



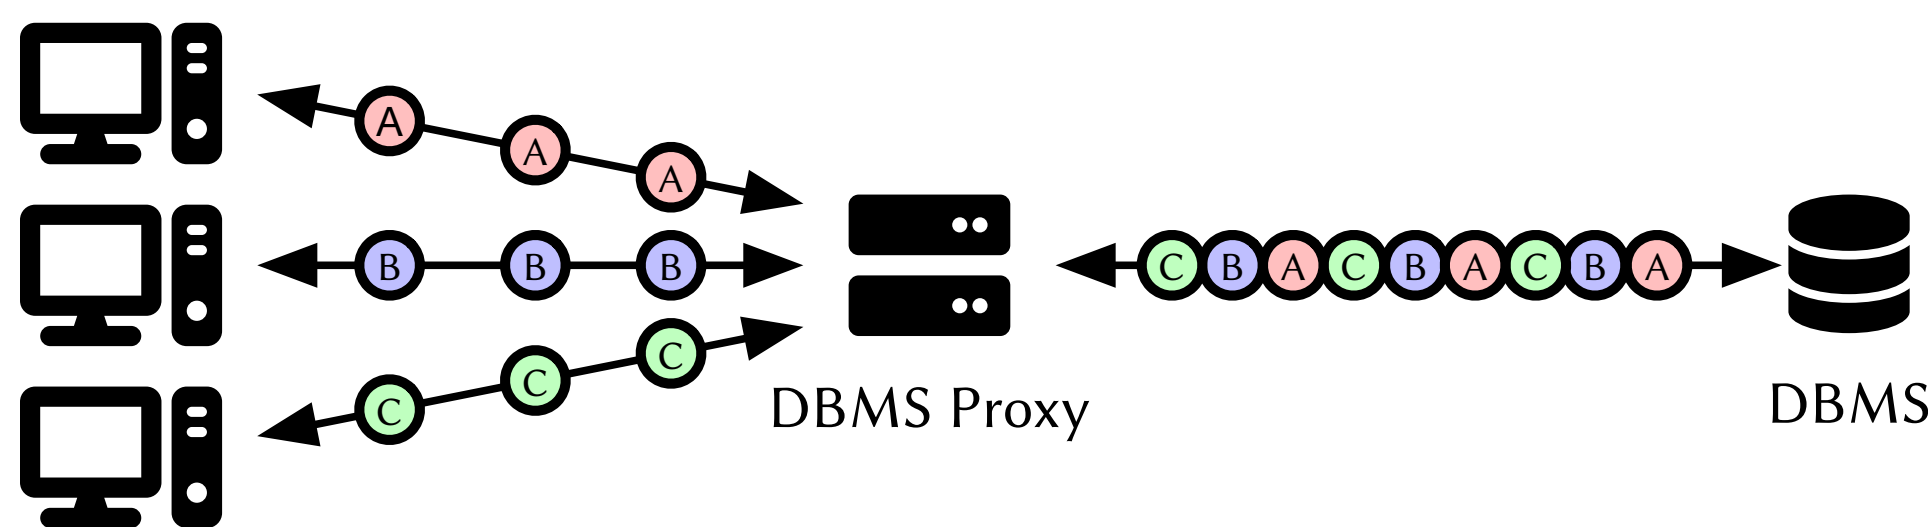
Tigger: A Database Proxy That Bounces With User-Bypass

Carnegie Mellon University

Matthew Butrovich, Karthik Ramanathan, John Rollinson*, Wan Shen Lim, William Zhang, Justine Sherry, Andrew Pavlo
Carnegie Mellon University, *Army Cyber Institute

