Beehive: A Flexible Network Stack for Direct-Attached Accelerators

Katie Lim
University of Washington

Matthew Giordano
University of Washington

Theano Stavrinos
University of Washington

Baris Kasikci
University of Washington

Tom Anderson
University of Washington

arXiv:2403.14770v1 [cs.AR] 21 Mar 2024

Abstract

Accelerators have become increasingly popular in datacenters due to their cost, performance, and energy benefits. Direct-attached accelerators, where the network stack is implemented in hardware and network traffic bypasses the main CPU, can further enhance these benefits. However, modern datacenter software network stacks are complex, with interleaved protocol layers, network management functions, as well as virtualization support. They also need to flexibly interpose new layers to support new use cases. By contrast, most hardware network stacks only support basic protocol compatibility and are often difficult to extend due to using fixed processing pipelines.

This paper proposes Beehive, a new, open-source hardware network stack for direct-attached FPGA accelerators designed to enable flexible and adaptive construction of complex protocol functionality. Our approach is based on a network-on-chip (NoC) substrate, automated tooling for the independent scale-up of protocol elements, compile-time deadlock analysis, and a flexible diagnostics and control plane. Our implementation interoperates with standard Linux TCP and UDP clients, allowing existing RPC clients to interface with the accelerator. We use three applications to illustrate the advantages of our approach: a throughput-oriented erasure coding application, an accelerator for distributed consensus operations that reduces the latency and energy cost of linearizability, and TCP live migration support for dynamic server consolidation.

1 Introduction

Hardware accelerators are becoming increasingly common in datacenters to reduce cost, improve performance, and reduce energy consumption relative to server CPUs. Typically, accelerators are hosted over the PCIe I/O bus, with the server CPU mediating all communication with the accelerator, illustrated in Box **(c)** in Figure 1. An emerging alternative model directly attaches the accelerator to the network, with its own network functionality implemented in hardware, illustrated in Box **(b)** in Figure 1. By bypassing the CPU, such a specialized hardware network stack reduces latency, performance variability, and overhead, freeing up the CPU for other purposes. Another advantage of this approach is that it simplifies scaling out accelerators to meet workload needs, independently of server hardware.

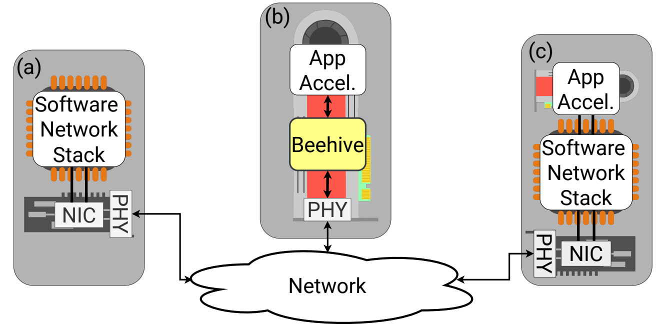

Figure 1. Box **(a)** represents the standard CPU server node. Box **(b)** represents Beehive's targeted, direct-attached setup where Beehive provides the network stack. Box **(c)** represents the traditional method of attaching an accelerator to the network using a CPU network stack. Beehive enables direct-attached accelerators to interoperate with other types of nodes on the network.

A key barrier to fully realizing the benefits of direct-attached hardware acceleration is the growing disparity between datacenter networking requirements and hardware network stack implementations. Modern datacenter networks have intricate host network stacks that are constantly evolving to meet customer and operator use cases [45]. Beyond core protocols, such as TCP/IP, modern applications require higher-level host network functionality like remote procedure call (RPC) processing, quality-of-service (QoS) management [17, 62], encryption [1, 19], and load balancing [18, 33]. Deployment flexibility necessitates management features like virtual networking [23, 26, 38], access control lists [46], congestion control [39, 42], traffic prioritization [29, 47], and load balancing [24, 49, 55, 58]. Deployment maintainability requires dynamic support for network monitoring [10, 63], reconfiguration [12, 37], and debugging [57].

