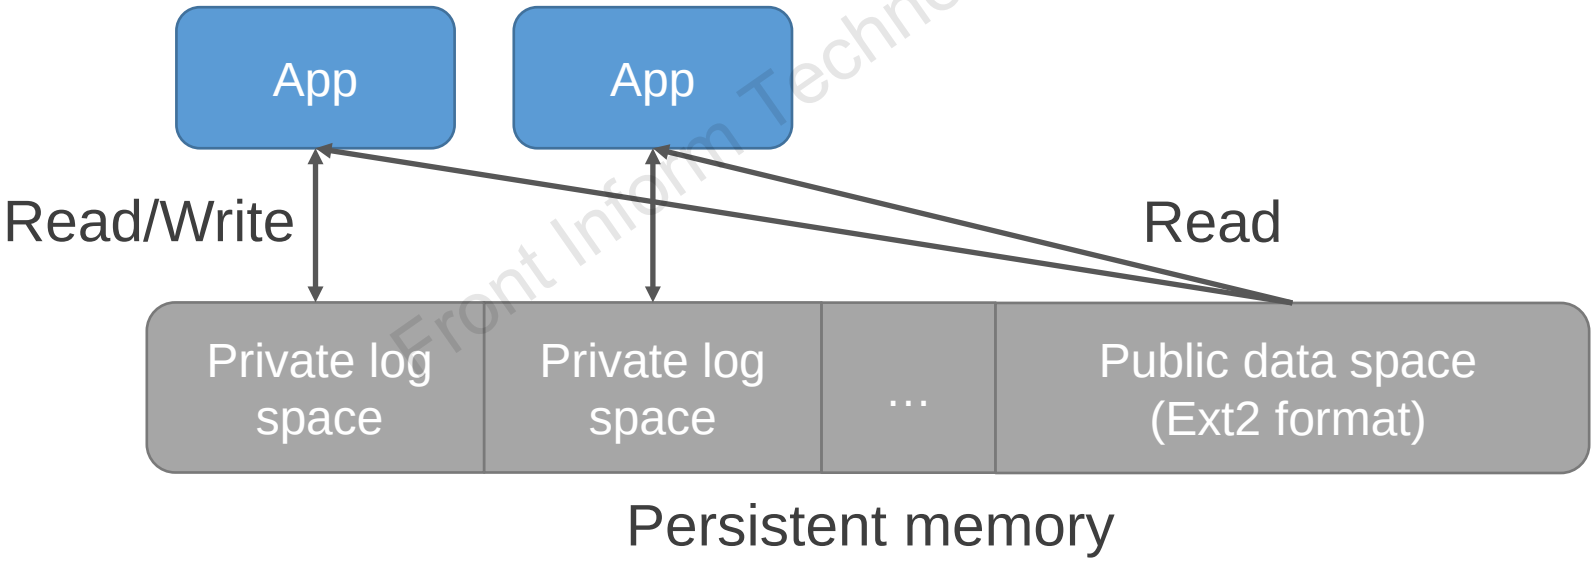

Main idea

2. NICFS uses SmartNIC to asynchronously patch logs to public data space, reducing write latency and reducing host CPU load. Meanwhile, this divides NICFS into two parts, NICFS-Frontend and NICFS-Backend.

