

fRAC: Remote Accelerator Calls on Hostless FPGAs

Ziyi Yang*
KAUST
Thuwal, Saudi Arabia
ziyi.yang@kaust.edu.sa

Krishnan B. Iyer*
KAUST
Thuwal, Saudi Arabia
krishnan.iyer@kaust.edu.sa

Yixi Chen
KAUST
Thuwal, Saudi Arabia
yixi.chen@kaust.edu.sa

Ran Shu†
Microsoft Research
Beijing, China

Zsolt István
TU Darmstadt
Darmstadt, Germany
zsolt.istvan@tu-darmstadt.de

Marco Canini
KAUST
Thuwal, Saudi Arabia
marco@kaust.edu.sa

Suhaib A. Fahmy
KAUST
Thuwal, Saudi Arabia
suhaib.fahmy@kaust.edu.sa

Abstract

Hardware accelerators are essential for meeting microsecond-scale demands in distributed applications, but invoking them over the network remains bottlenecked by host software. Host-mediated paths (RPC, RDMA) still rely on the CPU to reassemble requests and launch accelerators, which dominates end-to-end latency, while packet-level offload on SmartNICs lacks any notion of a complete request. We introduce the Remote Accelerator Call (RAC) paradigm, in which transport ingestion, request reassembly, dispatch, and accelerator invocation are handled entirely in hardware. We build fRAC, a hostless prototype that implements this network-to-accelerator interface as a lightweight shim on an FPGA hosting multiple virtualized accelerators, reconfigurable at runtime in a few milliseconds. Serverless functions backed by fine-grained accelerators are a natural fit, since host-mediated invocation overhead can match or exceed their short execution times. fRAC achieves under 7.8 μs end-to-end latency while sustaining application throughput of up to 90 Gbps. Compared to highly optimized DPDK-based CPU and DPU baselines, it achieves at least 80% lower latency under saturation, while consuming 85% less power.

CCS Concepts: • **Computer systems organization** → **Reconfigurable computing**; *Client-server architectures*; • **Hardware** → **Hardware accelerators**; • **Networks** → *In-network processing*.

Keywords: Field programmable gate arrays, hardware acceleration, serverless computing, in-network computing, remote procedure call.

*Both authors contributed equally to this work.
†Now at Broadcom.



This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGOPS ATC 2026, Hong Kong, China
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2818-1/2026/11
https://doi.org/10.1145/3828523.3857070

ACM Reference Format:
Ziyi Yang, Krishnan B. Iyer, Yixi Chen, Ran Shu, Zsolt István, Marco Canini, and Suhaib A. Fahmy. 2026. fRAC: Remote Accelerator Calls on Hostless FPGAs. In *The 2026 ACM SIGOPS Annual Technical Conference (SIGOPS ATC 2026), November 15-18, 2026, Hong Kong, China.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3828523.3857070

1 Introduction

Serverless platforms increasingly compose distributed applications from short-lived, compute-intensive functions. While hardware accelerators (GPUs, FPGAs, ASICs) are well suited to executing these workloads with high efficiency [13, 22, 30, 38], their integration into distributed environments remains bottlenecked by software-mediated invocation paths. Because serverless functions are inherently distributed, their inputs must be delivered over the network from clients or peer services. When integrating accelerators to process these networked requests, current architectures adapt existing communication paradigms. These include high-level RPC frameworks (standard in serverless platforms) and low-level RDMA (often used in high-performance accelerator clusters). Across all these approaches, moving data to accelerators and invoking them requires software intervention.

In standard RPC frameworks, host software parses requests, recovers request boundaries from fragmented packets, deserializes data, and schedules accelerator invocation [19, 39, 40]. Each of these steps consumes CPU cycles that scale with both payload size and request rate, and the resulting per-request software overhead is often comparable to or larger than the accelerator's computation time for microsecond-scale functions. Even when data moves over PCIe to host-attached GPUs, FPGAs, and ASICs, software is still needed to configure DMA, detect transfer completion, and launch kernels through runtimes such as CUDA [63] or XRT [2], each of which adds tens of microseconds of queue management and doorbell overhead per invocation [17, 22, 31, 54]. To reduce these data-copy overheads, RDMA-capable NICs can DMA payloads directly into accelerator memory, whether GPU DRAM via GPUDirect RDMA or FPGA buffers exposed by an accelerator runtime [2, 61, 79]. However, RDMA exposes

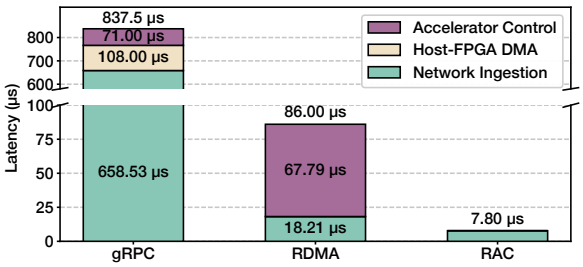


Figure 1. Latency breakdown of remote accelerator invocation for a 4KB payload across three paradigms. Host-Controlled (gRPC) is dominated by software network and RPC processing. RDMA with a host-side accelerator runtime optimizes the data path but leaves accelerator invocation under software control. RAC handles ingestion, dispatch, and invocation in hardware, eliminating the host-side segments altogether.

device or memory-region endpoints rather than request-granular accelerator endpoints. It can place data into memory, but does not natively dispatch complete requests across multiple accelerator instances or trigger selected accelerators when requests are ready. The host CPU still posts work requests, polls completion queues, and issues kernel launches. Accelerator deployment frameworks additionally place request routing, scheduling, and device management on the software invocation path, adding overhead beyond data movement itself [22, 54]. Host mediation also keeps the CPU and its host platform active, adding substantial power overhead to otherwise lightweight accelerator invocations. This host-managed model works well for the workloads that GPUs and TPUs are typically deployed for, such as large-batch, long-running computation (e.g., training and inference on large tensors), where the per-invocation control overhead is amortized over substantial compute time. For short-lived, microsecond-scale functions, however, this overhead remains on the critical path and becomes the dominant bottleneck once data movement is no longer the limit (Figure 1).

Packet-processing hardware (SmartNICs, programmable NICs, and switch-based offload) keeps traffic in the datapath, but operates at packet granularity. Each packet is treated independently, with no notion of application-level requests, requiring accelerators themselves to manage payload buffering and invocation on a per-accelerator basis [10, 33]. At the transport layer, offloads such as TCP Offload Engines (TOEs) [57] can reconstruct reliable, ordered byte streams in hardware, but these streams do not expose application request boundaries or accelerator-invocation semantics. Thus, across all these approaches, data movement may be optimized in hardware, but request formation and accelerator invocation remain on a software-run critical path.

A hardware path that connects transport ingestion directly to accelerator execution remains absent [25]. Such a path must provide four capabilities without placing software on the critical path. First, from the reliable, in-order byte stream provided by the transport layer, it must identify request boundaries and form complete application-level invocation requests. Second, it must dispatch complete requests to one of multiple accelerators while supporting multiple clients concurrently. Third, it must trigger accelerator execution when a request is ready, rather than waiting for software control. Fourth, it must virtualize accelerators rather than binding each network endpoint to a fixed accelerator. We call this paradigm Remote Accelerator Calls (RAC), where transport-to-accelerator invocation is handled directly in hardware. RAC is broadly applicable to accelerators that can be coupled with a transport endpoint on the same silicon.

We instantiate this paradigm in fRAC, a hardware prototype that implements these capabilities for network-attached FPGA accelerators. fRAC implements accelerator virtualization in hardware, exposing diverse accelerators through a common request-level invocation abstraction. Rather than being custom-built for a fixed accelerator, this abstraction lets flexible requests identify and invoke one of multiple accelerators through the same hardware interface. In contrast, RPC- and RDMA-based frameworks may optimize data movement to an accelerator, but they do not provide a hardware abstraction for selecting among and invoking multiple diverse accelerators. They rather require such invocation abstraction to be implemented in host software. While our prototype implementation is built atop TCP, fRAC is designed to be transport layer agnostic. Our primary focus is data movement, invocation and virtualization of accelerators. Motivated by the above, we make the following contributions:

- We define the Remote Accelerator Call (RAC) paradigm and implement the fRAC prototype that decouples network data transfer from accelerator invocation through a request-level invocation interface while preserving isolation across requests.
- We show that fRAC adds only modest FPGA resource overheads while delivering high performance, achieving application throughput of up to 90 Gbps with end-to-end invocation latency under 7.8 μs, and enabling accelerator reconfiguration in under 10 ms.
- We evaluate fRAC against highly optimized DPDK-based CPU and DPU implementations, demonstrating at least 80% lower latency than both platforms in high-throughput scenarios.
- We show that programmer effort to port existing accelerators to fRAC is minimal, as most designs already expose a suitable streaming interface.

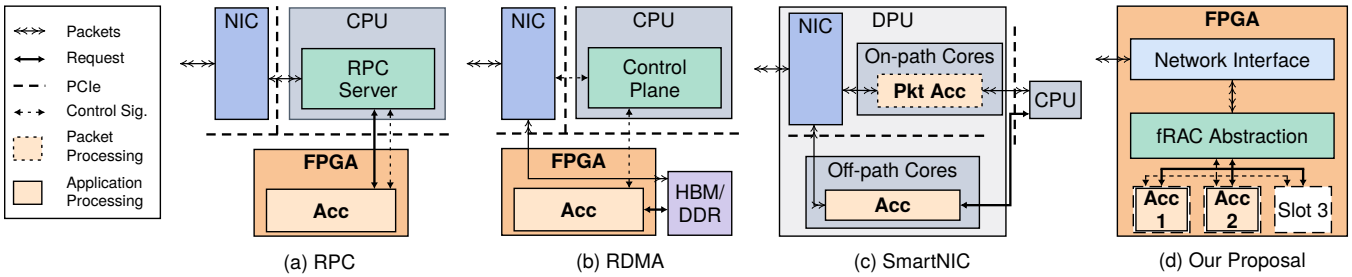


Figure 2. A comparison of network-to-accelerator invocation paths. **(a) RPC**: host software handles network processing, request parsing, scheduling, and accelerator invocation. **(b) RDMA**: the NIC can DMA data into registered memory regions reachable by the accelerator, but request dispatch and accelerator invocation remain host-managed. **(c) SmartNIC**: programmable datapaths either process traffic on-path at packet granularity (e.g., ClickNP [44], PANIC [48]) or support request-level processing off-path by traversing PCIe (e.g., FlexDriver [23]). Neither natively reassembles complete application requests for direct on-path accelerator invocation. **(d) fRAC**: a hostless FPGA path reassembles network fragments into complete serverless-function requests and invokes accelerators at application granularity.

2 Background and Motivation

Small, compute-intensive functions are increasingly offloaded to hardware accelerators to meet strict performance demands. However, integrating these into distributed environments requires an efficient network-to-accelerator path. Figure 2 contrasts RAC with prior offloading paradigms. Existing approaches rely on host-mediated control (e.g., RPC, RDMA) which add significant software latency, or use packet-level processing (e.g., SmartNICs) which cannot handle application-level requests efficiently. RAC eliminates the host CPU from the critical path by directly ingesting network data, reassembling complete requests, and triggering accelerators entirely in hardware. We now discuss the rationale for this design and defer to Section 6 for a detailed comparison with related work.