Background

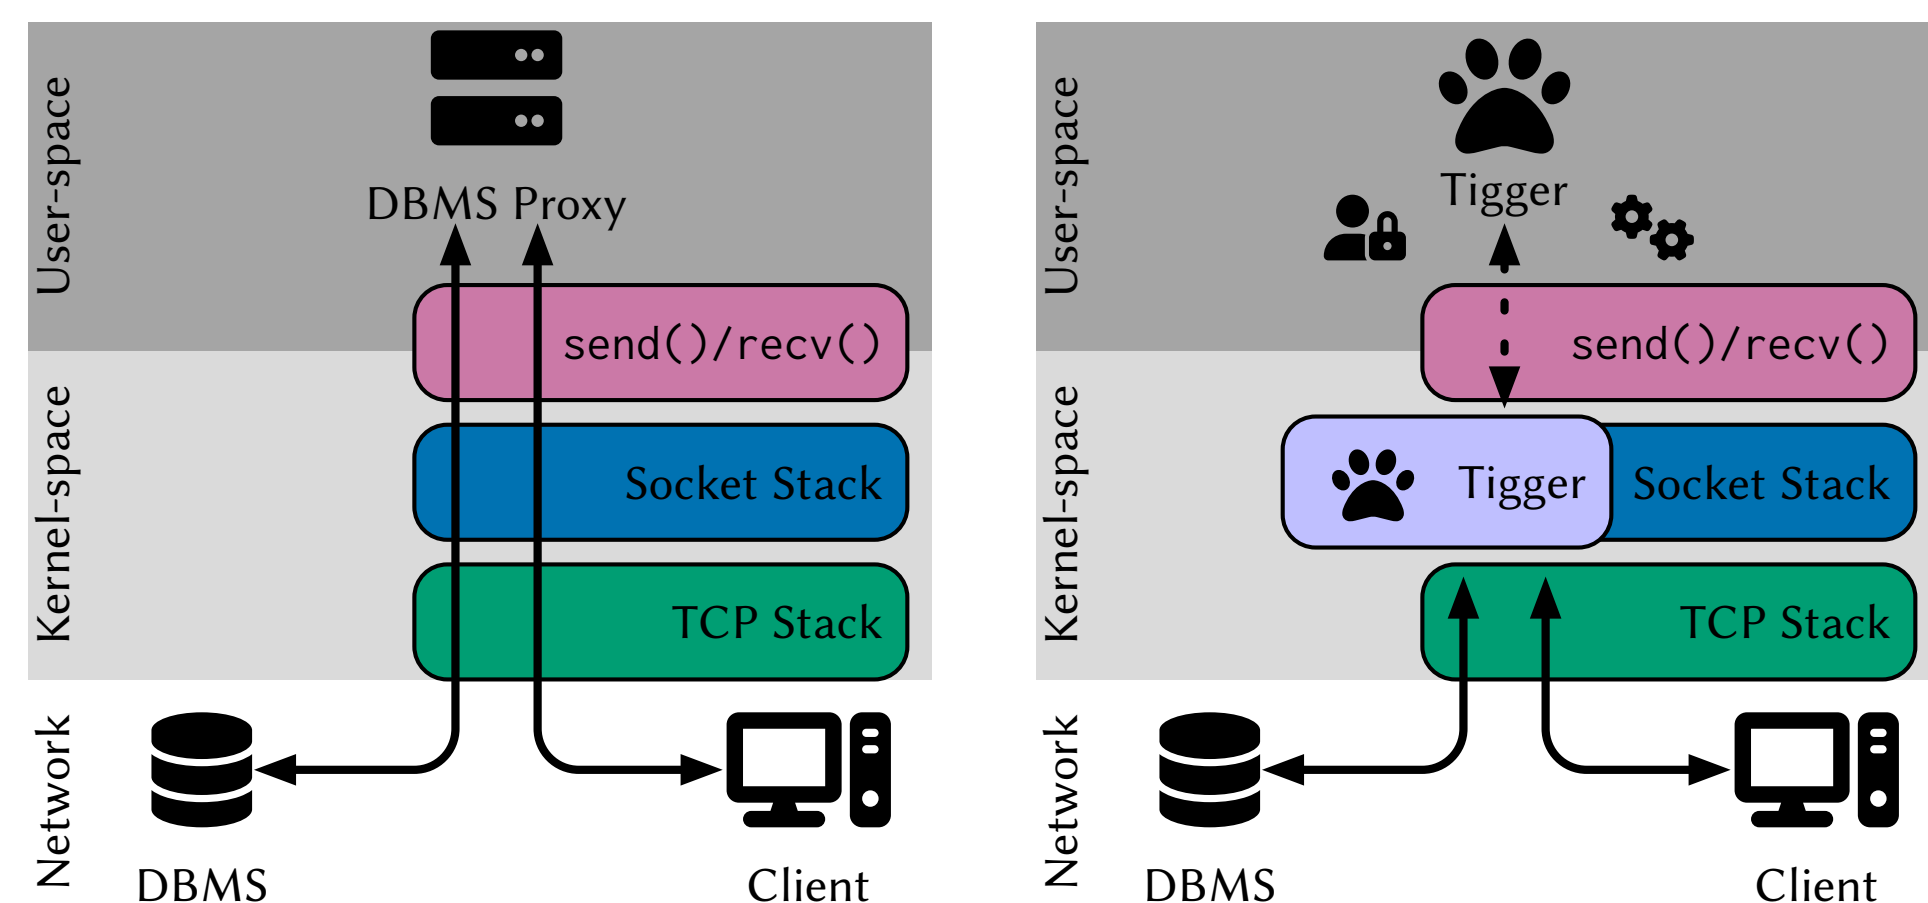
Developers often deploy database-specific network proxies whereby applications connect transparently to the proxy instead of directly connecting to the database management system (DBMS).