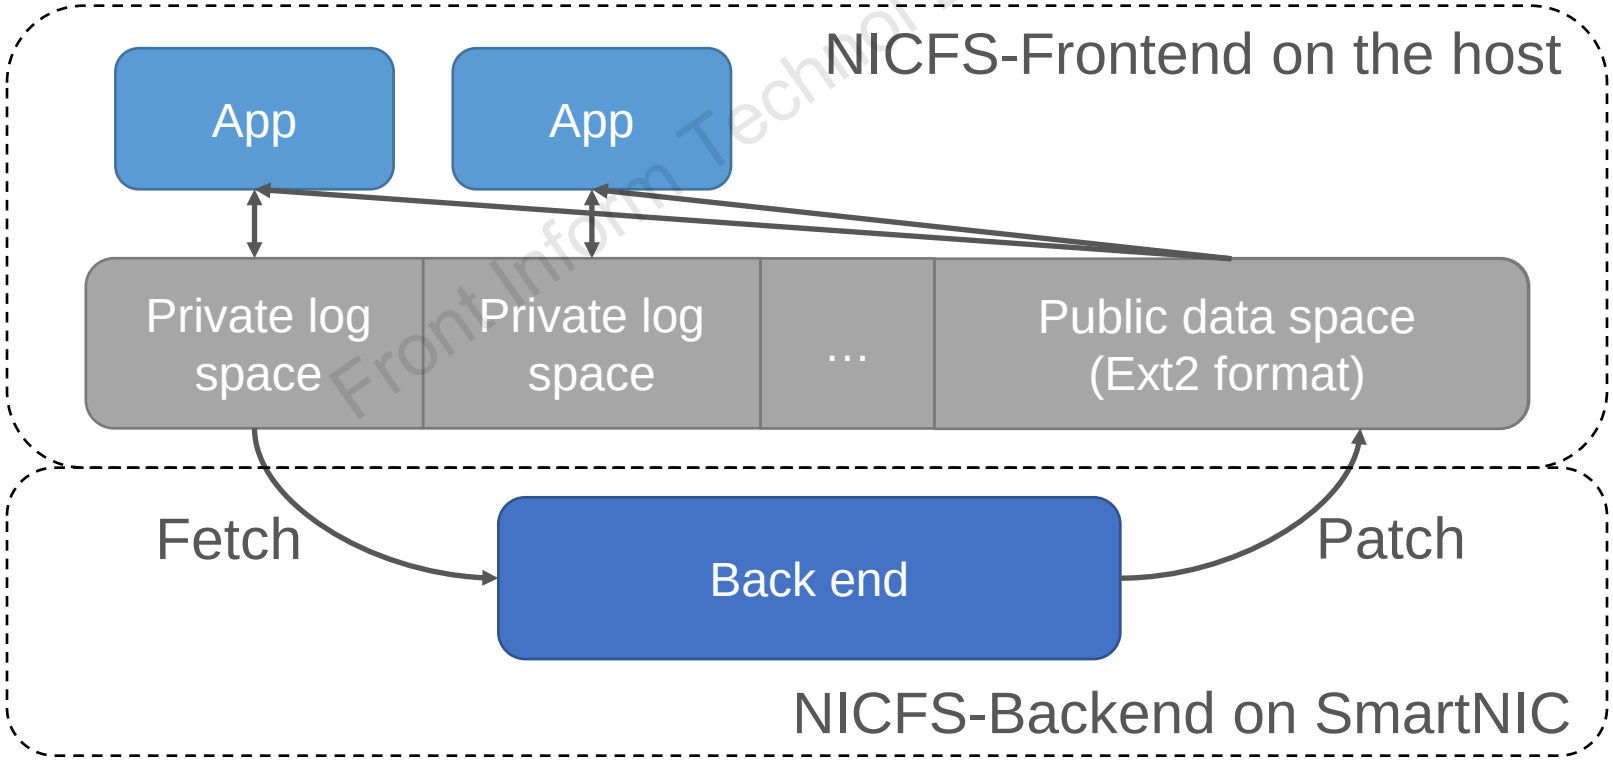

Main idea

3. To compensate for the poor performance of SmartNIC cores and enhance file system scalability, NICFS-Backend parallelizes log processing in pipeline and adopts efficient log merging and metadata caching strategies.