An example of a highly-flexible software network stack is Google's Snap networking system [45]. It is designed around composable message-passing engines, with modules for load balancing, network virtualization, network management, and custom transport protocols. New modules can be easily inserted anywhere in the stack, without re-engineering the

rest of the stack. By contrast, existing hardware network stacks are specialized and rigid since they are typically designed to support a single application with few protocol layers. Although some work focuses on flexible packet-level processing in hardware [41, 43], they operate at the packet handling layer with no support for higher-level transport and application protocols. Other work focuses on hardware transport offloads, but these systems lack a range of essential network functions [7, 16, 53].

This paper explores the design of an FPGA network stack for direct-attached accelerators that can better support the extensibility and complexity needed in a production environment. Flexibility is needed at multiple layers of the network stack: in packet processing (layer 3), transport and congestion control (layer 4), and the application layer (layer 7). Adding new functionality, scaling out a protocol layer to meet workload demand, or changing internal load balancing should be simple, as it is in software, without disrupting or re-engineering other layers. Our goal is to make a hardware network stack that is flexible and scalable while preserving the cost, performance, and energy benefits of direct-attached accelerators.

We propose and implement Beehive, an open-source hardware network stack architected as a collection of protocol layers and network functions that communicate via message-passing over a scalable network-on-chip (NoC). We provide automated tooling for managing scale out, load balancing, and compile-time deadlock analysis. To make our design concrete, we implement Ethernet, IP, UDP, TCP, network address translation (NAT), IP-in-IP encapsulation, and additional support for control and debugging of network functions. Our implementation interoperates with Linux TCP and UDP clients, allowing unmodified remote procedure call (RPC) clients to use our accelerator.

For our evaluation, we implement Beehive on an FPGA and show that it offers up to 31x higher per-core throughput than state-of-the-art CPU kernel-bypass networking stacks on small messages. We also demonstrate how Beehive can improve performance and energy consumption in three important use cases compared to CPU-only implementations.

First, modern data center storage systems typically use erasure coding to achieve better storage efficiency than replication with comparable fault tolerance. We implement an erasure coding accelerator in Beehive and show that, compared to a CPU-only version, the accelerator scales out to 62 Gbps using 20× less energy.

Second, many production distributed systems use consensus algorithms for data consistency even in case of failures. However, deployed systems often skip consensus for some reads in exchange for better performance. We show that accelerating a key piece of the consensus algorithm in hardware can reduce end-to-end median operation latency by 1.13×, with 1.14× better per-core throughput and 2× less energy than the CPU-only version. Because distributed consensus is dominated by small messages, it requires independent scale out and load balancing of different protocol layers, implemented without changing the protocol logic.

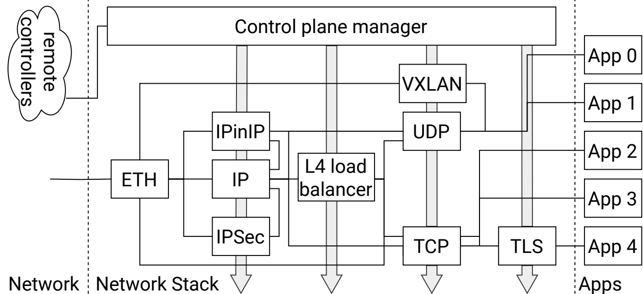

Figure 2. A high-level diagram of the type of network stack Beehive targets. Along with multiple transport protocols, this stack has IP-in-IP and VXLAN for network virtualization and a component for an L4 load balancer. The control plane potentially needs access to all components.

Finally, we demonstrate the advantages of Beehive's message-passing model to add support for TCP live migration by adding network address translation and management capabilities without changes to any other protocol layers. This design allows accelerator clients to be transparently migrated without resetting their TCP connection, with less than a millisecond of added end-to-end operation delay.

In summary, we contribute:

- Beehive, a design framework to build efficient and complex hardware network stacks for direct-attached accelerator deployments in modern datacenters.
- An open-source FPGA implementation of Beehive that includes tools and reusable components to build network stacks for accelerators that use different transport protocols, network virtualization, and L7 functionality.
- A demonstration of Beehive's ability to support scalability, flexibility, low latency, high throughput, and energy efficiency by integrating and evaluating an erasure coding accelerator, a consensus accelerator, and TCP migration functionality.