Clients

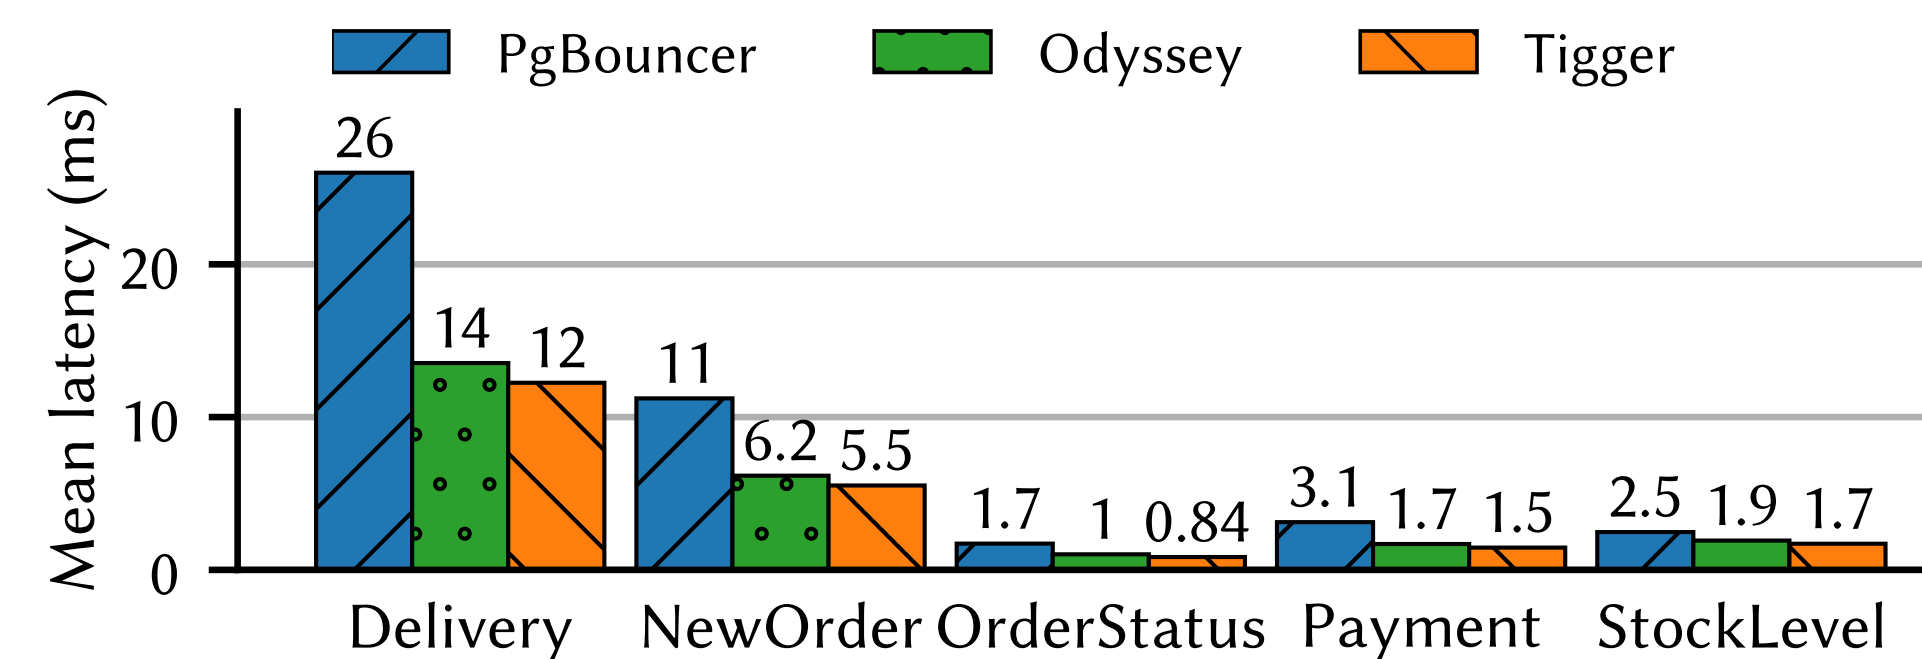
This indirection improves system performance because of connection pooling, load balancing, and other DBMS-specific optimizations. However, these Application Layer (L7) proxies are inefficient due to their reliance on OS system calls.

