

TAS: TCP Acceleration as an OS Service

Antoine Kaufmann
MPI-SWS

Timothy Stamler, Simon Peter
The University of Texas at Austin

Naveen Sharma, Thomas Anderson,
Arvind Krishnamurthy
University of Washington

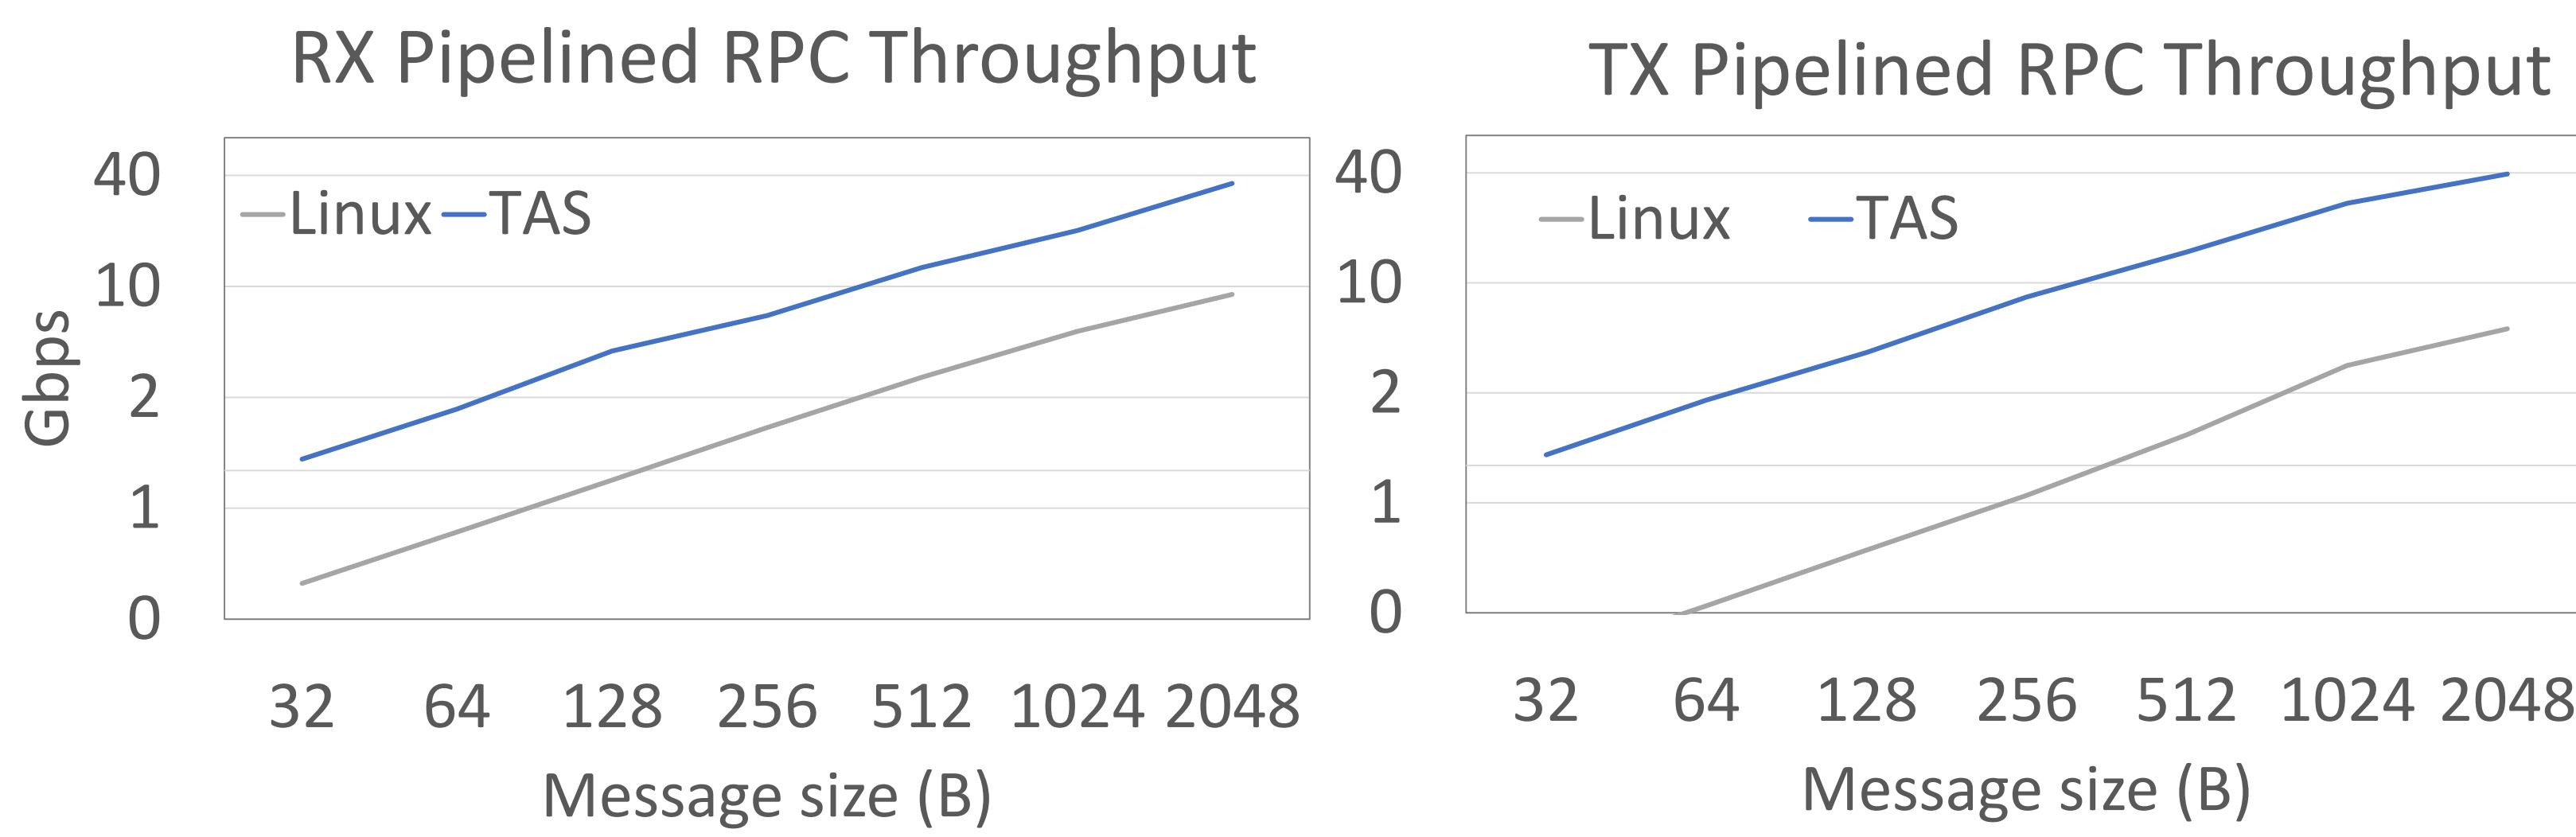
Faster Networks, Stagnant CPUs

- Network speeds rise beyond 100G. CPUs are stagnant.
 - TCP packet processing consumes an increasing portion of server CPU cycles
 - RPCs over TCP are common; applications want reliability
- | | kc | % |
|--------|-------|-----|
| Kernel | 15.77 | 94% |
| App | 1.07 | 6% |

Can we provide a high performance, flexible, protected TCP stack for datacenters in software?

	Features			Performance		
	Sockets	Protection	Cheap to Deploy	Kernel Bypass	Optimized Data Path	Dedicated Stack Cores
Linux	✓	✓	✓	✗	✗	✗
mTCP	✓	✗	✓	✓	✗	✗
Dune/IX	✗	✓	✗	✗	✓	✗
FlexSC	✓	✓	✓	✗	✗	✓
RDMA/TOE	✓	✓	✗	✓	✓	✗
TAS	✓	✓	✓	✓	✓	✓