2.1 The Cost of Software-Mediated Invocation

To understand why software remains a persistent bottleneck, we examine how the semantic gap between existing network abstractions and accelerator execution models manifests as a performance penalty. A complete accelerator invocation requires not only moving the request payload to the accelerator, but also explicitly triggering execution and polling for completion.

Figure 1 compares three invocation paradigms for a 4KB request and isolates where the bottleneck shifts as software is removed from the path. **Host-Controlled (gRPC)** pays for everything in software, including network ingestion, host-FPGA DMA, and kernel launch (over AMD XRT [2] in this case). Because RPC abstractions are designed for general-purpose software services, they incur serialization and context-switching overheads that dwarf the execution time of microsecond-scale hardware functions. **RDMA**[1] removes much of the software cost for networking and data movement. However, one-sided RDMA exposes only low-level read, write, and atomic operations [11]; richer behavior is possible only through elaborate verb chaining [69]. The host CPU must therefore still intervene to synchronize and launch execution through an accelerator runtime such as XRT [2] or CUDA [63]. This mismatch between data movement and accelerator invocation leaves software overhead as the dominant remaining bottleneck. Today, it can be avoided only by hardwiring a fixed-function streaming accelerator into the network datapath, sacrificing the ability to share, select among, or replace accelerators. **RAC** directly bridges this gap by providing an abstraction that is both expressive enough to support application-level invocation requests and lightweight enough to be implemented entirely in hardware, removing the software control path while simplifying accelerator deployment, as illustrated in Figure 2d.

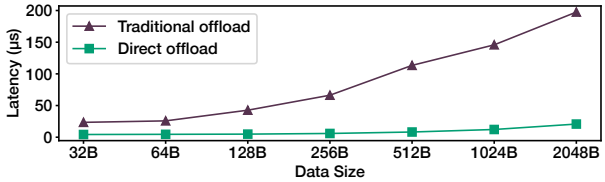


Figure 3. Median latency under host-controlled data movement and direct accelerator invocation, isolating the payload-size-dependent cost of the host PCIe path.

Figure 3 provides a complementary view of host-mediated data movement, isolating the PCIe penalty as payload size scales. Using a data sketch workload (Count-Min and Top-K) in ClickNP [44] on Microsoft Catapult [66], we compare traditional host-controlled data movement between two nodes implementing the two functions against direct accelerator chaining via the network. Relying on the host CPU to move data across PCIe imposes a latency penalty that scales poorly with larger payloads. Direct invocation bypasses the host PCIe path entirely, keeping data on the network or device and mitigating this scaling bottleneck. This reveals a fundamental limitation: as long as a host CPU bridges the network

[1] Our measured RDMA baseline uses a Mellanox NIC to DMA request data directly into XRT-allocated FPGA P2P buffers, while FPGA kernel invocation is still performed through XRT.

and the accelerator, invocation will bottleneck performance, especially when distributing across nodes. To unlock hardware acceleration for microsecond-scale serverless functions, the invocation model must bypass the host entirely, allowing the network to deliver data and trigger accelerator execution directly in hardware.

2.2 Fragmented Request Invocation

Application-level requests often have a granularity that exceeds the network MTU, leading to data being fragmented across multiple network packets. While transport protocols transparently segment and reassemble an in-order byte stream, they do not preserve message boundaries. Hardware that operates directly on such a stream therefore faces a semantic mismatch between network-level packets and application-level requests.

Hardware accelerators typically cannot context switch mid-request, since their intermediate state is distributed across pipeline registers and on-chip buffers that cannot be cheaply saved and restored. Once an accelerator begins consuming a request, it must remain dedicated to that request until it completes. If request fragments are forwarded directly as they arrive from the network, interleaved requests cause head-of-line blocking: fragments of other requests are ready, but the accelerator cannot accept them. For example, a streaming image filter that has begun processing one image stalls whenever a subsequent packet is delayed by jitter or congestion, creating "invocation bubbles". In a multi-tenant setting, a stalled accelerator is a locked resource: it cannot serve other clients while it waits for a delayed packet. Prior RPC-offload designs reassemble requests before dispatch but still involve the host: NEBULA [73] delivers them to CPU-run handlers, and RpcNIC [80] runs reconfigurable kernels on a PCIe-attached SmartNIC under host-software control.

Eliminating these bubbles requires decoupling network delivery from accelerator execution. Incoming fragments should be buffered, and requests dispatched to accelerators only once they have been fully reassembled. This absorbs network delay before invocation: once started, the accelerator streams the request without interruption and can immediately proceed to the next queued request. Although each accelerator could buffer its own inputs, it would need a separate buffer for every concurrently arriving request, provisioned for the worst case and duplicated across every accelerator. Our key insight is to lift reassembly out of individual accelerators into a shared, accelerator-agnostic layer, where buffer capacity is multiplexed across tenants and any streaming accelerator can be deployed without modification.

2.3 Serverless Functions with RAC

Serverless computing is a natural fit for the RAC model: applications are decomposed into short-lived, highly concurrent functions invoked on demand. Crucially, each invocation operates on a complete input, such as an image or a tensor block, rather than on individual network packets. As Table 1 shows, many such functions have efficient hardware implementations, and once accelerated, many complete in microseconds or milliseconds. For the shortest of these, host-mediated invocation can add latency comparable to the function itself, and for all of them, stalls caused by fragmented network delivery leave accelerators underutilized. These functions therefore require an execution model in which invocation granularity matches the input size and execution starts as soon as a complete input is available, which is precisely what RAC's complete-request invocation provides.

A serverless deployment separates orchestration from execution: a control plane places functions and manages workers, while workers serve invocation requests. RAC allows these workers to be hostless accelerator nodes: requests are reassembled and dispatched directly to concurrent accelerator pipelines, without host involvement, serving multiple clients in isolation and without invocation bubbles.

3 fRAC Design

fRAC is a lightweight framework that realizes RAC on network-attached FPGAs. FPGAs are a practical substrate for RAC: a single device exposes multiple accelerator slots, and dynamic partial reconfiguration allows accelerators to be swapped at runtime to adapt to changing demand [4, 75]. While most FPGA virtualization frameworks rely on host software to multiplex clients, manage execution, and reconfigure accelerators [14, 24, 41], fRAC performs these tasks on the FPGA itself.

Incoming network data is reassembled into complete application-level requests and buffered on-chip; each request is then routed directly to its target accelerator, which processes it to completion before serving the next. A reconfiguration controller swaps accelerators into slots on request from the control plane, keeping the network path active so that functions in other slots continue executing. Because fRAC is built atop a standard network stack (Figure 4), it requires no changes to existing infrastructure.

In a serverless deployment, fRAC operates as a hostless worker behind an existing control plane and gateway, which handle orchestration, request routing, and platform security; we therefore focus on the worker's datapath. fRAC hosts a library of hardware functions rather than arbitrary FaaS code. Supported functions expose a streaming interface, keep their per-invocation working set on-chip, and fit within the resources of a single accelerator slot; persistent state remains outside the FPGA.

fRAC provides request-level data isolation: each reassembly buffer holds only one in-flight request, and accelerators run each request to completion, so data from different requests is never interleaved. It also isolates stalls: only complete requests are dispatched, so a slow sender holds a buffer

Table 1. Examples of accelerators used in distributed applications and their FPGA implementations.

Category	Example Accelerators	Data Granularity	Runtime	FPGA Impl.
Audio/Video	Filtering, Resizing, Speech-to-text, Video encoding	Frame (≈ 200–500KB)	~2ms	[28, 50, 76]
Mathematical	Matrix Multiplication, SVD, Cholesky decomp., EMA	Matrix (≈ 2–32KB)	~10µs	[21, 51, 52, 59, 71]
Machine Learning	CNN inf., Clustering, Bayes classific., Aggregation	Tensor Block (≈ 3KB–3MB)	~5ms	[7, 16, 37, 65]
Data Analysis	Top-K selection, Count-Min Sketch	Stream data	~1µs/KB	[35, 42, 56]
ML Pre-Processing	Logit Transform, Min-Max Normalization	Tensor Block (≈ 3KB–3MB)	~3–30µs	[32, 58]

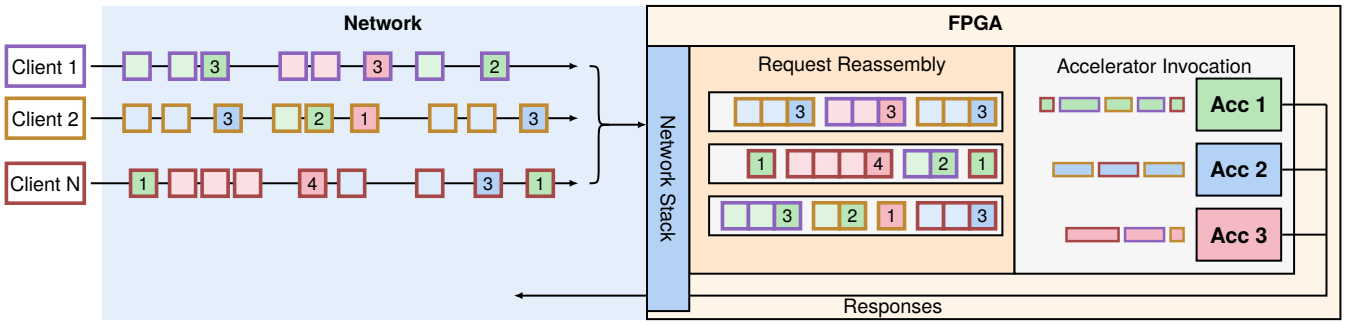


Figure 4. Overview of fRAC operating model. Multiple clients (outline color) issue requests addressing different accelerators (fill color), which are fragmented over multiple packets, the first of which determines the size. These fragments are reassembled into complete requests and dispatched to the corresponding accelerator. Reassembly is performed in buffers that are agnostic to client and accelerator for resource efficiency.

rather than an accelerator, and a long-running function delays only its own queue.

The remainder of this section addresses three design challenges: buffer management, accelerator scheduling, and accelerator reconfiguration. We validate each design decision in simulation[2] before presenting the implemented system in Section 4.

3.1 On-Chip Request Reassembly

To eliminate the invocation bubbles described in Section 2, fRAC reassembles requests on-chip before invoking any accelerator. Reconstructing the byte stream remains the job of the transport layer; fRAC instead recovers request boundaries within that stream, in hardware, across many concurrent connections, so that a complete request becomes the unit of accelerator invocation. The first fragment of each request carries a custom header specifying its total size and target function (cf. Figure 11). On arrival fRAC allocates a reassembly buffer for that request from a shared pool, using on-chip FPGA memory rather than external DRAM/HBM to minimize latency, and appends fragments to it. Once the request is complete, it is forwarded to the target accelerator's queue. Message framing itself is standard in software RPC frameworks; what fRAC contributes is performing it in the datapath of a hostless device and binding it directly to accelerator dispatch.

[2]The simulator setup is detailed in Appendix A (Table 5).

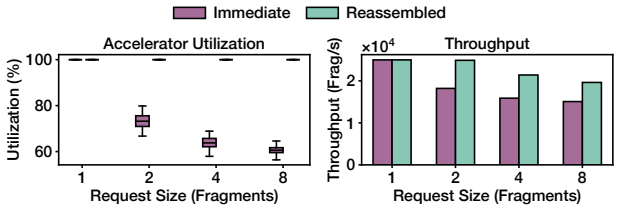


Figure 5. Comparison of immediate data ingestion at an accelerator vs. ingestion of reassembled requests.