2 Beehive Design Goals

Our overarching goal for Beehive is to build an open-source FPGA hardware design to support emerging applications for direct network-attached accelerators in a production environment. Figure 2 shows a high-level diagram of the type of network stack architecture we want to be able to support. Applications may only use some subset of these protocols and network functions. This requires:

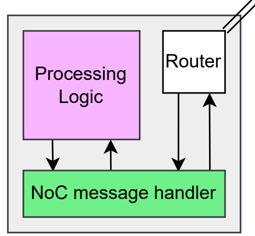

Figure 3. Architecture of a tile.

Standard client protocols. The vast majority of distributed applications that might benefit from the availability of hardware acceleration are designed to communicate using standard protocols such as IP, TCP, and remote procedure call (RPC). Our framework needs to be able to support unmodified client application and client host software communicating with the accelerator using these standard protocols.
Modularity. However, network stacks are not fixed. Requirements are constantly changing with new custom protocols (e.g. Google's Pony Express [45] or 1RMA [6]) and network functions. In order to facilitate rapid development and customization of the network stack, our framework must be modular, so we can compose or integrate new components with minimal to no modifications to existing components.
Scalability. Building a complex network stack potentially means supporting a variety of different components in the same design. Different components may be a bottleneck depending on the application workload. Thus, the architecture should be able to duplicate and scale out individual components, whether application or protocol logic, as needed.
Performance overhead and predictability. Since performance and performance predictability are key motivations to offload the network stack, the stack should be able to deliver end-to-end application bandwidth at 100 Gbps with minimal jitter if the accelerators have the capacity to support it.
Management flexibility. Components in a network stack need to be able to interact beyond just passing packet data. For example, components need to be able to expose interfaces to the control plane for telemetry and debugging [25]. The control plane may also need to update state used by a protocol or network function, such as configuring the load balancer used to parcel work across application accelerator instances. Such configurability should be possible even in large designs without extensive manual optimization.

3 Design

3.1 Beehive's Architecture

The basic component in Beehive is the tile, shown in Figure 3. Each tile has a network-on-chip (NoC) router, some logic that handles NoC message construction and deconstruction, and some processing logic, such as a protocol layer, network function, or application. Tile routers are connected together

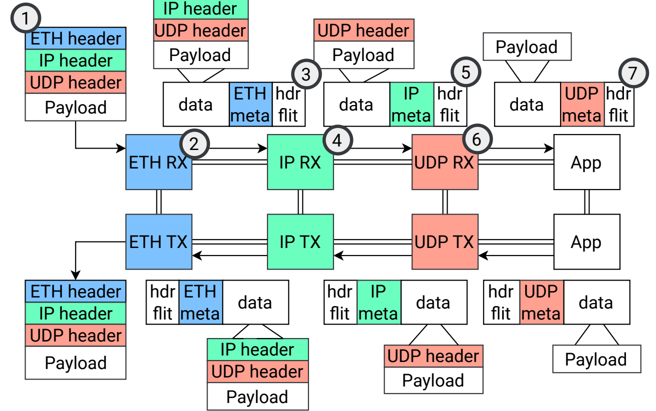

Figure 4. The flow through which a packet is processed or constructed in Beehive.

to form the NoC topology. We do not require a particular topology, although our prototype uses a 2D mesh. We assume the NoC is reliable, point-to-point ordered, and uses deterministic, deadlock-free routing.

A network packet is processed or constructed by passing NoC messages through a chain of tiles. A NoC message consists of one header flit followed by some number of body flits. The header flit typically contains data only relevant to NoC-level routing, such as source and destination tile coordinates or number of body flits. The body flits typically consist of both metadata flits containing packet header fields and a number of data flits carrying unprocessed packet payload.

Each tile hop is responsible for determining the next tile that a message should be sent to. This design is in contrast to earlier work which assumes that routes can be fully determined on packet arrival [43]. We discuss this decision in more detail in Section 3.4. This component may vary in complexity from a static CAM to more complex logic, such as content-based routing. The set of possible message chains is known ahead of time for deadlock analysis ,described in Section 3.5.

3.2 Processing a Packet

Figure 4 shows an example of a basic UDP stack in Beehive, with a UDP packet moving through the receive and send paths.