Microbenchmarks

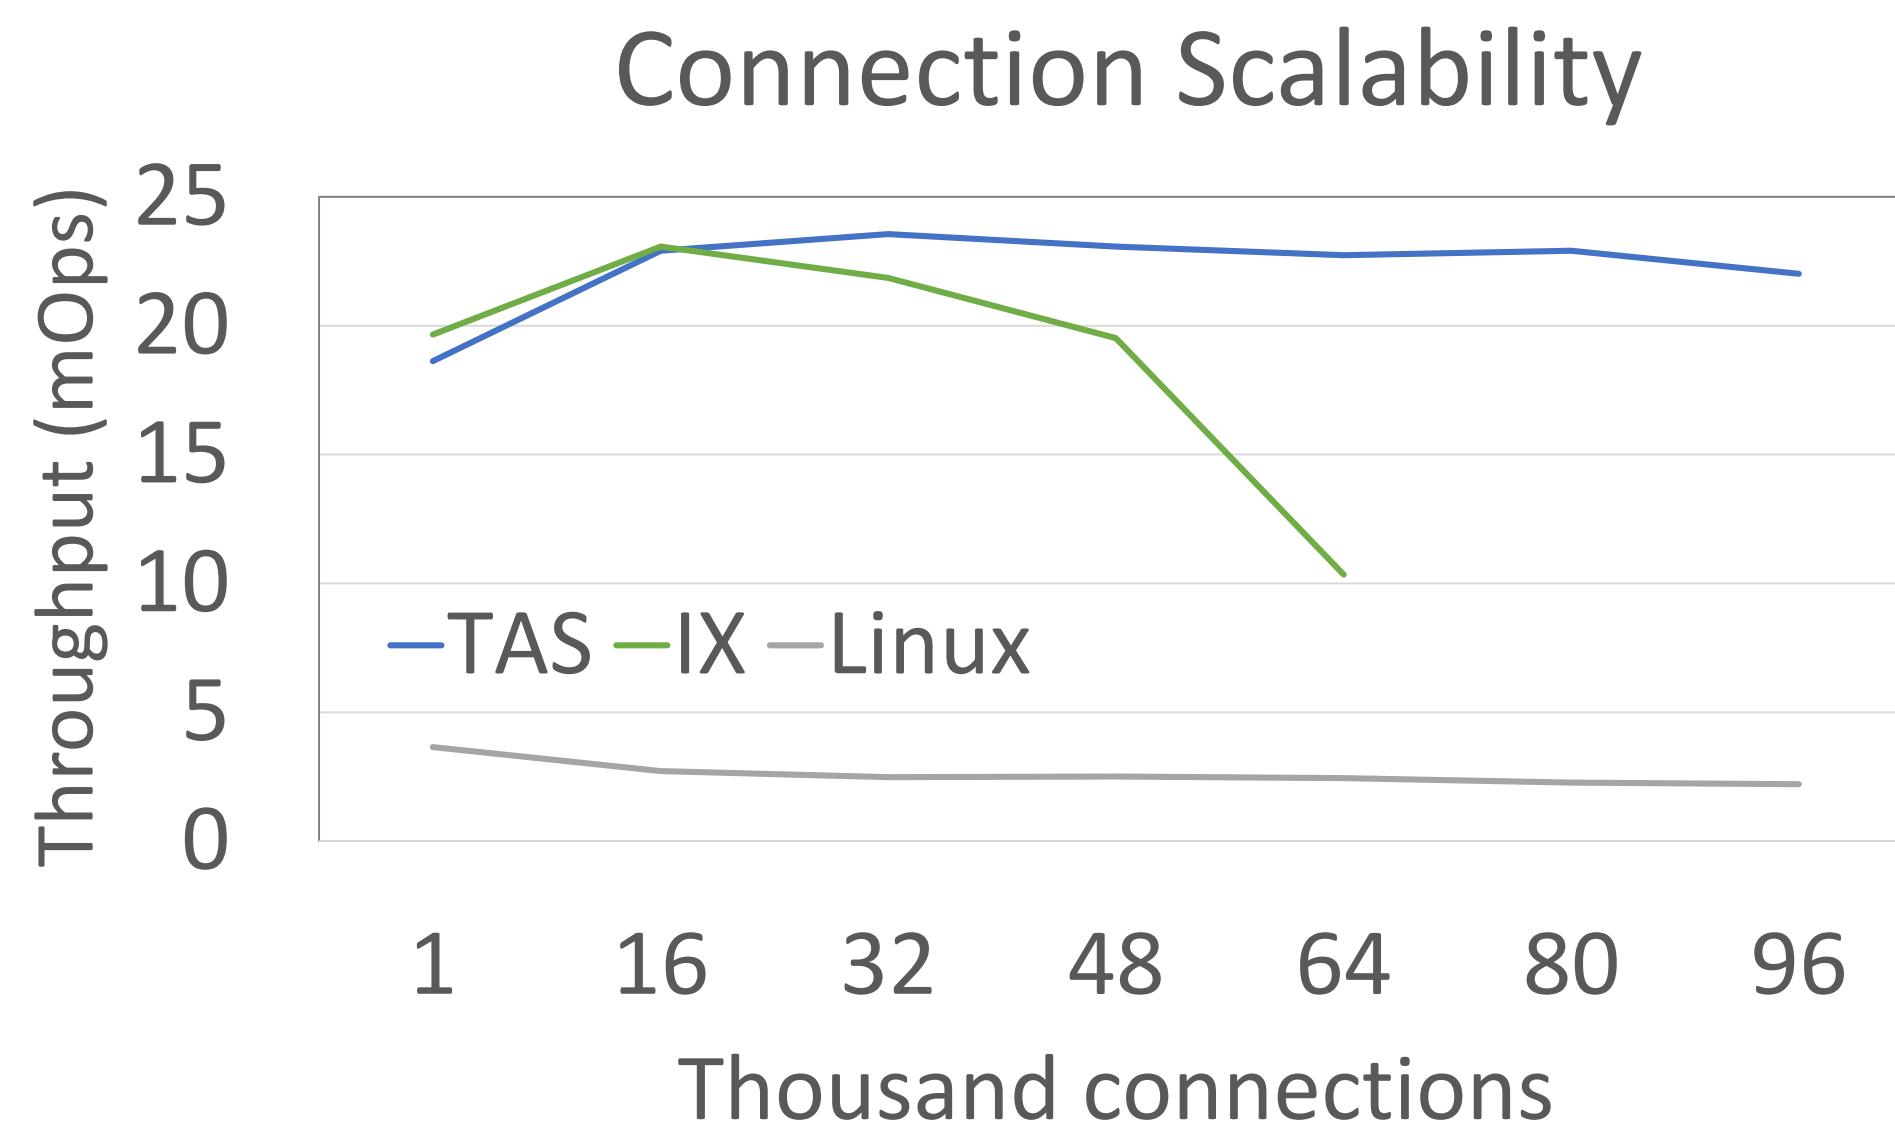


Unidirectional RPC test (no responses), 250 cycle workload