We validate the impact of request reassembly on accelerator utilization and multi-client support with two simulations. We align our setup with the function characteristics in Table 1, modeling an accelerator that processes requests composed of 1, 2, 4, or 8 fragments, with each fragment requiring 20 μs to process. Fragments belonging to the same request arrive with a random inter-fragment delay of 30–40 μs. The 20 μs processing time is representative of Table 1, and the slightly longer inter-fragment delay exposes invocation bubbles. Figure 5 (left) shows that the utilization of the accelerator, i.e., the time spent processing requests rather than stalling, decreases as request size increases when fragments are fed straight into the accelerator (Immediate in the plot). In contrast, enabling request reassembly (Reassembled) significantly improves utilization by eliminating stalling, since the accelerator only begins execution after receiving a complete request.

Eliminating stalling has the added benefit that an accelerator can service more requests from distinct clients. We

validate this with an experiment with two clients. Without request reassembly, any new request arriving from one client while the accelerator is busy serving a request from the other is dropped. Figure 5 (right) shows that as request size increases, immediate ingestion sees decreasing throughput as the accelerator spends 25–40% of its time stalled, while the design with reassembly can service queued complete requests back-to-back without stalling. Further analysis is in Appendix E.

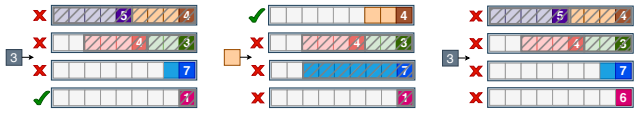
Hence, by presenting accelerators with only complete requests using reassembly buffers, fRAC improves both accelerator utilization and request throughput in multi-tenant settings.

3.2 Designing an Efficient Reassembly Buffer

We first ask *how should reassembly buffers be coupled with accelerators?* An intuitive approach is to adopt a per-accelerator reassembly buffer similar to the per-core shallow queue in [47]. However, this approach can lead to Head-of-Line (HoL) blocking when there are multiple clients issuing requests to the same accelerator. A single reassembly buffer can only reassemble one request at a time to maintain the contiguous order of fragments that constitute requests. This significantly limits scalability and leads to under-utilization of accelerators. Furthermore, it would be necessary to allocate suitable buffer space per accelerator based on the properties of the accelerator (e.g., request size distribution and execution time) and expected demand.

Rather than statically allocating buffers to specific accelerators, fRAC decouples them, enabling reassembly to be optimized independently of accelerator and request size mix, as is necessary for a virtualized system where the specific accelerator loaded in a slot might differ over time. As depicted in Figure 4, each accelerator has a separate queue that complete requests are forwarded to after reassembly, and these are processed atomically as soon as the accelerator becomes available.

fRAC instantiates multiple reassembly buffers that operate as follows. The first fragment of a new incoming request is dispatched to eligible buffers based on a policy, explored in Section 3.4. A buffer is *eligible* to accept the first fragment of a new request only if it is not currently reassembling a request, and only if the size of the new incoming request is no larger than the remaining space in the buffer. Once the first fragment of a request has been admitted to a buffer, that buffer is now allocated solely to the remaining fragments of that request; it is *ineligible* to serve any other request. Once a buffer completes assembling a request, it becomes eligible to accept a new request once more. If there are no eligible buffers, the incoming request fails (as with the request's remaining fragments). We omit the transport stack in our simulation to isolate the fRAC abstraction. In a full deployment, these "failed" fragments are absorbed by the reliable transport layer (e.g., TCP) and seamlessly trigger network retransmissions, rather than causing application-level failures. Assembled requests are held in the buffer only until they can be sent to accelerator queues.



(a) New request. **(b)** In-flight. **(c)** Failed request.

Figure 6. Behavior of the Reassembly Buffers. (a) The first fragment of each request identifies its size, and can be accepted only into an eligible buffer that is not reassembling another request and has sufficient capacity. (b) Subsequent fragments of a request are steered to the corresponding buffer. (c) If a new request does not find an eligible buffer, it fails. *Legend: Fully assembled requests are hatched. Eligible buffers in each scenario are marked with a tick and ineligible buffers with a cross.*

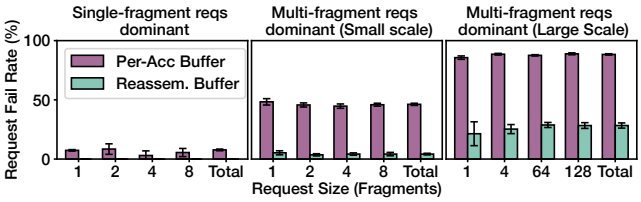


Figure 7. Request fail rates with per-accelerator buffers vs. a separate reassembly buffer.

Figure 6 shows fRAC's behavior in response to buffer states and incoming fragments. These policies are simple enough to be enforced at the hardware level using per-clock status signals. However, fRAC relies on the underlying transport protocol to guarantee in-order data delivery. If a connection is interrupted before a request is fully reassembled, fRAC must perform garbage collection to discard incomplete fragments from the reassembly buffer.

We simulate per-accelerator and independent reassembly buffers when servicing diverse requests in a setup with four buffers and three accelerators. Incoming requests are allocated in a round robin manner to *eligible* buffers. Three different workload mixes are used to explore the impact on request fail rate under a high saturating load: single-fragment requests dominate, and multi-fragment requests dominate (evaluated at both small and large request scales). Figure 7 shows that decoupling reassembly buffers from accelerators significantly reduces the request fail rate, dropping by approximately 40% and 60% for small-scale and large-scale multi-fragment workloads, respectively.

3.3 Dealing with Small Requests

Next we ask *should single-fragment requests be buffered along with larger requests?* Serverless workloads exhibit high request size variability (cf. Table 1). Small single-fragment requests do not require reassembly. However, if forced to

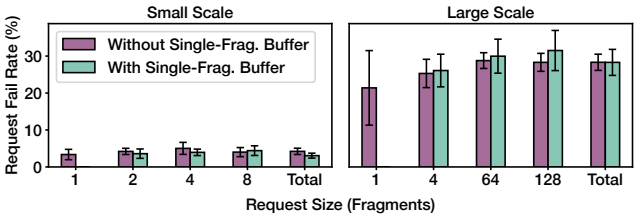


Figure 8. Request fail rates with and without a single fragment buffer.

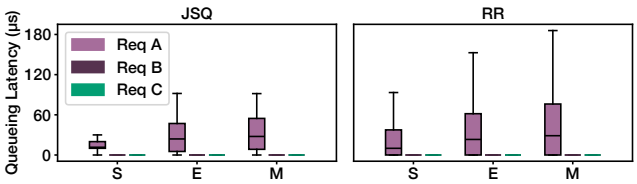


Figure 9. Queueing latency for accelerators A, B, and C under RR and JSQ for three request-size distributions: S (single-fragment dominant), E (even mix), M (multi-fragment dominant).

share buffers with larger requests that can occupy reassembly buffers for longer periods, they can suffer higher response latency, and the potential for request drops. Creating a distinct buffer for single-fragment requests allows them to be queued without interference by larger requests, ensuring accelerators can service them in the gaps between larger requests being reassembled. Figure 8 shows that adding a single-fragment buffer completely eliminates failures for single-fragment requests, solving request-level HoL blocking while still maintaining competitive request fail rates for multi-fragment requests.

3.4 Allocating Requests to Buffers

We now ask *how should a new incoming request be allocated to a specific reassembly buffer?* A naive round robin allocation between all buffers (Naive RR) would be problematic since buffers already allocated to an incomplete request are ineligible. Hence, fRAC adopts a simple variant, where scheduling is performed only among eligible buffers (Eligible RR). While more fine-grained strategies exist, like joining the shortest available buffer (JSB), fRAC favors Eligible RR for its hardware simplicity: it requires only a flag bit per buffer, whereas JSB would involve more complex comparisons between dynamic buffer states. Moreover, our experiments (details omitted for brevity) demonstrate that Eligible RR achieves availability comparable to JSB.

3.5 Allocating Requests to Multiple Accelerator Instances

Even with high accelerator utilization, an accelerator with long execution time and a high request volume can be oversaturated. To support scaling, fRAC supports hosting multiple instances of the same accelerator. Ideally, optimal buffer utilization could be achieved by sharing a single queue among multiple accelerator instances of the same type, allowing completed requests to be load-balanced across instances. However, this approach limits the framework's flexibility, as hardware queues must be pre-allocated, creating a tight coupling between accelerator types and queue instantiation. To avoid this constraint, fRAC decouples queues from accelerator types and maintains distinct request queues for each accelerator slot. Hence we ask *how should assembled requests be allocated to multiple queues for the same accelerator type?*

We simulate two policies: Round Robin (RR) and Join Shortest Queue (JSQ), which selects the accelerator queue with the least queued data. We simulate a workload with 95% of requests for accelerator A and instantiate 3 instances of accelerator A and run three request size distributions (cf. Appendix A, Table 5).

Figure 9 shows the queueing latency distribution for accelerators A, B, and C. While simple RR achieves comparable median latency to JSQ, it has higher worst-case latency and higher tail latency. This stems from RR assigning complete requests blindly in a fixed order, without considering request size, whereas JSQ accounts for the actual queue load. In fRAC, we adopt RR for its hardware simplicity, accepting its tolerable trade-off in tail performance in exchange for efficient and practical implementation.

3.6 On-Chip Buffering

A key design decision in fRAC is to use on-chip memory to implement the reassembly buffers and accelerator queues. This significantly reduces the overhead of reassembly and accelerator invocation, but also means the sizes of these buffers are constrained by limited on-chip memory. Modern FPGAs offer 10s of MB of on-chip memory (about 9MB of BRAM and 35MB of URAM in the case of the AMD Alveo U280 used in our prototype [3]). However, only a fraction of this (a few MBs) is available to fRAC, depending on how much of the FPGA is allocated to function acceleration. At the highest level, if an accelerator has an execution time of t_{Acc} and the input bandwidth is BW, then the buffers in aggregate should be able to buffer a data size of $BW \times t_{Acc}$. Assuming t_{Acc} of 100 μs and BW of 100 Gbps, that equates to 1.25 MB of buffer space.

In fRAC, we can use more resources for the reassembly buffers or the accelerator queues. Shorter accelerator queues mean complete requests may stay in reassembly buffers longer, causing fewer of them to be eligible, potentially resulting in request failures. Longer accelerator queues can result in underutilized queue space when fewer accelerators are active or being requested. Considering the ability to dynamically modify accelerators hosted in slots to balance incoming traffic demands, we choose to ensure accelerator queues can host multiple complete requests, while we size reassembly buffers to accommodate the size of the largest request plus some additional margin.