User-Bypass and Tigger

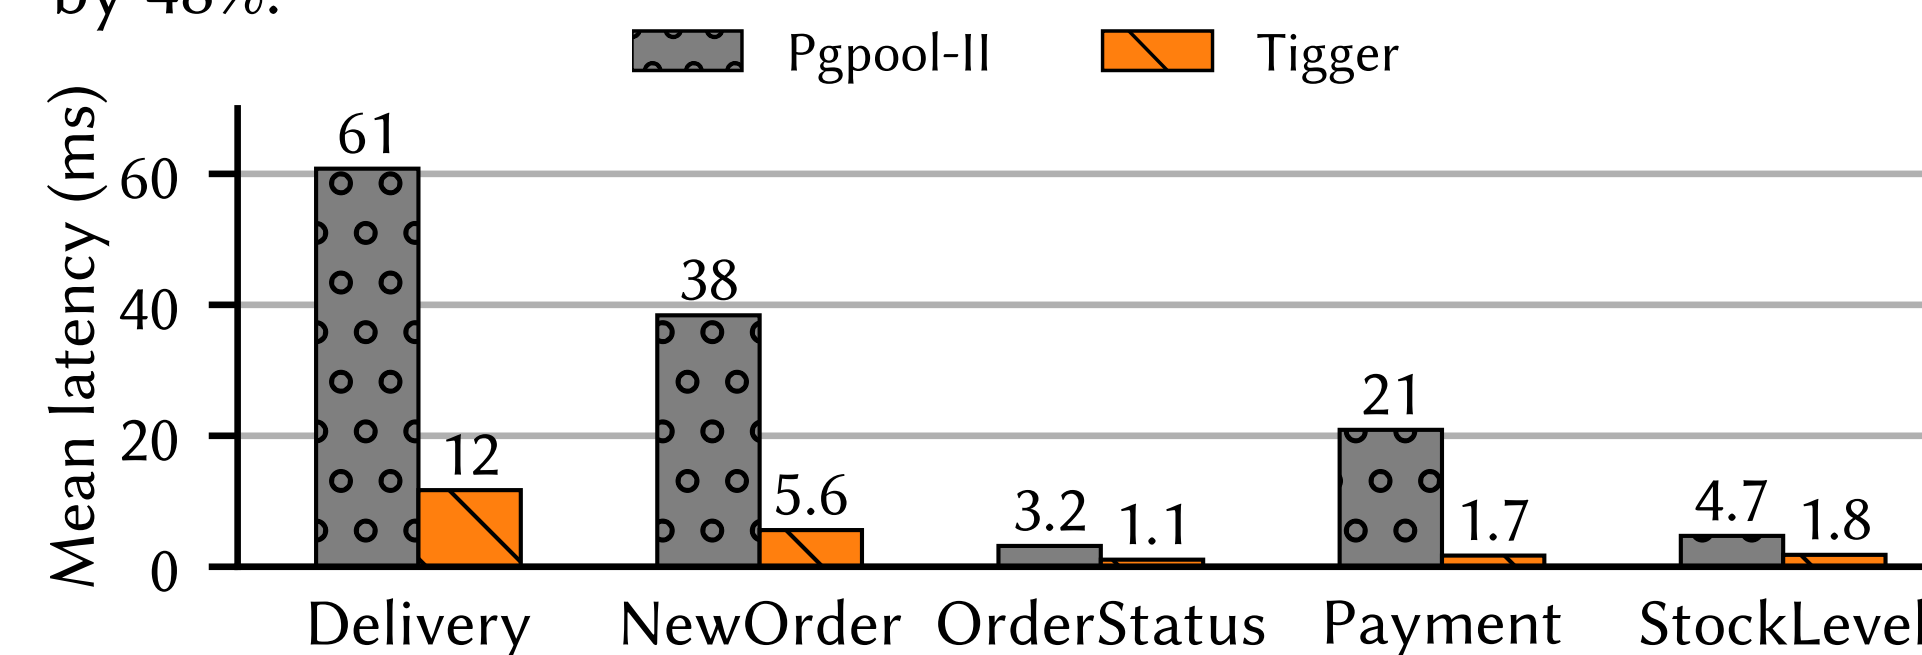


The **Tigger** DBMS proxy applies **user-bypass** to push application logic into kernel-space via Linux's eBPF infrastructure. This design eliminates the overheads associated with user-space proxies (e.g., copying buffers, scheduling user threads, and privilege escalation).

