

1. Record Nr.	UNINA9911049220303321
Autore	Yu Hua
Titolo	Big Memory Systems // by Yu Hua
Pubbl/distr/stampa	Singapore : , : Springer Nature Singapore : , : Imprint : Springer, , 2026
ISBN	981-9528-85-2
Edizione	[1st ed. 2026.]
Descrizione fisica	1 online resource (297 pages)
Collana	Computer Science Series
Disciplina	004.53
Soggetti	Computer storage devices Memory management (Computer science) Information storage and retrieval systems Computer systems Computer Memory Structure Information Storage and Retrieval Computer System Implementation
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	"1. Write Optimized and High Performance Persistence in Big Memory" -- "2. Lock free Concurrent Level Design for Persistent Memory" -- "3. GPU enabled Byte Granularity Persistence for Big Memory" -- "4. Scalable Learned Key Value Store for Disaggregated Memory" -- "5. Fast One sided RDMA based Transactions for Disaggregated Memory" -- "6. Multi Versioning Design for Distributed Transactions on Big Memory" -- "7. Fast and Cost Efficient Hashing Index Schemes for Cloud Systems" -- "8. Mitigating Asymmetric Read and Write Costs in Cuckoo based Designs".
Sommario/riassunto	This book presents groundbreaking research schemes that are shaping the future of big memory systems. Large-scale persistent and byte-addressable memory eliminates traditional storage bottlenecks by unifying memory and storage hierarchies. For engineers and architects building next-generation databases and distributed systems, these innovations deliver transformative performance: faster persistent writes through optimized flush mechanisms, the reduced contention via lock-free designs, and high throughput using GPU acceleration. Each chapter presents fully implemented architectures in real-world environments,

from learned indexes to RDMA-powered transactions, together with deployable code patterns and performance benchmarks. This book shows new devices are leveraging big memory to achieve low-latency persistence, hardware-accelerated data processing, and linear scalability in distributed environments. Designed for professionals with operating systems fundamentals, this book bridges cutting-edge research with practical implementation, where big memory's unique characteristics (persistence at DRAM speeds, massive capacity, and fine-grained access) demand fundamentally new architectural approaches. Learn how to achieve faster queries with learned indexes in the disaggregated memory, how to optimize cuckoo hashing for persistent memory's asymmetric costs, and why the latest GPUs incorporate these persistence techniques. This book also provides efficient and useful toolkits: the RDMA protocols have been adopted in storage tiers, while the lock-free designs improve real-time recommendation systems. Whether building cloud-native databases, low-latency recommendation systems, or memory-driven AI services, these solutions will help exploit the full potential of the big memory.