3.7 Accelerator Virtualization

To support accelerator virtualization, fRAC uses Partial Reconfiguration [4] to replace the accelerator loaded in a slot at runtime. This process is managed by a reconfiguration controller that coordinates slot reconfiguration on behalf of the control plane. As illustrated in Figure 10, when a reconfiguration request arrives, the controller first waits until the accelerator currently assigned to the target slot becomes idle. It then prevents the request selector from dispatching new requests to that slot, asserts the decoupling signals, and reconfigures the slot by streaming bitstream for the requested accelerator to the Internal Configuration Access Port (ICAP) [75]. Once reconfiguration succeeds, the controller updates the Selector's accelerator-to-slot allocation table so that subsequent requests for the newly loaded accelerator are dispatched to that slot. It then de-asserts the decouplers and reports the reconfiguration status back to the control plane. If an error occurs during reconfiguration, the slot remains decoupled and the control plane is notified of the failure. This allows fRAC to support transparent accelerator hot-swapping while requests targeting other slots continue to execute without interruption.

3.8 Putting It All Together

The design of fRAC builds atop an established network transport layer, leveraging ordered delivery of packet payloads to allow reassembly of application-level requests and accelerator invocation at that granularity. Reassembly buffers combine request fragments into complete requests, which are accompanied by required invocation parameters. These buffers are agnostic to the requesting client, the requested accelerator, and request size to provide high utilization in a dynamic workload scenario. A separate buffer for small single-fragment requests ensures they are serviced reliably even in the presence of larger requests. The Dispatcher implements the buffer joining policies outlined above, which have minimal overhead. Furthermore, the Reconfiguration Controller enables dynamic swapping of accelerators via Partial Reconfiguration, allowing the system to adapt to changing workloads at runtime without interrupting the network path. The complete fRAC architecture is shown in Figure 10.

4 Prototype Implementation

Our open-source[3] hardware prototype of fRAC is built on an AMD Alveo U280 FPGA with 100G network interfaces. The full request path is implemented within the FPGA. Packets enter through the 100G Ethernet interface, pass through the hardware TCP stack, are assembled by fRAC into complete requests, dispatched to the target accelerator, and returned as responses through the same network path.

Network Stack. We implement our prototype in Verilog, building atop the open-source TCP/IP stack in [29] for network transport. Since fRAC is a hostless framework, all transport-layer functionality is implemented directly in hardware. The hardware transport stack manages the complete TCP connection lifecycle, including establishment, per-connection transport state, reliable delivery, and teardown. No host participates in connection management. We choose this TCP/IP stack because it is reliable and easy to extend with custom logic. However, fRAC is transport-agnostic, meaning it can integrate with any protocol that guarantees packet ordering. Our design, depicted in Figure 10, is implemented in the *user kernel* portion of the network stack design. TCP headers of individual packets are processed by the TCP/IP stack, which extracts the connection ID and passes it along with packet payloads as the fragments to fRAC.

[3]fRAC codebase: https://github.com/accl-kaust/frac

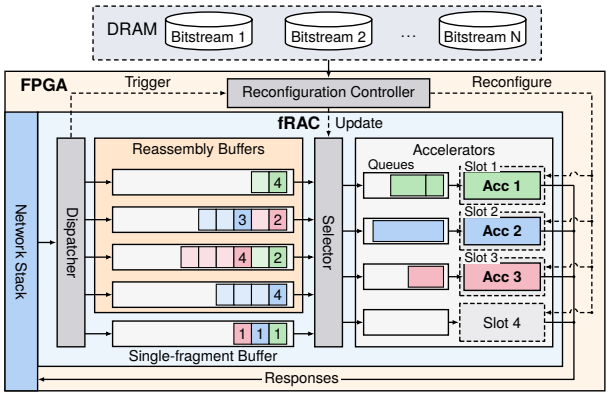


Figure 10. fRAC design showing packet payloads passed to the Dispatcher as fragments which are then reassembled in buffers based on information in the request headers. Assembled requests are then passed to the relevant accelerator and responses returned to the network stack. Accelerators can be replaced by the Reconfiguration Controller.

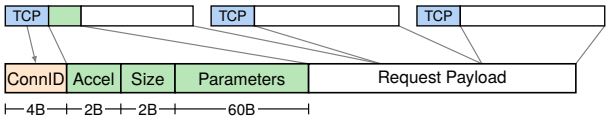


Figure 11. Request format showing the payloads of multiple TCP packets assembled into a request with its request header.

Request Format. Figure 11 illustrates the per-request header and its placement within TCP segments. We define a small 64B custom *request header* that is embedded in the first fragment in a request. This header identifies the `Accelerator` (2B) and the request `Size` (2B), plus 60B as the `Parameters` field carrying accelerator-specific parameters, e.g., the value K for the Top-K accelerator. `Parameters` is opaque to fRAC, which passes it to the accelerator with the payload data. The size of this field (60B) is standardized to simplify our implementation, but other designs are possible.

Reassembly Buffers. The reassembly buffers use the connection ID and `Size` in the first fragment of a request to determine how many fragments to assemble into a complete

request. The reassembly buffers are implemented as First-In-First-Out buffers (FIFOs) utilizing on-chip BRAM. This allows for minimal buffering overhead and low latency. The Dispatcher implements the buffer allocation policy to assign incoming fragments to reassembly buffers. The Selector maintains a table of the currently instantiated accelerator in each slot and monitors all reassembly buffers in every cycle to identify completely reassembled requests that are then forwarded to the corresponding accelerator queues. Requests are passed to the correct accelerator based on the `Accelerator` field in the header, along with the `Parameters`.

Our prototype implements four 4096×512b (0.25MB) reassembly buffers and an additional 16384×512b (1MB) single-fragment buffer. Each accelerator queue is 1024×512b (64KB), which significantly exceeds the largest request payload we support. These buffers can be increased in size in a full deployment as modern FPGAs have more available on-chip memory in the 10s of MBs, especially when including URAM resources. While some of these must be retained for accelerator implementation, there is ample availability for buffers an order of magnitude larger than those in this proof of concept.

Accelerators. We implement a variety of accelerators (A1–A4 in Appendix B, Table 6), and consider additional accelerators from the literature (A5–A7 in Table 6) when reasoning about partial reconfiguration and slot sizing.

• *Top-K (A1)*: processes a window of data, sorting and returning the highest K integers, with K being a dynamically set parameter per request (implemented based on [35]). This operation is commonly used in data analytics for ranking tasks, such as extracting the most significant keywords or trending terms, as demonstrated in [74].

• *Logit Transform (A2)*: applies a logit transformation over a floating point tensor block, performing subtraction, division, and a logarithm operation in a pipeline. This transformation plays a vital role in pre-processing within Deep Learning Recommendation Models (DLRMs), converting raw input features into normalized probability distributions for personalized recommendation predictions [60].

• *Min-Max Normalization (A3)*: normalizes a floating point tensor block, scaling each element by the range of the block with pipelined range determination, subtraction, and division. This is a crucial pre-processing step within DLRMs, standardizing raw input features into a consistent [0, 1] range to ensure uniform scaling across diverse data types [60].

• *CNN (A4)*: A quantized neural network accelerator designed for image classification using multiple convolutional layers (filters from 8 to 32 elements), max pooling, and dense layers, trained on the SVHN dataset via hls4ml [26]. It processes 24KB input images (64x64 48-bit/pixel) and produces a one-hot encoded 10-class output. This represents a computationally intensive application requiring request reassembly in all cases and having a much longer execution time.

Partial Reconfiguration. We must first define a set of Partially Reconfigurable Regions (PRRs) or slots, which must be spatially arranged subject to constraints relating to the underlying FPGA resources. To maximize the area available to accelerators, we must constrain the network stack and fRAC infrastructure portions to parts of the FPGA that include the required I/O, while leaving as much area available to accelerator slots as possible. This is highly challenging in a design that must achieve timing closure to function correctly with the I/O interfaces.

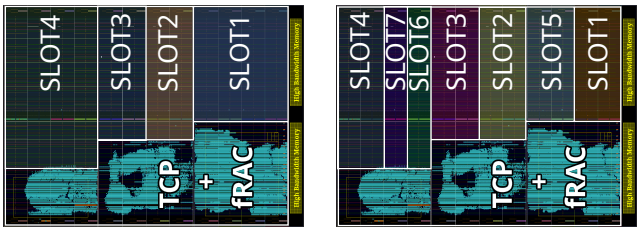
Since our design is hostless, we are able to remove hardware devoted to host PCIe communication, eliminating the XRT shell [2] and freeing additional FPGA area for accelerators. This optimization reduces the stack's resource usage by about 45% compared to the original design. Table 3 presents the resource usage of the original and optimized TCP/IP stack as well as fRAC.

The buffer allocations (including reassembly buffers and single-fragment buffers) are aligned with the size mentioned above. Figure 12 shows two example floorplans we designed with 4 or 7 accelerator slots to show slot configurations for different accelerators with different reconfiguration times, while respecting spatial layout constraints. Table 2 shows the resources available in each of the slots for the 4-slot configuration. Cross-referencing against the accelerators in Table 6 in Appendix B, we see that this arrangement can support a wide range of different accelerators being instantiated at

Table 2. PR slots and reconfiguration times (ms).

Slot	LUT	FF	BRAM	DSP	URAM	Time
1	213k	429k	288	1440	192	6.90
2	120k	245k	168	912	128	4.12
3	120k	245k	168	912	128	4.12
4	319k	641k	480	2304	256	9.92

Slots 1,4: all accelerators; Slots 2,3: A1–A3, A5–A7.



(a) 4-slot config. (b) 7-slot config.

Figure 12. Accelerator slot layouts on AMD Alveo U280.

Table 3. FPGA resources consumed by fRAC.

Component	LUT	FF	BRAM	URAM	DSP
Original TCP/IP Stack	260.4k (19.99%)	454.1k (17.43%)	443 (21.97%)	0 (0%)	0 (0%)
Optimized TCP/IP Stack	150.5k (11.5%)	130.2k (5.0%)	200.5 (10.0%)	16 (1.7%)	0 (0%)
fRAC Layer	5.1k (0.4%)	25.4k (1.0%)	250 (12.4%)	40 (4.2%)	0 (0%)

the same time, including more complex accelerators than we built for this prototype.

To enable fast accelerator switching, fRAC uses the high-throughput Internal Configuration Access Port (ICAP) [75]. When an accelerator is loaded, the Selector's slot allocation table is updated. We measured actual reconfiguration times, reported in Table 2, by issuing a request to reconfiguration controller and recording the time required for the controller to stream the corresponding accelerator bitstream from off-chip HBM to ICAP. The measured times are roughly 1–2 ms higher than theoretical estimates based on ICAP bandwidth, due to occasional stalls introduced while reading the bitstream. Compared to the 100ms to 1s cold-start delays commonly seen in software-based systems [9], fRAC can achieve reconfiguration in just a few milliseconds. This design could be extended further using nested partial reconfiguration [5] to allow the number and sizes of slots to be modified at runtime.

Integrating New Accelerators. A key strength of fRAC is that integrating new accelerators is straightforward. A 512-bit AXI-Stream interface is provided for data input and output. A simple wrapper module can translate these streams into the required AXI-Stream data width of an arbitrary accelerator through FIFOs. Another 32-bit AXI-Stream interface is provided for metadata input and output.

The accelerator is passed the 64B request header as the first piece of data, followed by the remainder of the payload. Depending on the specific design of the accelerator, it may not require any of this information, in which case it can process the following stream. It may use the Size field to configure internal execution, such as the number of iterations of a computation to perform, and it may use the custom-defined Parameters field to configure other datapath options, such as the value of K for a Top-K accelerator, the datatype for an arithmetic kernel, the matrix dimensions for a matrix multiplication, or the kernel for an image filter. The specification of these must, of course, be communicated to the clients so they can issue valid requests. The accelerator produces its output data stream and must indicate the size of the response on the last beat using the metadata stream. Some streaming accelerators may require additional logic to flush the pipeline after the completion of input ingestion to fit the run-to-completion model for a single request. The wrapper ensures the connection ID is passed with the response to route to the correct client.