TAS reduces data path operations and cache misses/contention

Echo server with 64B requests/responses

IX's larger state limits connection scalability

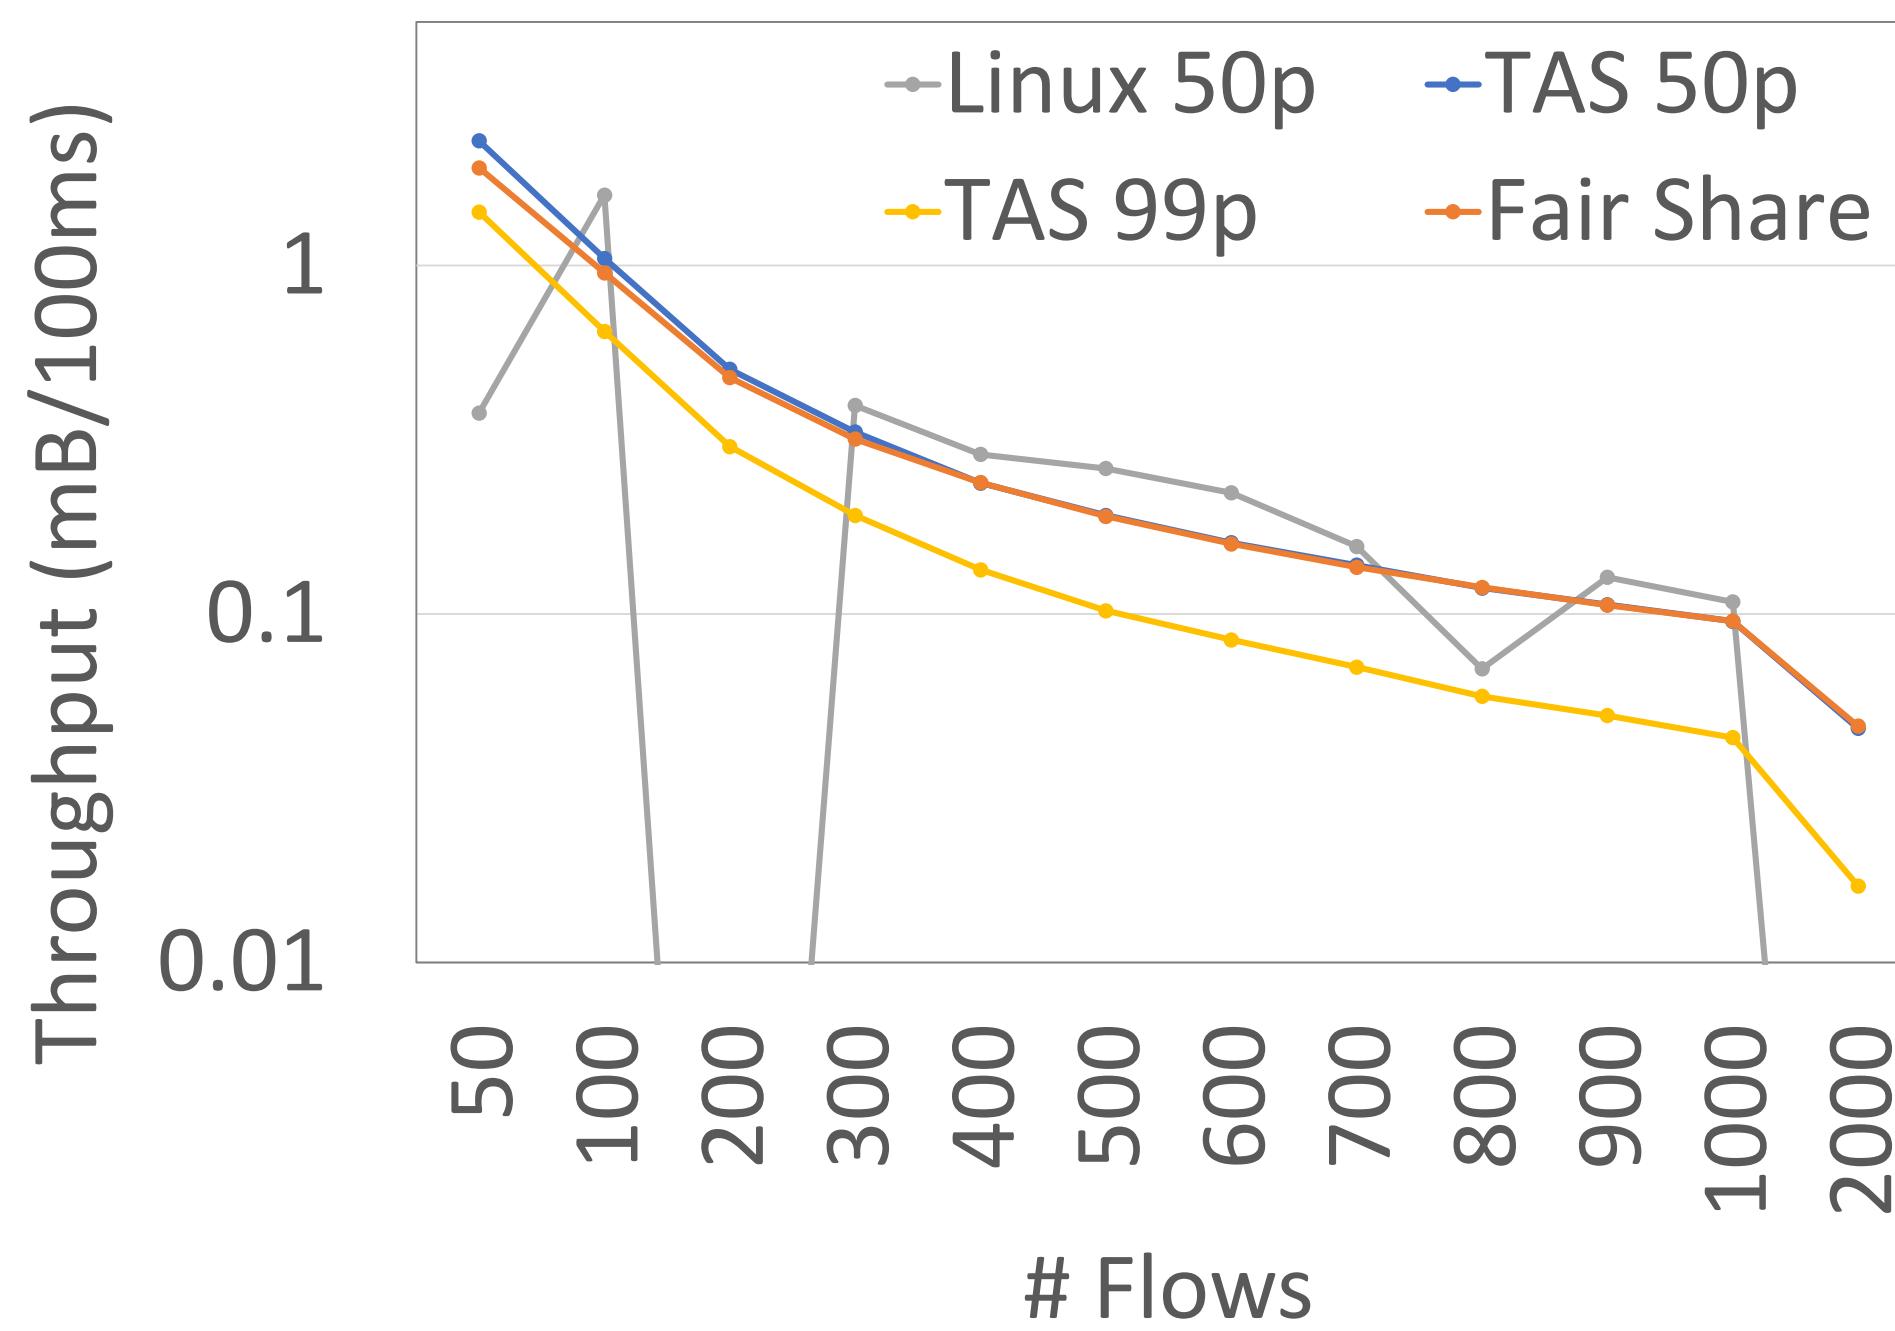


Congestion Fairness

Incast scenario