On the receive side, an Ethernet frame enters the Ethernet tile, which has ports for the I/O from the transceivers in addition to the ports connecting to other tiles. The processing logic within the tile parses and removes the Ethernet header, realigning the data. This is then turned into a NoC message consisting of a header flit, a metadata flit with the parsed Ethernet header, and some number of data flits containing the remaining packet data. The routing component in the Ethernet tile uses the type field in the Ethernet header to determine that the message should be passed to the IP tile. The

IP tile similarly parses the IP header, validates the header's checksum, and then creates a NoC message to be sent to the UDP layer based on the protocol field. Finally, the UDP tile parses the UDP header, validates the packet's checksum, and generates a NoC message to be sent to the application based on the port in the UDP header. The transmit path runs similarly, except instead of parsing headers from the data flits, headers are added into the data flits by each protocol tile. After the Ethernet tile adds on the Ethernet header, it is sent out the ports for I/O with the transceivers. This incremental composability is good for our goal of modularity as it makes it easier to insert new functionality between stages.

While there is only one possible destination for the tiles in this design, there can potentially be multiple endpoints, such as other protocols (e.g. TCP connected to IP), network services (e.g. network virtualization), or replicated tiles for higher bandwidth. With replicated tiles, there are multiple ways to decide on which tile should receive an incoming packet. The simplest method is to distribute packets between them in a round-robin fashion. However, more complex scheduling may be necessary if a tile holds state for particular flow. In this case, it is important that packets from the same flow always go to the same tile. This distribution can either be integrated within a tile or placed in a dedicated tile. We discuss examples about how we distribute packets to duplicated tiles in Section 5.

3.3 Message-Passing Interconnect

Being able to compose elements is essential for facilitating customization. We opt for a message passing model. This is beneficial for modularity, because defining a message-passing format allows us to standardize the physical interconnection between components, a recognized benefit in SoC design [22], and makes it easier to chain offloads together. ClickNP [41] and PANIC [43], two modular packet processing frameworks, have also used a message-based approach. The message passing can be done over dedicated connections, which is the approach used by ClickNP, or a NoC which is used by PANIC.

We prefer a NoC interconnect for two main reasons related to our goal of scalability. First, we can take advantage of the multiplexing that is provided by the NoC routers. Certain tiles may interact with a large number of other tiles, e.g. if we instantiate multiple copies of the same component. Direct connections can lead to large multiplexers and wires with significant fan-out. Although we could create specialized pipelined multiplexers and arbiters, these essentially look like NoC routers.

Second, we would like the interconnect wiring to remain stable whenever possible. In the ClickNP model, top-level wires are determined by the computational graph. If we wish to form a chain that links together two components that did not communicate before, we must add new interconnect wires, which are typically the longest wires. A NoC allows us to reuse physical wiring to chain any elements that exist in the design, although we must be careful with deadlock.

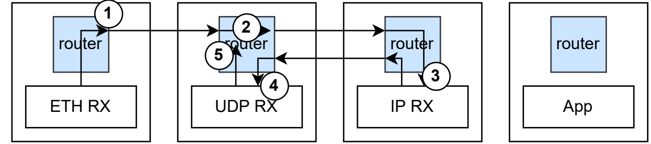

(a) This tile assignment deadlocks due to the order that the packet needs to be processed in versus the order of the NoC links it traverses.

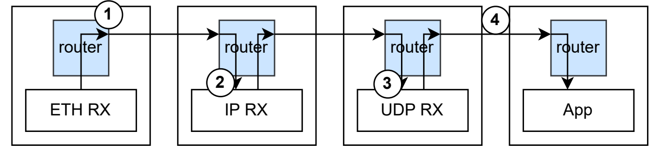

(b) This tile assignment is able to avoid deadlock since the packet acquires NoC resources in order

Figure 5. A demonstration of how tile assignment affects deadlock. We can take advantage of protocol layer ordering, so a packet being processed always acquires NoC resources in the same order.