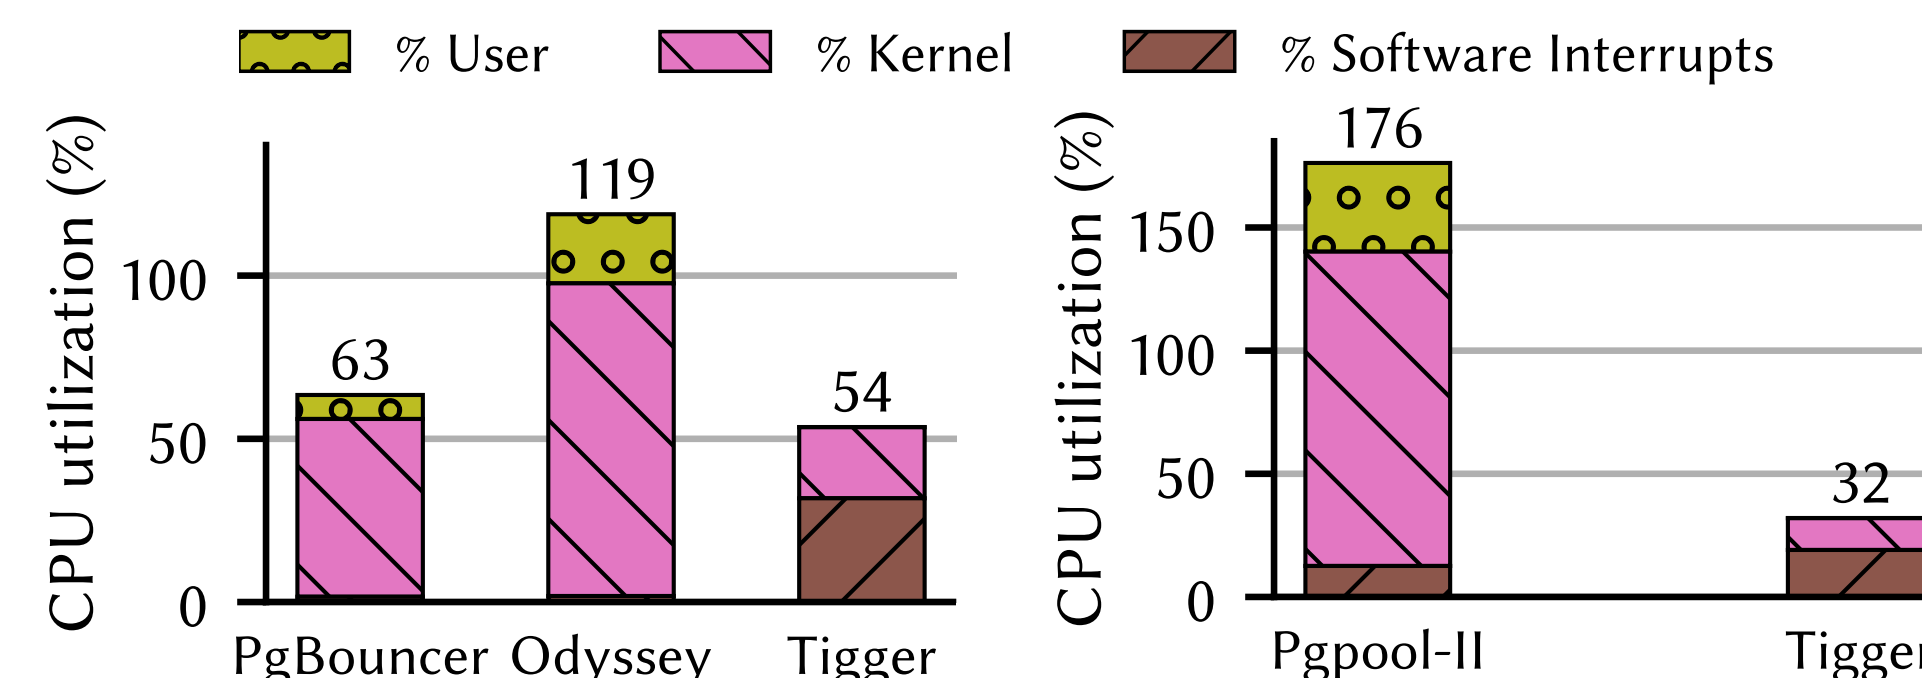
Evaluation



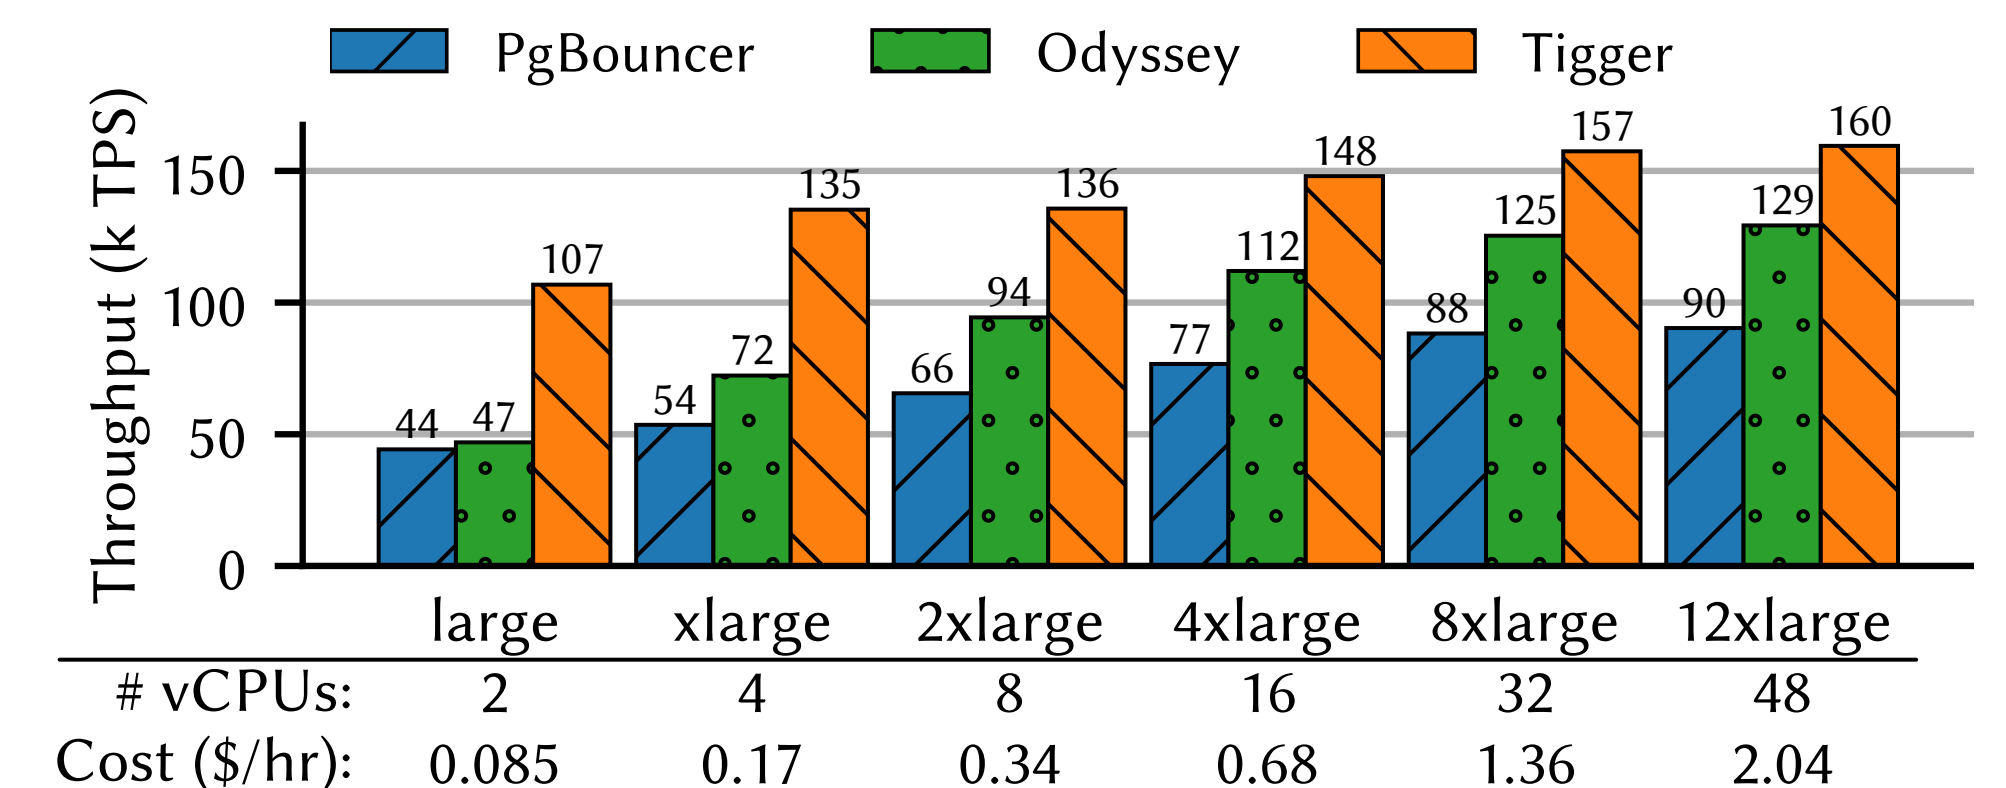
For transaction pooling with TPC-C, Odyssey's parallelism lowers latencies by 40%. Tigger's user-bypass design improves across every transaction profile, lowering latencies by 48%.



For workload replication with TPC-C, Tigger's reduction in system calls, and thereby buffer copies, lowers transaction latencies by 77%.



The CPU utilizations for the previous experiments show how Tigger's user-bypass design reduces User and Kernel CPU time. There is an increase in Software Interrupts time because Tigger's eBPF programs run on these kernel threads.



With 2 vCPUs, Tigger more than doubles the throughput of other proxies. Odyssey requires 8x the vCPUs (at 8x the cost) to match that baseline performance from Tigger.

Future Work

Ephemeral I/Os: Identify I/O reads that the DBMS discards shortly after use. Emphasize reads where minimal logic is applied to the data (e.g., inspecting page headers, following pointers in index structures).

Shuffle Nodes: Explore if nodes in distributed query processing (e.g., Apache Uniffle) benefit from user-bypass.

Conclusion

We introduce a user-bypass technique to safely avoid system calls, which copy buffers to and from kernel-space. User-bypass pushes application logic into the Linux kernel using eBPF. We show how to apply user-bypass with our PostgreSQL-compatible proxy Tigger. When compared against other proxies, Tigger offers the best performance and lowest operating cost. Our evaluation shows the value of user-bypass to reduce data movement between kernel-space and user-space and avoid costly system calls. In the future, DBMS components that perform ephemeral I/Os should be designed with user-bypass in mind.