Consider a CNN accelerator generated with hls4ml [26]. Both input and output interfaces use AXI-Stream ready/valid handshakes. It ingests one 48-bit pixel every 223 cycles due to resource-shared pipelines. After reading 4096 inputs (64×64 pixels), it produces the classification output 913,408 clock cycles after the first input. To integrate it into fRAC, we first assign it an accelerator ID to be used by the Selector. We wrap the module with a FIFO-based adapter that bridges the 512-bit fRAC stream to the 48-bit interface of the IP. A counter in the wrapper ensures sufficient input padding is applied to flush the outputs for each complete input. Simpler accelerators such as Logit Transform and Min-Max Normalization exhibit more straightforward behavior and do not require flushing. As long as an accelerator exposes a streaming interface and its execution semantics are understood, it can easily be wrapped to integrate into the fRAC framework without changes to its core logic. Notably, this CNN accelerator embeds all weights and biases in the logic using on-chip memory. While the accelerators we implemented do not use off-chip memory, connectivity to external HBM channels through AXI interconnect is supported. Table 6 in Appendix B shows the hardware implementation results of the accelerators we implemented (A1–A4) as well as additional accelerators from the literature (A5–A7), which we use only as comparative reference designs.

fRAC also provides abstract shells for the layouts shown in Figure 12, allowing new accelerators to be synthesized and implemented without rebuilding the full fRAC stack. This can save hours of vendor-tool compilation time during accelerator development. In contrast, under a typical PR flow, adding a new accelerator would require regenerating all accelerator bitstreams and updating the bitstream database stored in local memory. The Abstract Shell flow [6] avoids this restriction, allowing a new accelerator to be implemented without invalidating the existing static design or previously generated accelerator bitstreams. This separation also gives cloud infrastructure maintainers a practical way to support accelerator development, since they can provide the abstract shells without exposing the complete fRAC design.

5 Evaluation

We explore whether fRAC: (i) *exhibits minimal overhead over the transport layer it builds upon, and sustains low latency under high throughput conditions;* (ii) *handles mixed accelerator requests while ensuring performance isolation;* (iii) *provides efficient reassembly to increase accelerator throughput;* (iv) *outperforms optimized software implementations in scalability and delivers more consistent latency;* (v) *maintains stable performance under a real-world serverless trace.* A full evaluation artifact for these experiments is available.[4]

Choice of Baselines. We do not include mainstream serverless platforms such as OpenFaaS [64] and AWS Lambda [1] in our quantitative comparison as their end-to-end invocation latencies are typically in the millisecond range [36], dominated by software platform overheads. Similarly, host-controlled FPGA frameworks have high overhead. Molecule [22] incurs 50–100 μs to DMA just 4KB between CPU and FPGA. F3's FPGA API call overhead is 10.8 μs [54], over an order of magnitude higher than fRAC's 0.16 μs framework overhead in Figure 13. λ-NIC reports median latencies above

[4] https://zenodo.org/records/23200153

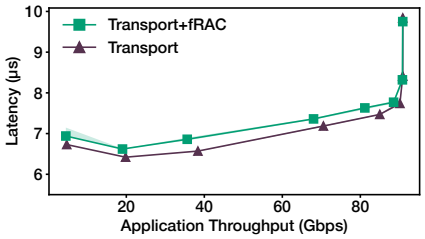


Figure 13. Median latency vs application throughput with/without fRAC (shaded 25–75th pct.).

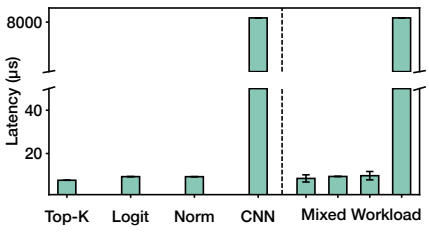


Figure 14. Comparison of latency between isolated baselines and a workload with mixed requests.

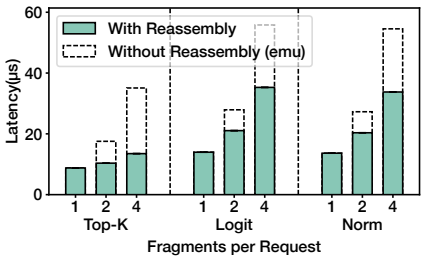


Figure 15. Impact of request size on latency.

50 μs even for simple workloads [15]. We adopt more demanding baselines by comparing against a highly optimized DPDK-based Libtpa server and a DPU implementation.

Testbed and Experimental Setup. Clients run on a machine with AMD EPYC 7763 64-core processor, 512 GB DDR4 RAM, NVIDIA ConnectX-6 NIC, running Ubuntu 20.04. The client uses the Libtpa TCP stack based on DPDK (see Appendix C). We heavily adapted Libtpa [12] for our evaluations. Each client is pinned to a hardware thread, all within the same NUMA node connected to the NIC. fRAC is deployed on an AMD Alveo U280 FPGA. The FPGA receives requests from its QSFP network interfaces, processes them, and returns responses back to the client directly. For the CPU comparison, we use another machine with identical configuration running a DPDK-based server using Libtpa. For the DPU baseline, we deploy an NVIDIA BlueField-3 SmartNIC [62] in off-path mode with identical configuration, as off-path is the only mode that supports general-purpose application offloading and all function processing runs solely on the DPU Arm cores, with no involvement of the host CPU. All interfaces are connected to an EdgeCore DCS810 switch at 100 Gbps, with an MTU configuration of 9000B. All fRAC experiments in this section run on the Alveo U280 FPGA. Persistent TCP connections are established before measurement, so latencies exclude connection establishment and teardown. We use saturation experiments to expose throughput ceilings and queueing behavior, and trace replay to evaluate time-varying production workloads.

Low Invocation Latency. We first evaluate the latency overhead of the fRAC abstraction under increasing throughput. We implement an echo function directly atop the underlying TCP stack, and as a function with the fRAC abstraction. For this experiment only, the client and the FPGA are connected directly with a cable without a switch, so that the measured latency reflects the limit of the fRAC prototype itself rather than switch-induced delay. We compare the latencies under increasing load with a varying number of clients. Each client opens a TCP connection and transmits single-fragment requests with a 4096B payload for 30 seconds in a closed loop, meaning each client issues its next request only after receiving the response to its previous request, keeping at most one request outstanding. The first 10 seconds of results are discarded. With 24 concurrent clients, we achieve approximately 90 Gbps of application throughput (i.e., considering only the payload). The 24-client count is simply an artifact of the number of clients needed to reach maximum throughput in this high-load configuration. As shown in Figure 13, fRAC introduces minimal overhead, adding only about 0.16 μs compared to echoing data in the transport layer. Even at peak load, overall latency remains below 8 μs, demonstrating excellent *scalability* under saturating network conditions.

Performance Isolation. To demonstrate that fRAC suitably isolates requests for different accelerators, even when their execution times vary significantly (cf. Table 1 and Table 6), we deploy four accelerators (Top-K, Logit Transform, Min-Max Normalization, and CNN) in separate slots. We first conduct a single-client test for each accelerator to observe its baseline latency in isolation. To evaluate performance isolation under contention, we next run a mixed workload with four concurrent clients each sending requests exclusively to one of the four for five seconds. Logit Transform and Min-Max Normalization clients send single-fragment requests of 1024B, while CNN workload requests are six 4096B fragments, forming one input image. Figure 14 shows that latency remains consistent for isolated and mixed workloads, confirming that performance isolation is maintained even when accelerators have different execution times. Error bars may not be visible due to low variance.

Accelerator Utilization. We extend the previous experiment with clients sending requests of 1, 2, and 4 fragments (each 1024B) in a closed loop to demonstrate the benefit of request reassembly. We measure latency from last fragment sent to response received. Run to completion means processing multi-fragment requests requires multiple invocations without reassembly. Figure 15 shows that, with reassembly, latency increases sub-linearly with request. For Logit Transform and Min-Max Normalization, an assembled 2-fragment request takes 76% of the time of processing its fragments separately, and a 4-fragment request 64%. For a lightweight workload like Top-K, where communication overhead exceeds processing time, the advantage is larger. Error bars may not be visible due to low variance.

Scalability Comparison with CPU and DPU. To demonstrate that fRAC provides better scalability and more stable

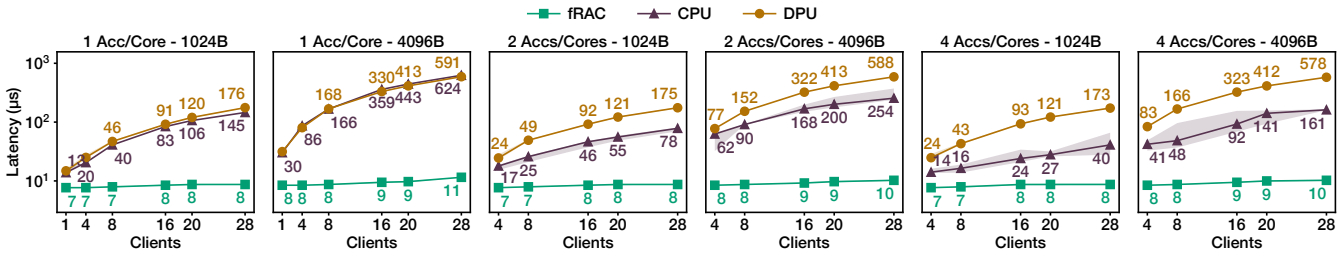


Figure 16. Latency comparison for the Top-K function on fRAC, CPU, and DPU, showing median, 25th, and 75th percentiles.

latency than function execution on CPU or DPU, we use high-performance DPDK Libtpa (Appendix C) as a baseline. DPDK is not a standard stack for serverless applications; most frameworks typically exhibit end-to-end latencies ranging from hundreds of μs to tens of ms [36], as demonstrated in host-managed offload [7, 22, 54]. However, this more demanding baseline highlights the performance gains achieved by eliminating host-side control in fRAC. We instantiate fRAC with 1, 2, and 4 Top-K accelerators and compare against function invocation on server CPU and the Arm cores of the NVIDIA DPU, both using Libtpa (Appendix C), pinned to 1, 2, and 4 CPU cores on each platform for fair comparison, ensuring equivalent parallel processing. We scale the number of clients up to 28 to saturate network bandwidth (this is the minimum required to reach maximum DPDK throughput in our closed-loop setup, fRAC can scale to hundreds of clients). Each client sends single-fragment requests of 1024B or 4096B. While this workload does not exercise request reassembly, it serves to highlight latency differences among the platforms. Figure 16 shows that fRAC reduces latency by 46–95% for 1024B requests and 73–98% for 4096B requests compared to the CPU, and 53–95% and 73–98% compared to the DPU. Latency on the CPU and DPU increases linearly with the number of clients. In contrast, fRAC latency remains nearly flat regardless of the number of clients and accelerators: for example, with a single accelerator and 4096B requests, latency rises only from 8 to 11 μs from 1 to 28 clients, while CPU latency grows from 30 to 624 μs. Figure 17 shows fRAC's much lower tail latency under full load (28 clients). We also ran the CNN workload on fRAC and the CPU server: each 24 KB request (six 4096B fragments) takes 8.6 ms on fRAC versus 40 ms with C-compiled TensorFlow on the CPU.