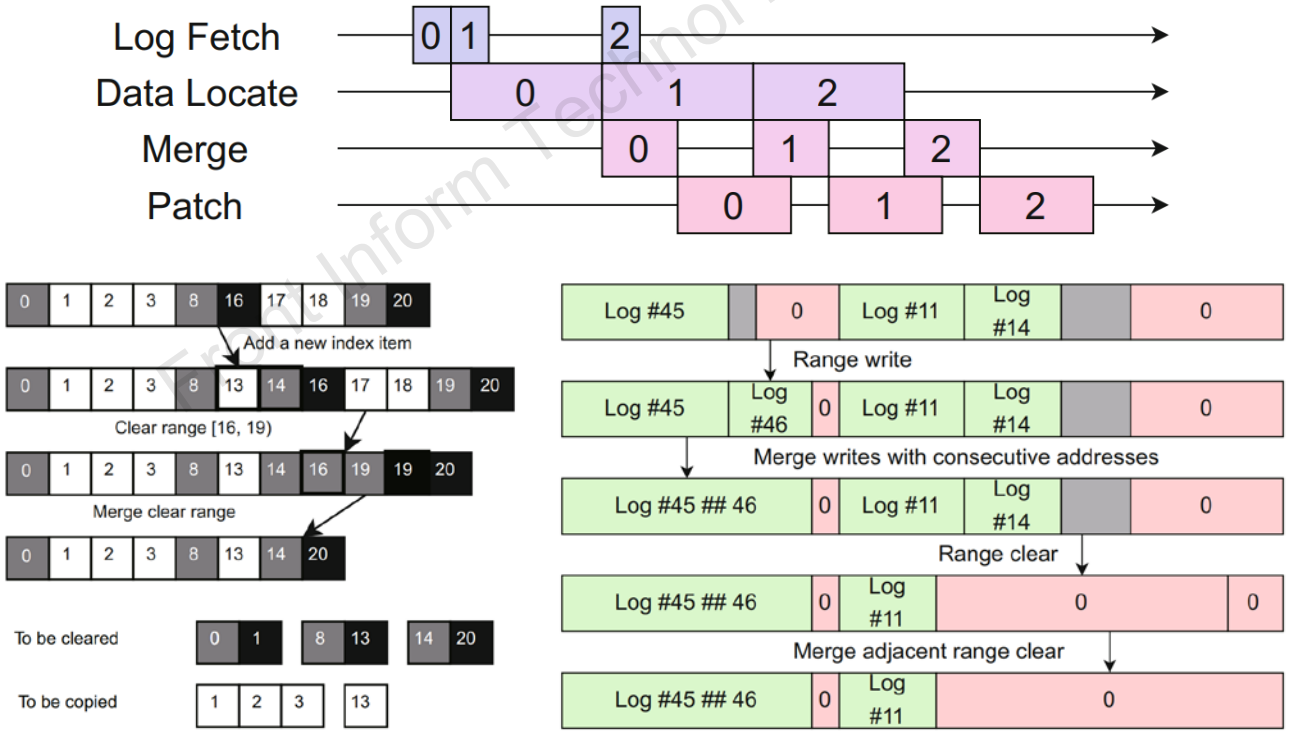

Overall architecture

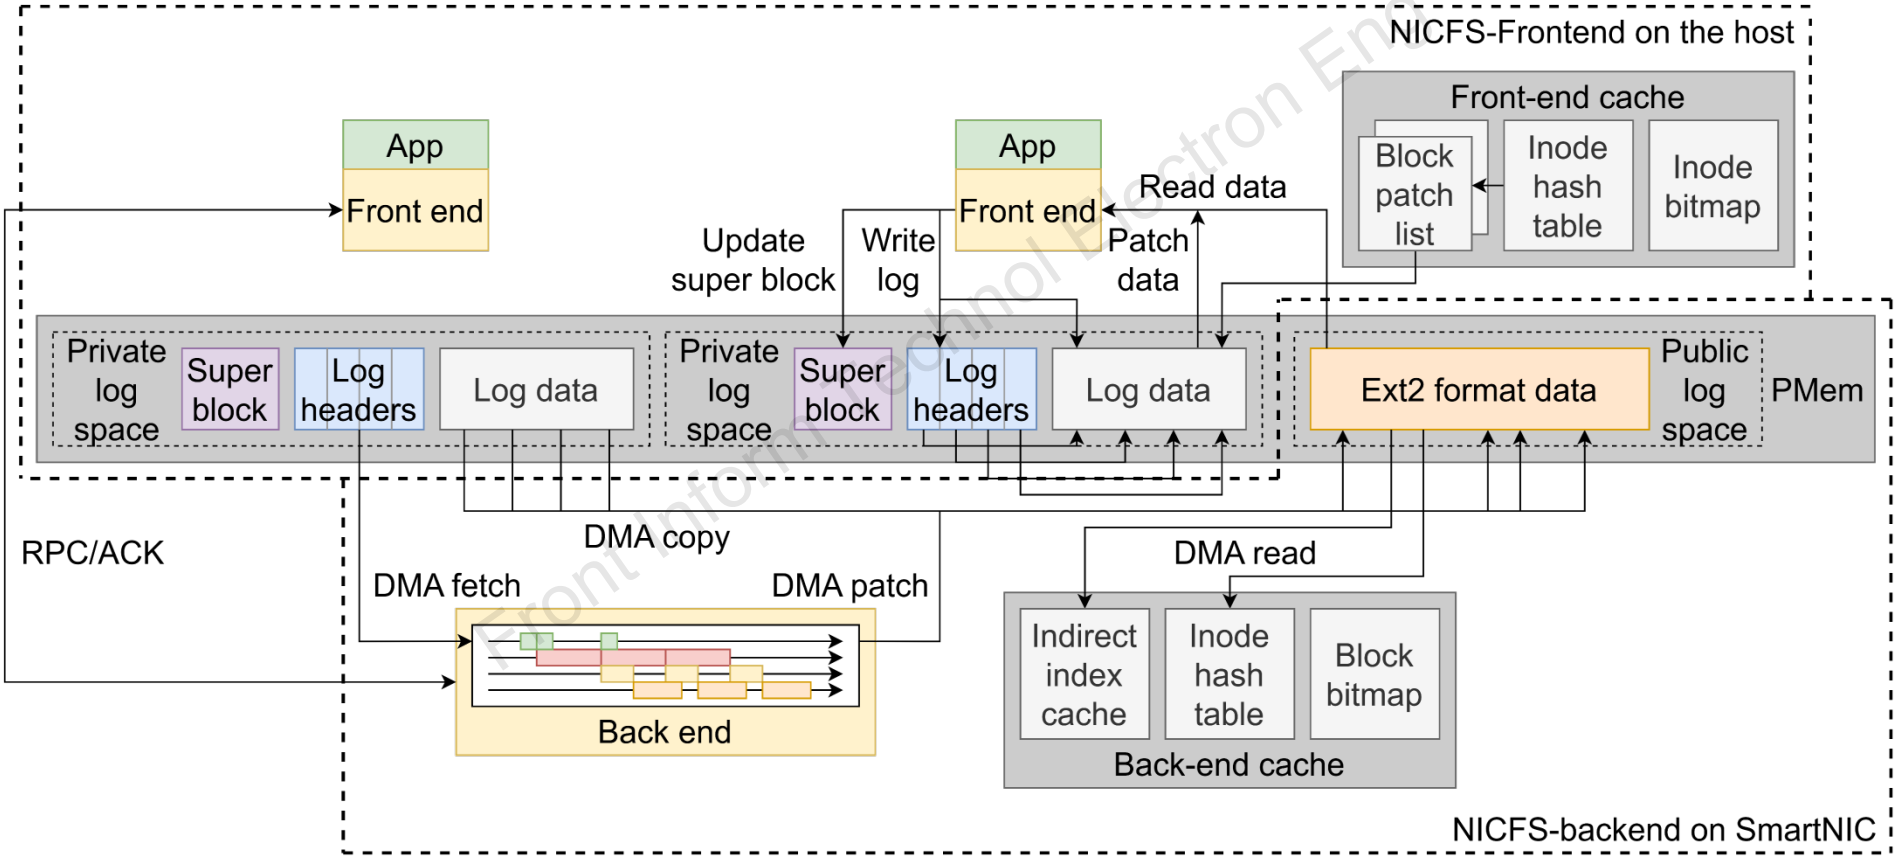

Method

The experiments were performed on a single machine with the following configurations. We compared the performance differences between NICFS and the classic Ext4 file system for sequential and random reads and writes at different granularities, and verified the scalability of NICFS and the effectiveness of each design.