TAS rate-based CC provides fair share

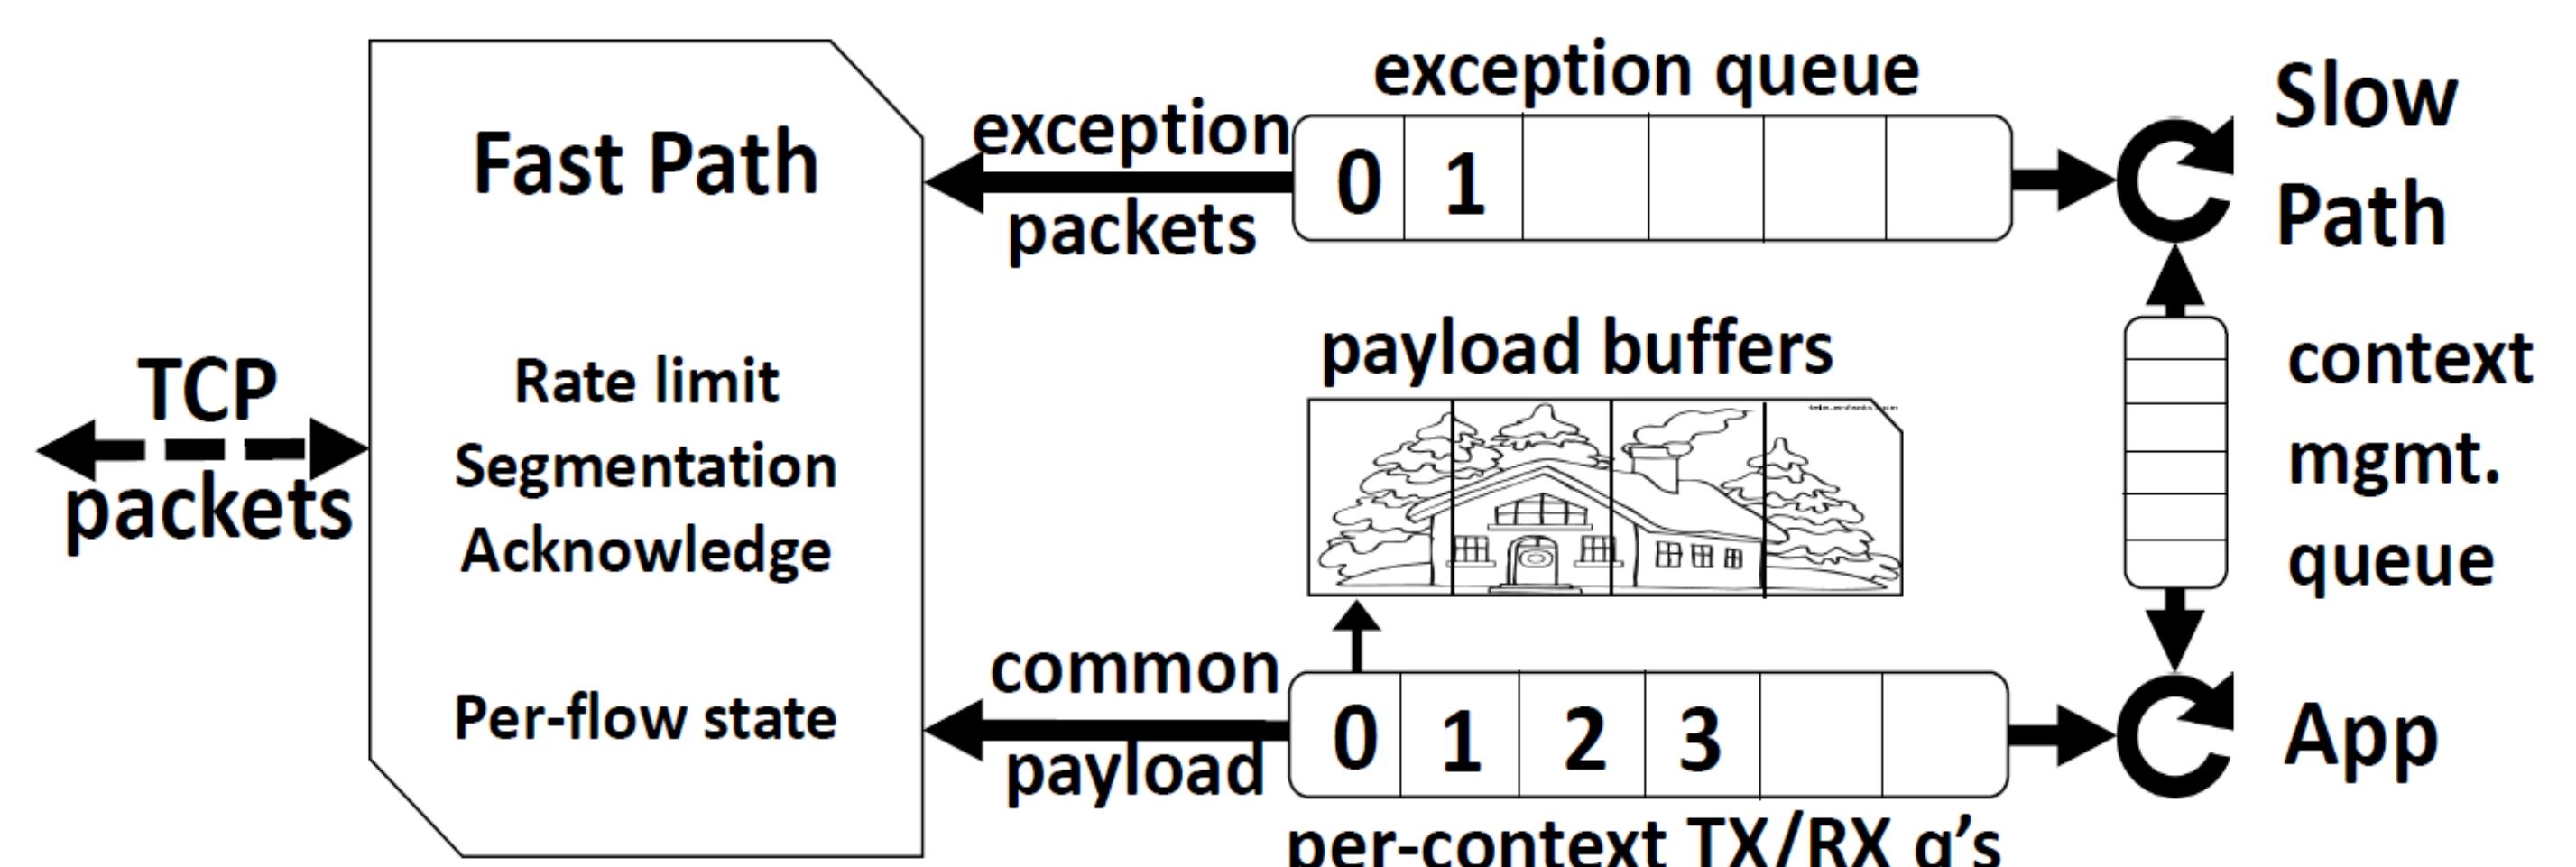
Linux is unpredictable; 99p is 0 for all cases