3.4 Tile Chain Routing

In addition to NoC-level routing, Beehive has routing at the network packet level to determine the sequence of tiles that need to be chained together. We considered two routing methods: node-table routing, where each tile determines the correct next tile, and source routing, where the chain of tiles is completely determined when the first NoC message in the chain is created, such as when a packet is first received from the network. We chose to use node-table routing, because certain classes of traffic we want to support for our goal of interoperability require per-flow state or non-trivial protocol processing to fully determine the chain of tiles. Specifically, we consider routing for traffic that is either encrypted or is for layer 7.

Encryption may obfuscate parts of packet payloads that are needed to fully route a packet, which would require the ingress tile to handle the decryption. An application request can span multiple packets. Which application tile should receive an RPC may depend on the RPC header or even the contents of the request. Further, the packets of one request can potentially be reordered or interleaved with other requests. To properly route these requests, an ingress tile would need to reassemble the stream, further complicating the implementation. In both cases, the ingress tile would need to implement significant, high-level protocol logic which is detrimental for modularity.

3.5 Deadlock

As with any NoC-based design, avoiding message-based deadlock must be a consideration. We note that NoC deadlock detection, avoidance, and recovery is a complex problem with a whole body of research behind it.

NoCs can deadlock in two ways: at the routing level and at the message passing level. To prevent routing-level deadlocks, we employ dimension-ordered routing [21]. Message passing deadlocks are a bigger concern in Beehive, because we enable each tile to route to any other tile at runtime. This means that our routing resources can get exhausted. The deadlock in Figure 5a is an example of this, in which the UDP RX tile must route east twice in one chain and it cannot route east a second time.

We apply resource acquisition ordering to solve this problem. Resource ordering can be imposed by taking advantage of the fact that protocol layers and services are composed in certain orders. For any individual packet, the path is determined at runtime, but we assume that all possible paths through the network stack for supported packet types are known when the network stack is compiled. As a simple example, consider the examples in Figure 5 of different topologies for the receive path of a UDP stack. Beehive's NoC uses wormhole, dimension-ordered routing. The packet should be processed by Ethernet, IP, UDP, and then the application. With the tile layout in Figure 5a, the route from the Ethernet to IP tile passes through the UDP tile's router. As the UDP tile attempts to pass the packet along to the application, it must reacquire a NoC link it already used (5) and is thus deadlocked. If tiles are laid out as in Figure 5b, no resources need to be reacquired, and the packet can be processed successfully.

We statically analyze all message paths in our prototypes at compile-time to avoid deadlock by creating a resource dependency graph that takes into account every possible path through the network stack. If a message path is found that could cause deadlock, the designer should modify the tile layout to one that does not.

Repeated protocol headers (e.g. two IP headers in the IP-in-IP protocol) break resource ordering. In Beehive, we choose to duplicate tiles (e.g. two IP RX tiles). If tiles are too expensive to duplicate, a potential solution is adding buffers to break dependencies [40, 56]. These buffers give space for the NoC to drain into, freeing routing resources.

3.6 Control Plane Interfaces

For manageability, network operators need to be able to reconfigure protocol components from an external controller over a transport-layer connection. In Beehive, we choose to use a NoC as well for the control plane rather than a dedicated control bus. This is because control plane management can also benefit from a structured interconnect for scalability reasons.

First, for complex designs with a large number of components, it becomes costly to run dedicated, ad-hoc wires to every tile. Second, we want configuration to be over a reliable transport. This requires the control plane to use the transport layer, and a NoC enables this without physically coupling the component to the transport layer. This also enables us to add specific control plane management tiles to orchestrate state modifications. We describe a specific example in Section 4.5.

Because the control plane has lower performance requirements, in Beehive we use a separate, lower-width NoC. This also prevents control plane traffic from contending for the same resources as long dataplane chains in the deadlock dependency graph, so there is more flexibility in placement.

3.7 Application Interfaces

Many application accelerators process requests at a coarser granularity than a packet, so they need the ability to communicate with the transport protocol layer and request data from a particular flow rather than being pushed packets in the order they arrive. While we could use dedicated wires for this communication, it can also benefit from the use of the NoC.