Power Consumption. Serving the same CNN client load, the CPU server and NIC consume 188–250 W from idle to full load (104–110 W from the CPU alone), while the hostless fRAC prototype consumes a total of 28–31 W, a 6–8× reduction in platform power.

Realistic Workload. We use the Azure Function Invocation Trace [81], which reports two weeks of serverless function invocations. We select the one-hour interval with the highest invocation density and replay it as a sequence of accelerator requests. Pre-processing applied to the trace for our experiments is described in Appendix D. Figure 18 shows the resulting latency CDF per accelerator (color) and request size in fragments (line style), demonstrating that fRAC maintains stable latency under a real-world workload.

Table 4. Comparison of fRAC with existing work. *TP*: hardware transport-protocol support (*Partial*: NIC offloads or a custom/RPC transport, not a complete in-hardware stack). *RG*: request-level granularity (*Specific*: application-specific messages, not a general request abstraction). *MT*: multi-tenancy. *AO*: accelerator offload (*Packet*: packet-level network functions only; *Static*/*Dynamic*: application accelerators without/with runtime reconfiguration). *HD*: hostless deployment.

	TP	RG	MT	AO	HD
RingLeader [47]	✗	✗	✗	✗	✗
PANIC [48]	✗	✗	✓	Packet	✗
SuperNIC [49]	✗	✗	✓	Packet	✗
ClickNP [44]	✗	✗	✗	Packet	✗
Beehive [46]	✓	✗	✗	Packet	✓
TOEs [57]	✓	✗	✗	✗	✗
Pigasus [82]	✓	✗	✓	Static	✗
HFT [43]	✓	Specific	✗	Static	✗
BlastFunction [7]	✗	✗	✓	Static	✗
Molecule [22]	✗	✗	✓	Static	✗
F3 [54]	✗	✗	✓	Dynamic	✗
Harmonia [45]	✗	✗	✓	Static	✗
Ensō [70]	✓	✗	✓	✗	✗
nanoPU [34]	✓	✓	✓	✗	✗
RpcNIC [80]	Partial	✓	✓	Dynamic	✗
Split gRPC [68]	✗	✓	✓	✗	✗
FlexDriver [23]	✗	✓	✓	Static	✗
NEBULA [73]	Partial	✓	✓	✗	✗
Arca [72]	✗	✓	✓	✗	✗
Rain [53]	✓	✓	✓	✗	✗
Coyote v1 & v2 [41, 67]	✓	✗	✓	Dynamic	✗
Strega [55]	✓	✓	✓	Static	✗
fRAC	✓	✓	✓	**Dynamic**	✓

6 Comparison with Related Work

Table 4 summarizes how fRAC compares to prior works across hardware transport support, request granularity, multi-tenancy, accelerator offload, and hostless deployment.

Packet-level SmartNICs. RingLeader [47] and PANIC [48] extend the Corundum [27] FPGA NIC for intra-server orchestration and limited network function offload, respectively. SuperNIC [49] offloads graphs of network tasks. ClickNP [44] provides modular network function offload without protocol

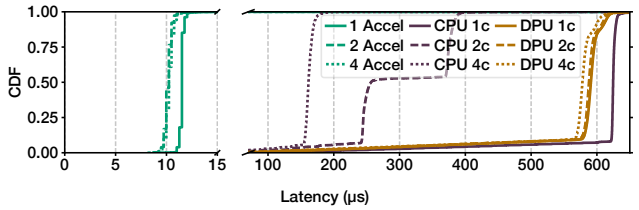


Figure 17. Latency CDF for saturating 4KB requests.

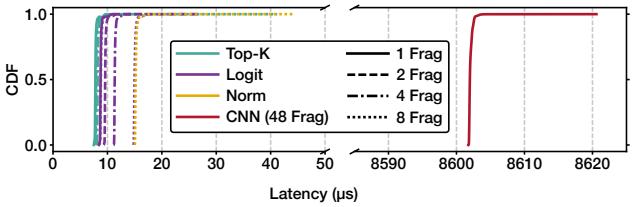


Figure 18. Latency CDF for the Azure trace replay.

support. All target datacenter traffic routing to host applications rather than application-level function acceleration.

Application-Specific Processing. High-Frequency Trading (HFT) accelerators on FPGA hardwire exchange-specific parsing and trading logic into the network datapath [43]. They optimize a fixed application pipeline, whereas fRAC routes flexible requests between clients and accelerators.

Transport Offload and Stream Processing. TCP Offload Engines (TOEs) move connection management and in-order byte-stream reconstruction from the host CPU to the NIC [57]. Pigasus [82] couples per-flow TCP reassembly to a fixed intrusion detection pipeline. In contrast, fRAC inserts a request-level hardware shim between the transport and accelerators. A shared, connection- and accelerator-agnostic buffer pool decouples data ingestion from execution, routing complete requests to virtualized, reconfigurable accelerators.

Host-Managed FPGA Virtualization. Frameworks managing data movement and reconfiguration in software [14, 24] incur significant overhead for networked applications, as seen in BlastFunction [7], Molecule [22], F3 [54], and Harmonia [45]. Coyote [41, 67] is a heavyweight OS-FPGA hybrid system providing transport support and reconfigurable accelerators. Strega [55] extends EasyNet [29] for request-level HTTP calls but remains host-managed and static. RpcNIC [80] runs reconfigurable SmartNIC kernels but relies on the host CPU for control and scheduling. Split gRPC [68] partitions RPC processing between the host and SmartNIC for isolation, but retains host-side execution and orchestration.

Streaming and Request-level Processing. Ensō [70] and nanoPU [34] reassemble packets for efficient NIC-to-software streaming without accelerating the workload itself. Beehive [46] offloads transport protocols but lacks multi-tenancy and application-level abstractions. NEBULA [73] employs a stateless request/reply protocol eliminating network/PCIe overheads, but lacks accelerator offload. FlexDriver [23] supports request-level granularity but separates the network interface from accelerators, forcing data over PCIe. Rain [53] schedules microsecond-scale requests in-network via programmable switches, focusing on cluster-level load balancing rather than endpoint accelerator offload. Arca [72] executes fine-grained serverless functions on CPUs using continuations rather than invoking hardware accelerators.

In contrast, fRAC invokes virtualized, runtime-reconfigurable accelerators on complete requests, entirely within the FPGA and without a host.

7 Discussion and Conclusions

We have proposed the Remote Accelerator Call (RAC) paradigm, which decouples network data transfer from accelerator control through a request-level invocation abstraction. Our fRAC prototype demonstrates that network-facing FPGAs can directly form, dispatch, and launch accelerator work, completely bypassing host-software bottlenecks. We validated that this hostless approach achieves microsecond-scale latency and high throughput, even under heavy load, outperforming host-managed approaches at significantly reduced power consumption.

fRAC's design is intentionally modular: the number and size of reassembly buffers and accelerator slots can be tuned to different deployments. The simple per-buffer signaling scheme allows the Dispatcher to change how buffers are shared without modifying accelerator logic. On the client side, a lightweight RPC-style library hides transport details, allowing applications to seamlessly leverage fRAC's hardware interface. We also demonstrate that fRAC can reconfigure accelerator slots within a few milliseconds, allowing the set of deployed accelerators to change at runtime to adapt to evolving demand. We have demonstrated an extension to the request-level interface design to support chained function pipelines within and across fRAC nodes in [78], demonstrating scalability to more complex application graphs.

Exploring integration of fRAC with alternative transport protocols is a natural next step, enabling better coupled use of hardware buffer resources, and even lower latency. Developing a control plane that uses live hardware queue signals to dynamically scale accelerator instances on demand while balancing of load across instances would further improve workload adaptation. We also plan to extend fRAC to support more complex accelerators that rely on off-chip memory. In the longer term, we envisage an fRAC-like invocation abstraction as a hardened interface in future network-attached devices.

Acknowledgments

This publication is based upon work supported by King Abdullah University of Science and Technology (KAUST) under Award No. ORFS-CRG11-2022-5017. We thank Gianni Antichi and R. Srikant for helpful discussions, and the anonymous reviewers for their comments.

A Simulator Setup

Since compiling each hardware design iteration takes hours, we explore fRAC's key design decisions with a Python simulator built on SimPy, which produced all results in Section 3. It models independent clients issuing varied requests to one multi-accelerator fRAC instance, via three components:

Generator. The generator simulates clients issuing requests. We abstract away network transport, assuming that fragments arrive in order but that larger requests span several fragments in distinct packets. We use the Yahoo! Cloud Serving Benchmark (YCSB) [18] generator, configured by three properties. First, the mix of request types is set with CoreWorkload, mapping YCSB operations to accelerator requests. Second, request sizes follow one of three distributions, namely single-fragment dominant (S: 95%, 3%, 1%, and 1% of 1-, 2-, 4-, and 8-fragment requests), even mix (E: 25% each), and multi-fragment dominant (M: 10%, 30%, 30%, and 30%). Third, 100 requests are sent closed-loop, so a client issues a new request only after receiving the previous response. Fragments arrive with a random 5–15 μs delay, and a dropped fragment is retried after a random 80–120 μs delay to avoid herd behavior. Parameters are listed in Table 5.

Reassembly Buffer. This component reassembles request fragments into complete requests following the fRAC design in Section 3.2, with a configurable number of buffers.

Accelerators. This component emulates multiple accelerator types and instances, each optionally with its own input queue of complete requests. Each is modeled by its per-fragment execution time, which scales with input size.

B FPGA Resources and Performance

Table 6 summarizes the accelerators used in our evaluation and partial reconfiguration analysis. The latency of A1–A3 is measured by counting clock cycles in simulation with a 4096 B input, while A4 takes a fixed input size of 24 KB.

C Evaluation Setup

We benchmark fRAC and the CPU server baseline with Libtpa [12], a DPDK-based TCP stack for high-throughput tests, via its tperf application, which we extensively modified. Clients run closed-loop tests; multiple clients maximize throughput. The server assembles each connection's requests and invokes the function once all packets of a request have arrived, mirroring fRAC. The CNN workload uses the same model for the fRAC accelerator and the software function, which runs inference with the TensorFlow C library and AVX extensions; model parameters are loaded when a connection is established and CNN is requested, so later requests on a persistent connection skip model loading.

Table 5. Simulation parameters for each experiment. Mix, exec. time, and instances are per accelerator (A:B:C). PA: per-accelerator buffers; NRR/ERR: Naive/Eligible RR.

	Generator			Buffers		Accelerators		
	Cl.	Mix (%)	Sizes	No.	Policy	Policy	Exec. (μs)	Inst.
§3.2	8	33:33:33	S, E, M	3	PA, RR	–	10:10:10	1:1:1
§3.3	8	33:33:33	M	4	ERR	–	10:10:10	1:1:1
§3.4	4–8	33:33:33	M	4	NRR, ERR, JSB	RR	10/1000: 5/10:1	4:4:4
§3.5	8	95:2.5:2.5	S, E, M	8	ERR	RR, JSQ	10/1000: 5/10:1	3:1:1

Table 6. Accelerator resource and performance summary.