Design Decisions

1. Separate into Fast and Slow Path

- Fast path: in-order, dataplane packets
- Slow path: connection setup/teardown, congestion control algorithm, timeouts



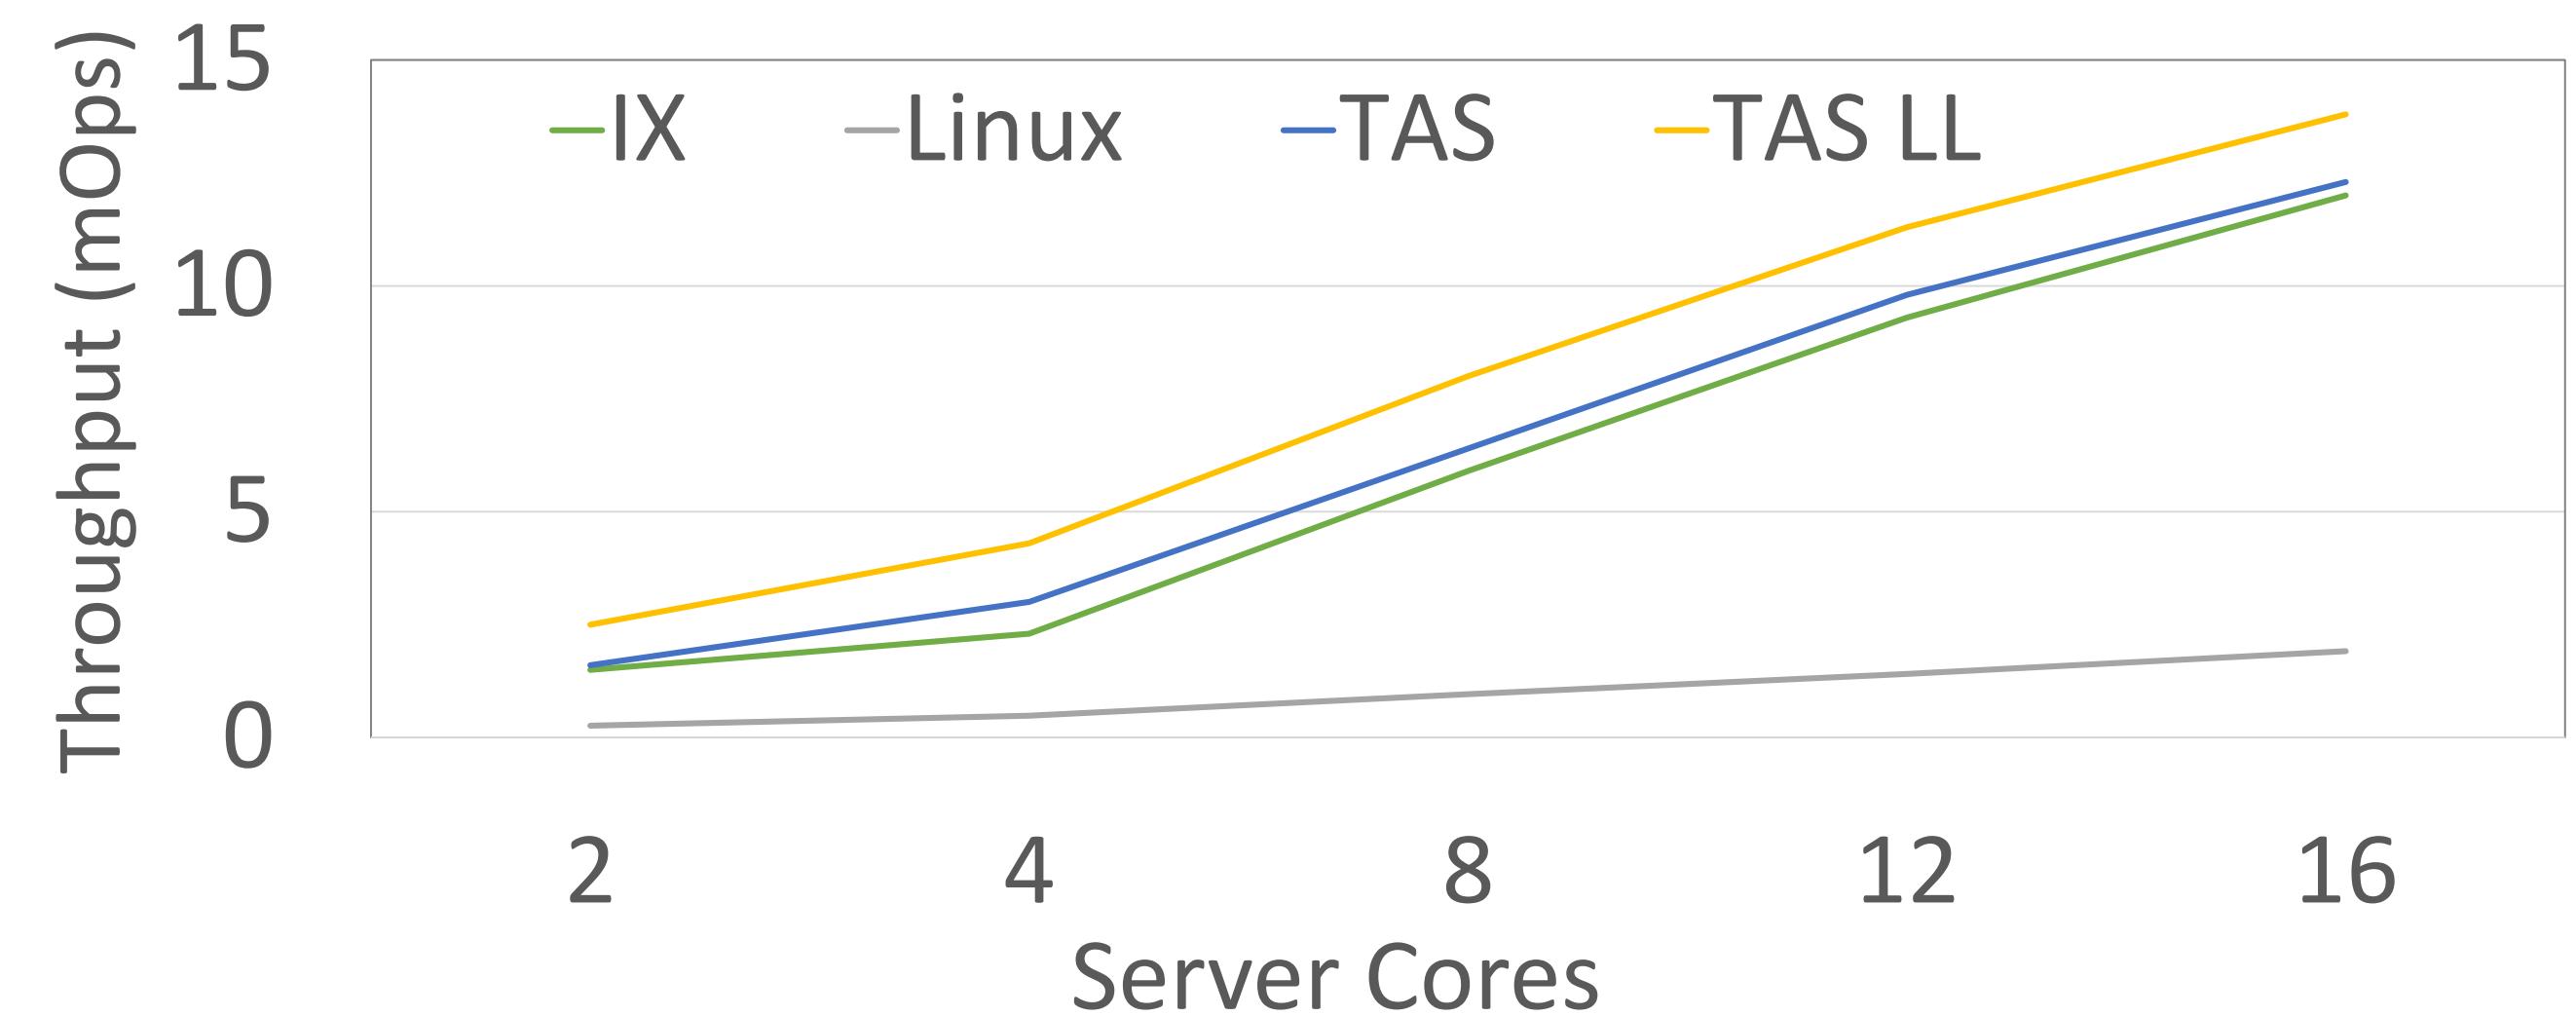
2. Minimize connection state in fast path