To support duplicated application tiles connected to the same transport layer, the NoC provides a convenient structure to multiplex between them in a scalable manner. The modularity provided by message passing on the NoC also allows an application to easily interface with any protocol in the network stack while reusing existing wires if, for example, an application wanted to switch from TCP to a custom reliable transport protocol. Finally, the standard interface of the NoC enables easy insertion of filters on the application's NoC message, so network operators can enforce policies, such as dropping network traffic to or from non-whitelisted nodes. We describe the application NoC interface to our TCP layer in Section 4.4.

4 Implementation

To demonstrate the Beehive approach, we built a set of core protocol tiles, network functions, and applications. For protocols, we implement tiles for Ethernet, IPv4, UDP, and TCP. For network functions, we implement an IPinIP encapsulation layer and a NAT layer for network virtualization. For applications, we implement a Reed-Solomon encoder and an accelerator for a viewstamped replication node. These applications are described in more detail in Section 5.

We also describe our tooling that we developed to lower the effort required to maintain multiple designs and integrate new components. All of Beehive is implemented in standard SystemVerilog and was tested on an Alveo U200 communicating with standard CPU clients using a Linux or kernel-bypass network stack. We embed our Beehive prototype within Corundum [28], an open-source 100 Gbps NIC, in

the application slot to provide FPGA-specific infrastructure, such as the Ethernet MAC.

4.1 Network-on-chip (NoC)

We use the 2D mesh NoC from OpenPiton [9] with some modifications. The NoC is wormhole-routed, uses dimension-ordered routing, and is full-duplex. We widen the NoC to 512 bits to match the width of the Xilinx MAC IP core and increase the flit width to 512 bits. The header flit format is inherited from OpenPiton. The maximum payload size for a NoC message is 256 MiB.

4.2 Protocol tiles

Protocols have one tile each for transmit and for receive processing. Each tile can be replicated if more throughput is needed for that element. Protocols are implemented as streaming components, so they begin to transmit the next NoC message as soon as possible rather than storing the entire NoC message before forwarding. Since each router has one input interface and one output interface for the tile, each side will utilize an entire router's bandwidth if running at 100 Gbps. Since the packet-level protocol layers do not share state between their transmit and receive sides, this is a straightforward split.

The Ethernet, IP, and UDP tiles construct or remove the appropriate headers and calculate checksums, as shown in Figure 4. The Ethernet receive processor can handle VLAN tagged packets. Our IP layer does not support IP fragmentation as our intended use case is for internal datacenter services.

Protocol tiles also have optional hash tables that use the 4 tuple as the key for load balancing to downstream replicated tiles. We set up initial packet-level routing within the tiles at compile time when we build the FPGA image. The hash table can be rewritten during runtime. Any packet that does not have an entry for a next hop (e.g.traffic with an unsupported protocol) is dropped to filter out unwanted traffic.

4.3 Buffer tiles

In Beehive, we also have buffer tiles that hold large blocks of memory. In our current prototype, these buffers are large BRAMs, but the backing buffer could also be DRAM. These buffer tiles are accessible to any other tile in the system via NoC messages. This allows us to have shared buffers between tiles, so that multiple tiles can share state when needed.

4.4 TCP engine

In order to evaluate how Beehive can support reliable transport, we prototype a TCP engine that implements server-side TCP. It can receive connection setup requests, generate sequence and ACK numbers, and support fast retransmit and window-based flow control [11]. Currently, it does not support selective acknowledgments, initiating connections, or congestion control. Full TCP offload functionality has been demonstrated by previous work [53] and could be integrated into Beehive.

While the TCP engine has an RX router and a TX router like the other protocol tiles, the send and receive paths in TCP must share state. For example, the transmit path needs to know for which packets it has received acknowledgments. We choose to support sharing by running dedicated wires between the tiles. Every receive path only has one corresponding transmit path, so wires do not fan out. Although dedicated wires cause tiles to be closely coupled, the frequent NoC messages needed for state updates would lead to a practical coupling.

On the completion of the 3-way handshake, the TCP engine sends a NoC message to notify an application tile based on the destination port for the connection. On the receive side, the TCP engine implements an interface that lets an application specify the size of the request it should be notified for with a NoC message. When enough data has arrived to satisfy that request, the TCP engine sends a notification message back to the application with the buffer address where the data requested has been stored. The application then retrieves the data from the buffer for processing before sending another message to the TCP stack when it has finished using the data.

