

Yacine Taleb - 3rd Year Ph.D Student @Inria

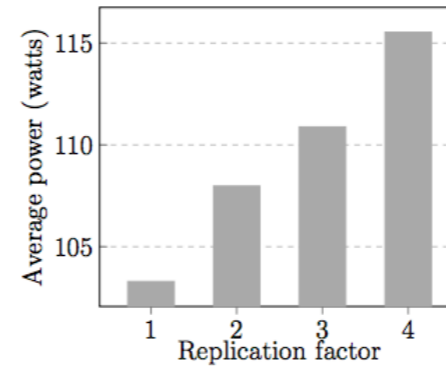

Advisors: Gabriel Antoniu (Inria) and Toni Cortes (BSC)

Efficient Replication and Fault-tolerance

Key-words: Systems, Replication, Fault-tolerance, Efficiency, In-memory, Stream Storage



Ph.D research summary

Replication is expensive in in-memory systems




Ph.D research summary

Replication is expensive in in-memory systems



ICDCS'17



Why ?
Network, **CPU**, ...



E.g. RAMCloud, FaRM, Redis


Ph.D research summary

Replication is expensive in in-memory systems




ICDCS'17

Why ?
Network, **CPU**, ...





E.g. RAMCloud, FaRM, Redis

It could benefit from RDMA




Ph.D research summary

Replication is expensive in in-memory systems



ICDCS'17

Why ?
Network, **CPU**, ...