Unit	Description
OS	Ubuntu 20.04.1 LTS with Linux kernel version 5.4.0-26-generic
CPU	Intel Xeon Gold 6240 at 2.60 GHz with 72 cores
Memory	Samsung M393A4K40DB2-CVF 64 GB DDR4-2993 DRAM
PMem	Intel Optane persistent memory 100 series 128 GB module
SmartNIC	NVIDIA MBF2H516A-EENOT BlueField-2 DPU with 8 Cortex-A72 cores and 16 GB DDR4-3200 DRAM

OS, operating system; CPU, central processing unit; DDR, double data rate; DPU, data processing unit; DRAM, dynamic random access memory; LTS, long-term support; PMem, persistent memory

Major results

Evaluation results showed that NICFS outperformed Ext4 by about 21%/10% and about 19%/50% on large and small reads/ writes, respectively.

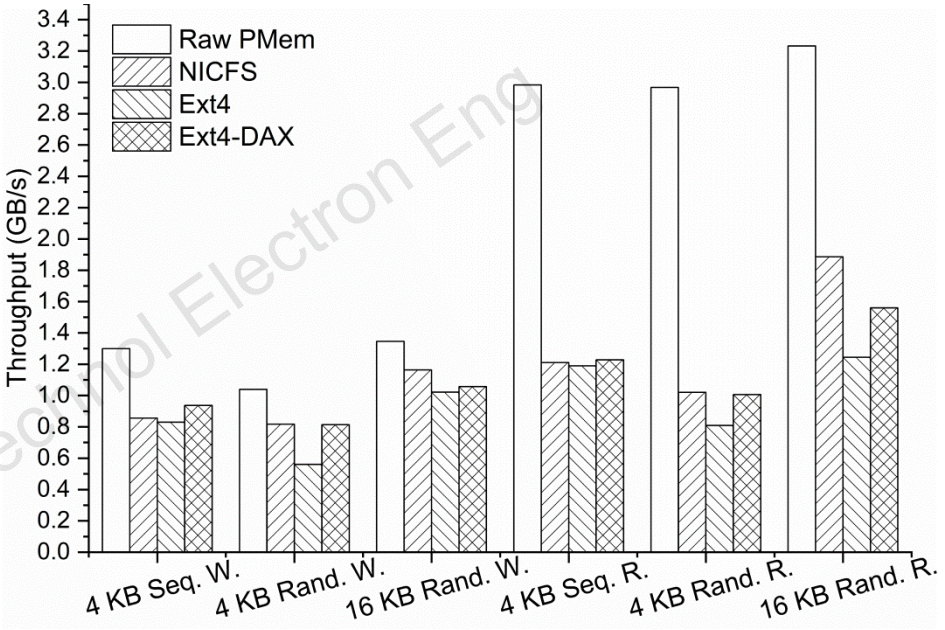

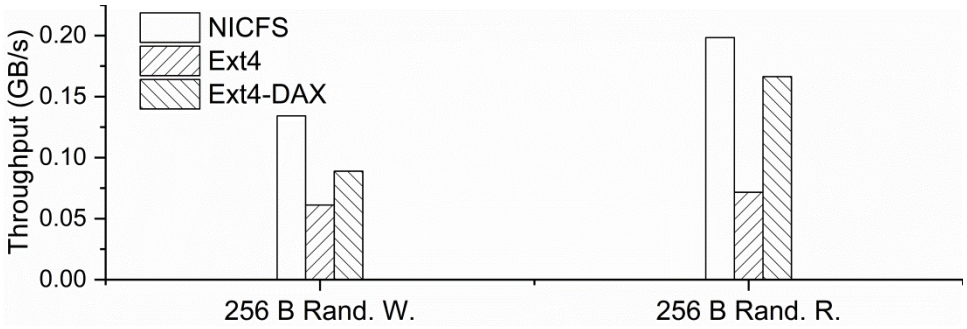

Conclusions

- We proposed NICFS, a file system based on persistent memory for persistent storage and SmartNIC for offloading file processing.
- A set of data structures and algorithms were designed for logging, caching, merging, processing, and parallelizing suitable for NICFS.
- The performance of NICFS was slightly better than that of Ext4 with kernel page cache on regular reads/writes, and significantly better than that of Ext4 on small data reads/writes.
- The processing power of SmartNIC should match the performance of persistent memory; otherwise, the weak core performance of SmartNIC may become a hindrance to scalability.