Function	Input	Resources					Performance	
		LUT	FF	BRAM	URAM	DSP	Clock (MHz)	Latency (μs)
Implemented								
A1 Top-K	1–32KB	9.1k	7.3k	0	0	0	200	0.41
A2 Logit Transform	1–32KB	3.1k	7.4k	0	0	6	200	5.4
A3 Min-Max Norm.	1–32KB	4.1k	6.2k	7.5	0	2	200	5.6
A4 CNN	24KB	41.9k	60.9k	79.5	140	219	200	8600
From the literature								
A5 MM [20]	32KB	60k	88.8k	23	–	640	200	80
A6 SVD [77]	128B	38k	57.8k	12	–	174	300	10
A7 PQC [8]	2KB	53k	28k	29	–	16	256	58

D Azure Trace Processing

The Azure Function Invocation Trace [81] contains two weeks of real-world serverless invocations. As it spans many functions, we cluster invocations into accelerator classes: k-means on invocations shorter than 1s yields three clusters, mapped to our Top-K, Logit Transform, and Min-Max Normalization accelerators; longer invocations map to CNN tasks. For request size granularity, each sub-second invocation is assigned 1, 2, 4, or 8 fragments according to its duration, using the 25th, 50th, and 75th percentiles of the corresponding cluster's duration distribution as the boundaries between the four bins; CNN invocations use 48 fragments of 512 B each, matching our CNN accelerator's fixed input size. Finally, we replay the densest hour, issuing requests at their inferred invocation times (end_timestamp minus duration).

E Theoretical Analysis

We estimate how long an accelerator idles when fragments are fed to it as they arrive (Immediate in Figure 5). Consider a request of k fragments separated by gaps $X_i \sim \mathrm{Exp}(\lambda)$, and a run-to-completion accelerator that spends τ on each fragment. If the accelerator is free when fragment 1 arrives, it idles for $(X_i - \tau)^+$ after each fragment $i < k$, which by memorylessness is $e^{-\lambda\tau}/\lambda$ on average. The total idle time is thus $\mathbb{E}[I] \le (k-1)e^{-\lambda\tau}/\lambda$, with equality for $k = 2$, and the request holds the accelerator for $I + k\tau$, a share $\eta = \mathbb{E}[I]/(\mathbb{E}[I] + k\tau)$ of which is idle. For $k = 2$, $\eta = e^{-\lambda\tau}/(e^{-\lambda\tau} + 2\lambda\tau)$, which is 33% for the mean gap (35 μs) and τ (20 μs) of Figure 5. Meanwhile, requests from other clients are dropped or wait behind the incomplete one. Reassembly removes this idle time: the buffers absorb the gaps, and each request holds the accelerator only for $k\tau$.

References

[1] Amazon Web Services. 2025. AWS Lambda. https://aws.amazon.com/lambda/ Accessed: 2025-11-27.

[2] AMD Inc. 2024. Xilinx Runtime (XRT). https://github.com/Xilinx/XRT Accessed: 2024-12-02.

[3] AMD Inc. 2025. Alveo U280 Data Center Accelerator Card Data Sheet (DS963). https://docs.amd.com/r/en-US/ds963-u280/Summary Accessed: 2025-06-21.

[4] AMD Inc. 2025. Dynamic Function eXchange. https://www.amd.com/en/products/adaptive-socs-and-fpgas/technologies/dynamic-function-exchange.html Accessed: 2026-06-09.

[5] AMD Inc. 2025. Nested Dynamic Function eXchange. https://docs.amd.com/r/en-US/ug909-vivado-partial-reconfiguration/Nested-Dynamic-Function-eXchange Accessed: 2026-06-09.

[6] AMD Inc. 2026. Abstract Shell Design Flow. https://docs.amd.com/r/en-US/ug909-vivado-partial-reconfiguration/Abstract-Shell-Design-Flow Accessed: 2026-06-09.

[7] Marco Bacis, Rolando Brondolin, and Marco D. Santambrogio. 2020. BlastFunction: An FPGA-as-a-Service System for Accelerated Serverless Computing. In *Design, Automation & Test in Europe Conference (DATE)*. 852–857.

[8] Luke Beckwith, Duc Tri Nguyen, and Kris Gaj. 2021. High-Performance Hardware Implementation of CRYSTALS-Dilithium. In *International Conference on Field-Programmable Technology (FPT)*. 1–10.

[9] James Beswick. 2021. Operating Lambda: Performance Optimization – Part 1. AWS Compute Blog. https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/ Accessed: 2025-03-07.

[10] Marcel Blöcher, Lin Wang, Patrick Eugster, and Max Schmidt. 2021. Switches for HIRE: Resource Scheduling for Data Center In-Network Computing. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 268–285.

[11] Matthew Burke, Sowmya Dharanipragada, Shannon Joyner, Adriana Szekeres, Jacob Nelson, Irene Zhang, and Dan R. K. Ports. 2021. PRISM: Rethinking the RDMA Interface for Distributed Systems. In *ACM Symposium on Operating Systems Principles (SOSP)*. 228–242.

[12] ByteDance. 2024. Libtpa: A DPDK-based userspace TCP stack. https://github.com/bytedance/libtpa Accessed: 2025-03-07.

[13] Adrian M. Caulfield, Eric S. Chung, Andrew Putnam, Hari Angepat, Jeremy Fowers, Michael Haselman, Stephen Heil, Matt Humphrey, Puneet Kaur, Joo-Young Kim, Daniel Lo, Todd Massengill, Kalin Ovtcharov, Michael Papamichael, Lisa Woods, Sitaram Lanka, Derek Chiou, and Doug Burger. 2016. A Cloud-Scale Acceleration Architecture. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1–13.

[14] Fei Chen, Yi Shan, Yu Zhang, Yu Wang, Hubertus Franke, Xiaotao Chang, and Kun Wang. 2014. Enabling FPGAs in the Cloud. In *ACM International Conference on Computing Frontiers (CF)*. 3:1–3:10.

[15] Sean Choi, Muhammad Shahbaz, Balaji Prabhakar, and Mendel Rosenblum. 2020. λ-NIC: Interactive Serverless Compute on Programmable SmartNICs. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*. 67–77.

[16] Kuan-Yu Chou and Yon-Ping Chen. 2020. Real-Time and Low-Memory Multi-Faces Detection System Design With Naive Bayes Classifier Implemented on FPGA. *IEEE Transactions on Circuits and Systems for Video Technology* 30, 11 (2020), 4380–4389.

[17] Ryan A. Cooke and Suhaib A. Fahmy. 2020. Characterizing Latency Overheads in the Deployment of FPGA Accelerators. In *International Conference on Field-Programmable Logic and Applications (FPL)*. 347–352.

[18] Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking Cloud Serving Systems with YCSB. In *ACM Symposium on Cloud Computing (SoCC)*. 143–154.

[19] Daniel Crankshaw, Xin Wang, Giulio Zhou, Michael J. Franklin, Joseph E. Gonzalez, and Ion Stoica. 2017. Clipper: A Low-Latency Online Prediction Serving System. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 613–627.

[20] Johannes de Fine Licht, Maciej Besta, Simon Meierhans, and Torsten Hoefler. 2021. Transformations of High-Level Synthesis Codes for High-Performance Computing. *IEEE Transactions on Parallel and Distributed Systems* 32, 5 (2021), 1014–1029.

[21] Johannes de Fine Licht, Grzegorz Kwasniewski, and Torsten Hoefler. 2020. Flexible Communication Avoiding Matrix Multiplication on FPGA with High-Level Synthesis. In *ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*. 244–254.

[22] Dong Du, Qingyuan Liu, Xueqiang Jiang, Yubin Xia, Binyu Zang, and Haibo Chen. 2022. Serverless Computing on Heterogeneous Computers. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 797–813.

[23] Haggai Eran, Maxim Fudim, Gabi Malka, Gal Shalom, Noam Cohen, Amit Hermony, Dotan Levi, Liran Liss, and Mark Silberstein. 2022. FlexDriver: A Network Driver for Your Accelerator. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 1115–1129.

[24] Suhaib A. Fahmy, Kizheppatt Vipin, and Shanker Shreejith. 2015. Virtualized FPGA Accelerators for Efficient Cloud Computing. In *IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. 430–435.

[25] Suhaib A. Fahmy, Ziyi Yang, Yixi Chen, Gustavo Alonso, Zsolt István, and Marco Canini. 2025. FPGAs Are the Hero In-Network Computing Needs. In *ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)*. 131–139.

[26] FastML Team. 2025. fastmachinelearning/hls4ml. https://github.com/fastmachinelearning/hls4ml Version 1.1.0.

[27] Alex Forencich, Alex C. Snoeren, George Porter, and George Papen. 2020. Corundum: An Open-Source 100-Gbps NIC. In *IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. 38–46.

[28] Song Han, Junlong Kang, Huizi Mao, Yiming Hu, Xin Li, Yubin Li, Dongliang Xie, Hong Luo, Song Yao, Yu Wang, Huazhong Yang, and William (Bill) J. Dally. 2017. ESE: Efficient Speech Recognition Engine with Sparse LSTM on FPGA. In *ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*. 75–84.

[29] Zhenhao He, Dario Korolija, and Gustavo Alonso. 2021. EasyNet: 100 Gbps Network for HLS. In *International Conference on Field-Programmable Logic and Applications (FPL)*. 197–203.

[30] John L. Hennessy and David A. Patterson. 2019. A New Golden Age for Computer Architecture. *Commun. ACM* 62, 2 (2019), 48–60.

[31] Wentao Hou, Jie Zhang, Zeke Wang, and Ming Liu. 2024. Understanding Routable PCIe Performance for Composable Infrastructures. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 297–312.

[32] Zhongyun Hua, Binghang Zhou, and Yicong Zhou. 2018. Sine-Transform-Based Chaotic System With FPGA Implementation. *IEEE Transactions on Industrial Electronics* 65, 3 (2018), 2557–2566.

[33] Stephen Ibanez, Gordon Brebner, Nick McKeown, and Noa Zilberman. 2019. The P4->NetFPGA Workflow for Line-Rate Packet Processing. In *ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*. 1–9.

[34] Stephen Ibanez, Alex Mallery, Serhat Arslan, Theo Jepsen, Muhammad Shahbaz, Changhoon Kim, and Nick McKeown. 2021. The nanoPU: A Nanosecond Network Stack for Datacenters. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 239–256.

[35] Zsolt István, Louis Woods, and Gustavo Alonso. 2014. Histograms as a Side Effect of Data Movement for Big Data. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*. 1567–1578.

[36] Zhipeng Jia and Emmett Witchel. 2021. Nightcore: Efficient and Scalable Serverless Computing for Latency-Sensitive, Interactive Microservices. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 152–166.

[37] Petar Jokic, Stephane Emery, and Luca Benini. 2018. BinaryEye: A 20 kfps Streaming Camera System on FPGA with Real-Time On-Device Image Recognition Using Binary Neural Networks. In *IEEE International Symposium on Industrial Embedded Systems (SIES)*. 1–7.

[38] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. 2017. In-Datacenter Performance Analysis of a Tensor Processing Unit. *ACM SIGARCH Computer Architecture News* 45, 2 (2017), 1–12.

[39] Anuj Kalia, Michael Kaminsky, and David Andersen. 2019. Datacenter RPCs Can Be General and Fast. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 1–16.