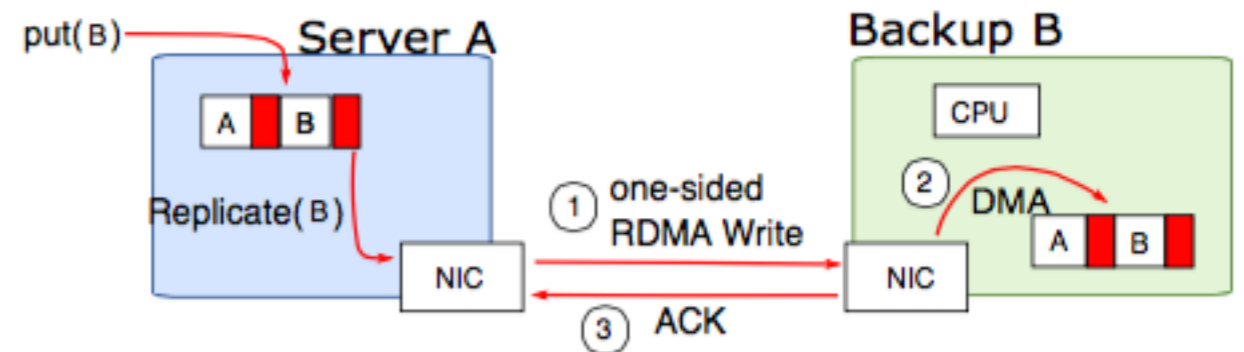
E.g. RAMCloud, FaRM, Redis

It could benefit from RDMA

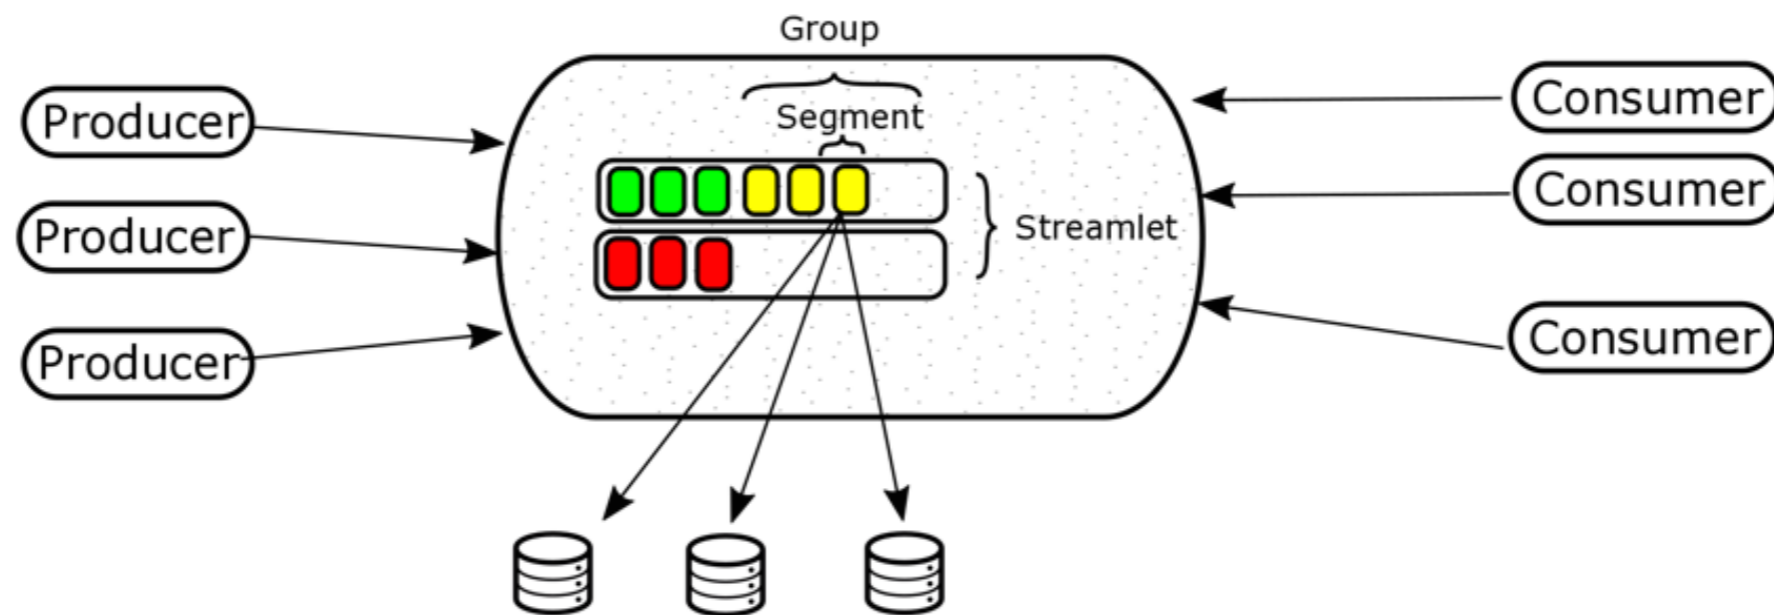
But... **unsafe!**

Tailwind guarantees atomicity

ATC'18



KerA: High-throughput and Durable Stream Store

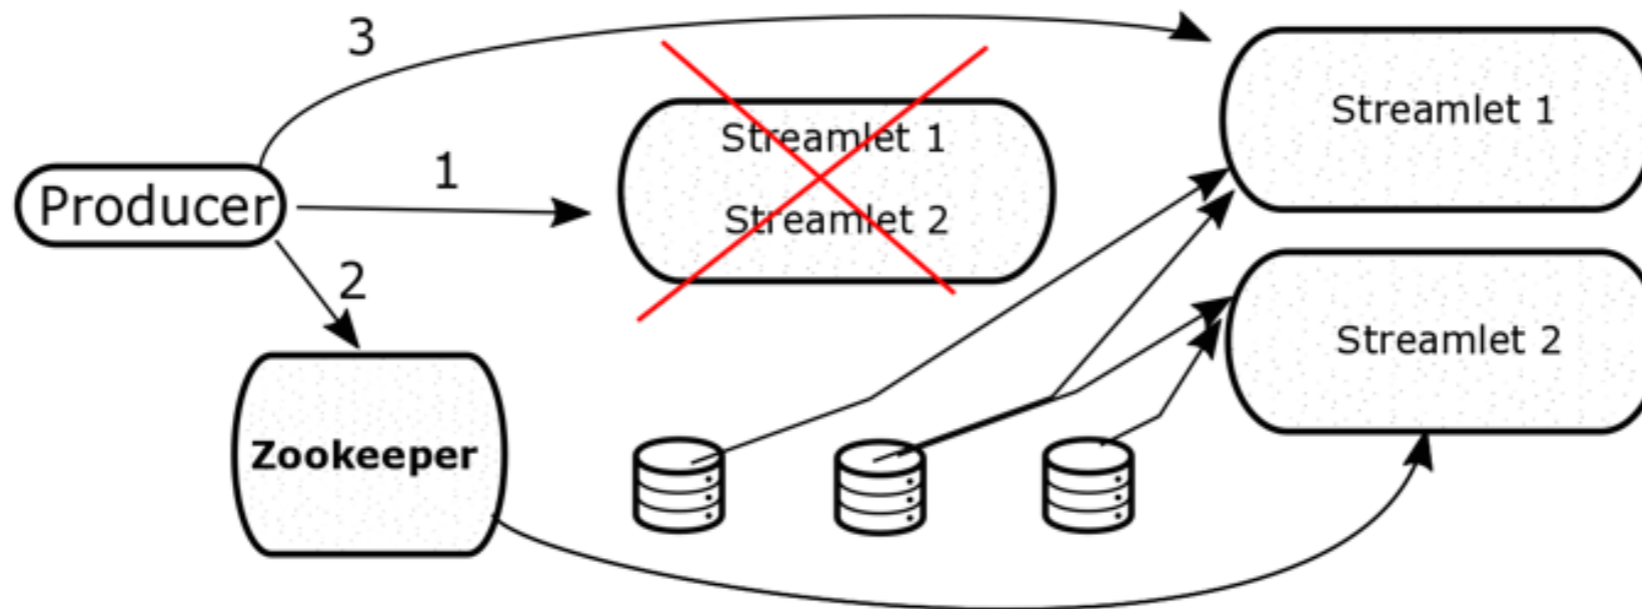


A new stream-partitioning scheme -> **Scale**-up/down

Fine grain -> better **throughput**

Co-locating stream processing and storage (kv-store)

Durability in KerA



KerA is durable and guarantees Exactly-once!

Crash? Needs 2 RTTs to resume service

Early results: **5X** better throughput than Kafka (Producers/Consumers)
10X with Kernel-bypass