- Shrink and localize connection state
- Improve cache efficiency and scalability

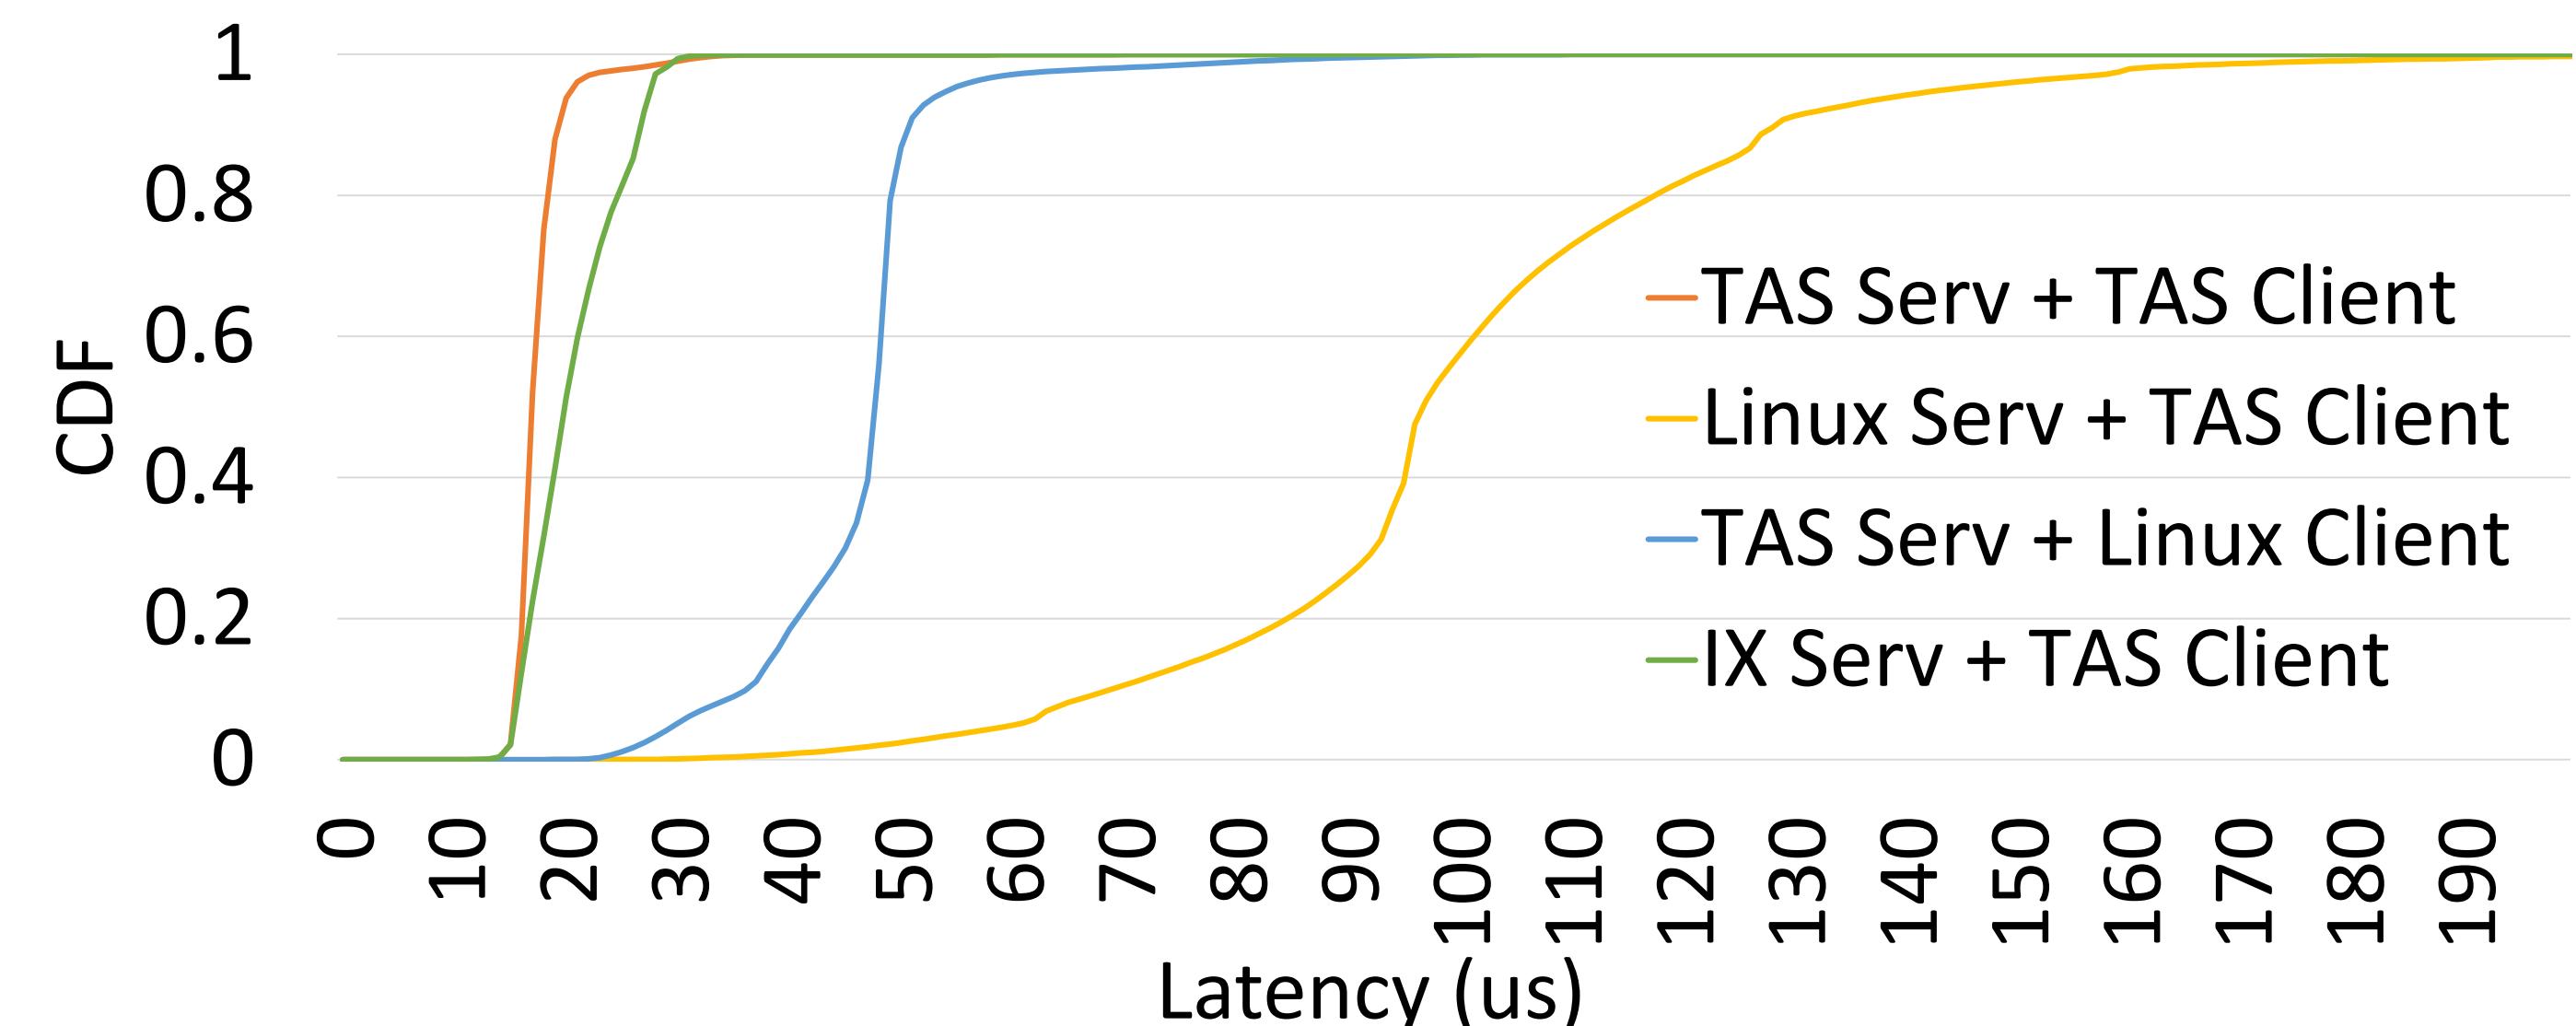
3. Dedicate cores for packet processing

- Run stack on separate cores from application
- Avoid cache pollution, gain workload proportionality

Key Value Store



IX doesn't provide sockets, requires KVS modification



IX 90p latency is 50% higher, max 230% higher (batching)

KVS Microarchitectural Analysis

Counters	Linux		TAS	
	App	Stack	App	Stack
CPU Cycles	1.1k	15.7k	0.7k	1.9k
Instructions	12.7k		3.9k	
CPI	1.32		0.66	
Frontend Bound	173	2600	102	248
Backend Bound	388	9046	353	684