[40] Sagar Karandikar, Chris Leary, Chris Kennelly, Jerry Zhao, Dinesh Parimi, Borivoje Nikolic, Krste Asanovic, and Parthasarathy Ranganathan. 2021. A Hardware Accelerator for Protocol Buffers. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 462–478.

[41] Dario Korolija, Timothy Roscoe, and Gustavo Alonso. 2020. Do OS Abstractions Make Sense on FPGAs?. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 991–1010.

[42] Amit Kulkarni, Monica Chiosa, Thomas B. Preußer, Kaan Kara, David Sidler, and Gustavo Alonso. 2020. HyperLogLog Sketch Acceleration on FPGA. In *International Conference on Field-Programmable Logic and Applications (FPL)*. 47–56.

[43] Christian Leber, Benjamin Geib, and Heiner Litz. 2011. High Frequency Trading Acceleration Using FPGAs. In *International Conference on Field-Programmable Logic and Applications (FPL)*. 317–322.

[44] Bojie Li, Kun Tan, Layong (Larry) Luo, Yanqing Peng, Renqian Luo, Ningyi Xu, Yongqiang Xiong, Peng Cheng, and Enhong Chen. 2016. ClickNP: Highly Flexible and High Performance Network Processing with Reconfigurable Hardware. In *ACM SIGCOMM Conference (SIGCOMM)*. 1–14.

[45] Luyang Li, Heng Pan, Xinchen Wan, Kai Lv, Zilong Wang, Qian Zhao, Feng Ning, Qingsong Ning, Shideng Zhang, Zhenyu Li, Layong Luo, and Gaogang Xie. 2025. Harmonia: A Unified Framework for Heterogeneous FPGA Acceleration in the Cloud. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 498–514.

[46] Katie Lim, Matthew Giordano, Theano Stavrinos, Irene Zhang, Jacob Nelson, Baris Kasikci, and Thomas Anderson. 2024. Beehive: A Flexible Network Stack for Direct-Attached Accelerators. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 393–408.

[47] Jiaxin Lin, Adney Cardoza, Tarannum Khan, Yeonju Ro, Brent E. Stephens, Hassan Wassel, and Aditya Akella. 2023. RingLeader: Efficiently Offloading Intra-Server Orchestration to NICs. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 1293–1308.

[48] Jiaxin Lin, Kiran Patel, Brent E. Stephens, Anirudh Sivaraman, and Aditya Akella. 2020. PANIC: A High-Performance Programmable NIC for Multi-Tenant Networks. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 243–259.

[49] Will Lin, Yizhou Shan, Ryan Kosta, Arvind Krishnamurthy, and Yiying Zhang. 2024. SuperNIC: An FPGA-Based, Cloud-Oriented SmartNIC. In *ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*. 130–141.

[50] Weili Liu. 2020. Voice Control System Based on Zynq FPGA. *Journal of Physics: Conference Series* 1631, 1 (2020), 012177.

[51] Jun Luo, Qijun Huang, Sheng Chang, Xiaoying Song, and Yun Shang. 2013. High Throughput Cholesky Decomposition Based on FPGA. In *International Congress on Image and Signal Processing (CISP)*. 1649–1653.

[52] Weiwei Ma, M. E. Kaye, D. M. Luke, and R. Doraiswami. 2006. An FPGA-Based Singular Value Decomposition Processor. In *Canadian Conference on Electrical and Computer Engineering (CCECE)*. 1047–1050.

[53] Zhihuang Ma, Xingming Cui, Xiaoliang Chen, and Zuqing Zhu. 2026. Rain: RDMA-Assisted In-Network Scheduling for Microsecond-Scale Workloads. *Proceedings of the ACM on Networking* 4, CoNEXT2 (2026), 22:1–22:21.

[54] Charalampos Mainas, Martin Lambeck, Bruno Scheufler, Laurent Bindschaedler, Atsushi Koshiba, and Pramod Bhatotia. 2025. F3: An FPGA-Accelerated FaaS Framework. In *ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*. 1–16.

[55] Fabio Maschi and Gustavo Alonso. 2024. Strega: An HTTP Server for FPGAs. *ACM Transactions on Reconfigurable Technology and Systems* 17, 1 (2024), 3:1–3:27.

[56] Naoyuki Matsumoto, Koji Nakano, and Yasuaki Ito. 2015. Optimal Parallel Hardware K-Sorter and Top K-Sorter, with FPGA Implementations. In *International Symposium on Parallel and Distributed Computing (ISPDC)*. 138–147.

[57] Jeffrey C. Mogul. 2003. TCP Offload Is a Dumb Idea Whose Time Has Come. In *USENIX Workshop on Hot Topics in Operating Systems (HotOS)*. 25–30.

[58] Mokhles A. Mohsin and Darshika G. Perera. 2018. An FPGA-Based Hardware Accelerator for K-Nearest Neighbor Classification for Machine Learning on Mobile Devices. In *International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies (HEART)*. 6:1–6:7.

[59] Sergio D. Muñoz and Javier Hormigo. 2015. High-Throughput FPGA Implementation of QR Decomposition. *IEEE Transactions on Circuits and Systems II: Express Briefs* 62, 9 (2015), 861–865.

[60] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Mallevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep Learning Recommendation Model for Personalization and Recommendation Systems. arXiv:1906.00091

[61] Stanko Novakovic, Alexandros Daglis, Edouard Bugnion, Babak Falsafi, and Boris Grot. 2014. Scale-Out NUMA. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 3–18.

[62] NVIDIA Corporation. 2025. NVIDIA BlueField-3 DPU. https://www.nvidia.com/en-us/networking/products/data-processing-unit/ Accessed: 2025-10-20.

[63] NVIDIA Corporation. 2026. *NVIDIA CUDA C++ Programming Guide*. https://docs.nvidia.com/cuda/cuda-c-programming-guide/

[64] OpenFaaS Ltd. 2025. OpenFaaS. https://www.openfaas.com/ Accessed: 2025-11-27.

[65] Yuliang Pu, Jun Peng, Letian Huang, and John Chen. 2015. An Efficient KNN Algorithm Implemented on FPGA Based Heterogeneous Computing System Using OpenCL. In *IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. 167–170.

[66] Andrew Putnam, Adrian M. Caulfield, Eric S. Chung, Derek Chiou, Kypros Constantinides, John Demme, Hadi Esmaeilzadeh, Jeremy Fowers, Gopi Prashanth Gopal, Jan Gray, Michael Haselman, Scott Hauck, Stephen Heil, Amir Hormati, Joo-Young Kim, Sitaram Lanka, James Larus, Eric Peterson, Simon Pope, Aaron Smith, Jason Thong, Phillip Yi Xiao, and Doug Burger. 2015. A Reconfigurable Fabric for Accelerating Large-Scale Datacenter Services. *IEEE Micro* 35, 3 (2015), 10–22.

[67] Benjamin Ramhorst, Dario Korolija, Maximilian Jakob Heer, Jonas Dann, Luhao Liu, and Gustavo Alonso. 2025. Coyote v2: Raising the Level of Abstraction for Data Center FPGAs. In *ACM Symposium on Operating Systems Principles (SOSP)*. 639–654.

[68] Esteban Ramos, Robert Soulé, Peter Alvaro, Pietro Bressana, Edmund Chen, Uri Cummings, Rui Li, James Tsai, and Rajit Manohar. 2024. Split gRPC: An Isolation Architecture for RPC Software Stacks. In *ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)*. 81–87.

[69] Waleed Reda, Marco Canini, Dejan Kostić, and Simon Peter. 2022. RDMA Is Turing Complete, We Just Did Not Know It Yet!. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 71–85.

[70] Hugo Sadok, Nirav Atre, Zhipeng Zhao, Daniel S. Berger, James C. Hoe, Aurojit Panda, Justine Sherry, and Ren Wang. 2023. Ensō: A Streaming Interface for NIC-Application Communication. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 1005–1025.

[71] S. K. Shome, S. R. K. Vadali, U. Datta, S. Sen, and A. Mukherjee. 2012. Performance Evaluation of Different Averaging Based Filter Designs Using Digital Signal Processor and Its Synthesis on FPGA. *International Journal of Signal Processing, Image Processing and Pattern Recognition* 5, 3 (2012), 75–92.

[72] Akshay Srivatsan, Yuhan Deng, Katherine Mohr, Emma Sudo, Sebastian Ingino, Francis Chua, and Keith Winstein. 2026. Continuation-Centric Computing with Arca. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 1989–2005.

[73] Mark Sutherland, Siddharth Gupta, Babak Falsafi, Virendra Marathe, Dionisios Pnevmatikatos, and Alexandros Daglis. 2020. The NEBULA RPC-Optimized Architecture. In *ACM/IEEE International Symposium on Computer Architecture (ISCA)*. 199–212.

[74] Ciprian-Octavian Truică and Jérôme Darmont. 2017. T^2K^2: The Twitter Top-K Keywords Benchmark. In *European Conference on Advances in Databases and Information Systems (ADBIS), Short Papers and Workshops*. 21–28.

[75] Kizheppatt Vipin and Suhaib A. Fahmy. 2018. FPGA Dynamic and Partial Reconfiguration: A Survey of Architectures, Methods, and Applications. *Comput. Surveys* 51, 4 (2018), 72:1–72:39.

[76] Zishen Wan, Yuyang Zhang, Arijit Raychowdhury, Bo Yu, Yanjun Zhang, and Shaoshan Liu. 2021. An Energy-Efficient Quad-Camera Visual System for Autonomous Machines on FPGA Platform. In *IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. 1–4.

[77] Xilinx Inc. 2019. Singular Value Decomposition (SVD). https://xilinx.github.io/Vitis_Libraries/solver/2019.2/guide_L2/L2_api.html Vitis Solver Library 2019.2. Accessed: 2024-12-02.

[78] Ziyi Yang, Zsolt István, Marco Canini, and Suhaib A. Fahmy. 2026. Application-Level Function Pipelines on In-Network FPGA Accelerators. In *IFIP Networking Conference (IFIP Networking)*. 73–82.

[79] Yifan Yuan, Jinghan Huang, Yan Sun, Tianchen Wang, Jacob Nelson, Dan R. K. Ports, Yipeng Wang, Ren Wang, Charlie Tai, and Nam Sung Kim. 2023. Rambda: RDMA-Driven Acceleration Framework for Memory-Intensive μs-scale Datacenter Applications. In *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 499–515.

[80] Jie Zhang, Hongjing Huang, Xuzheng Chen, Xiang Li, Jieru Zhao, Ming Liu, and Zeke Wang. 2025. RpcNIC: Enabling Efficient Datacenter RPC Offloading on PCIe-Attached SmartNICs. In *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 1379–1394.

[81] Yanqi Zhang, Íñigo Goiri, Gohar Irfan Chaudhry, Rodrigo Fonseca, Sameh Elnikety, Christina Delimitrou, and Ricardo Bianchini. 2021. Faster and Cheaper Serverless Computing on Harvested Resources. In *ACM Symposium on Operating Systems Principles (SOSP)*. 724–739.

[82] Zhipeng Zhao, Hugo Sadok, Nirav Atre, James C. Hoe, Vyas Sekar, and Justine Sherry. 2020. Achieving 100Gbps Intrusion Prevention on a Single Server. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 1083–1100.