The transmit side is similar. The TCP engine implements an interface where the application can request space in its transmit buffer of a certain size. The TCP engine sends a notification when there is room in that buffer with the buffer address where the data should be stored. The application then copies the data into the buffer and notifies the TCP engine.

4.5 Network function tiles

We implement both IPinIP encapsulation and an IP NAT; the client-side TCP migration we use for our evaluation assumes a NAT instead of encapsulation. For both tiles, the control plane can dynamically update the table that maps virtual IPs to physical IPs. This mapping is changed when the client migrates.

Taking advantage of Beehive's control plane interconnect, we implement an internal controller as a separate tile that receives an RPC over TCP from an external controller. The internal controller then sends NoC messages to the IP encapsulation or NAT tiles with the information needed to update their tables. Finally, the internal controller sends a confirmation response to the external controller over the TCP connection.

4.6 Debugging and logging

In Beehive, tiles may keep logs, and we provide a UDP-based protocol to readback logs. Each log is associated with a particular UDP port and exposes its own interface on the NoC for readback. The UDP receive layer is responsible for directing packets to the appropriate log interfaces. The log read

interface keeps a small buffer for requests and drops requests when it is full. The client program reads out the log an entry at a time and resend requests for any entries for which it does not receive a response.

We also have tiles that log information about TCP packet headers to help provide more visibility into the FPGA's execution. This log can later be replayed in a cycle-accurate simulation. We also found that having a cycle accurate trace is necessary for proper replay, because the TCP engine may behave differently depending on the timing of events (e.g. dropping different packets).

These tiles have two interfaces available over the NoC. One is used to forward packets to and from the TCP engine. This path logs the header information with a cycle timestamp. The other interface allows the logs to be read out over the network in response to a request sent over UDP. Because the logging tiles are embedded within the fabric, they can record the exact timing that packets entered and exited the TCP engine. During simulation replay, Beehive's modularity allows us to easily replace the logging tiles with an interface to our testing replay framework.

4.7 Tooling

We developed a set of tools that lower the engineering effort to create new designs, such as generating portions of the Verilog (e.g. top-level wiring for NoCs) or performing compile-time deadlock analysis. The design configuration is passed to these tools via an XML file. The XML file contains the expected dimensions for the design as well as an element for each NoC tile endpoint. At minimum, this element contains tags specifying a name to use for the endpoint as well as its X and Y coordinates. It may also contain fields with information for generating the tables used for determining the correct next hops.

Given the dimensions in the XML file, we generate declarations of all the top-level wires between tiles. We also generate the subset of the port connections for a tile that correspond to wires between NoC routers and connect the appropriate wires for the tile configuration. We choose not to generate the whole tile instantiation, because certain tiles need to maintain additional ports for I/O, such as the Ethernet MAC.

The XML file also enables us to check whether the high-level topology of the NoC is sound. For example, we check two tiles have the same X and Y coordinates, and all NoC coordinates are within the expected dimensions of the design. Because a 2D mesh must be a rectangle, this also gives us the opportunity to automatically generate empty tiles that just contain a router. We also use the information about the NoC topology and next hops in the XML file to generate a resource dependency graph that we can analyze for cycles.

5 Integrating with Beehive

5.1 Erasure Coding

Erasure codes such as Reed-Solomon (RS) are commonly used in distributed storage systems to achieve high resilience to disk failures with modest storage overhead [31, 35, 52]. An RS encoder turns input data into data plus redundancy bits at a pre-set ratio, striped across storage servers. If some storage elements fail, the remaining blocks from the stripe can be combined with these extra blocks to regenerate the missing blocks.

We configure our system to use an (8,2) code (8 data blocks and 2 redundancy blocks) to emulate a storage system that could tolerate up to two disk failures. We integrate an RS encoding accelerator operating on 4KB requests into Beehive as a UDP application, instantiating four copies of the application to scale out. The accelerator is stateless, so any request can go to any copy. We introduce a front-end round-robin scheduler tile to distribute work among the RS tiles. Each RS tile also logs metadata to calculate bandwidth.

5.2 Consensus Witness

Consensus algorithms are an essential part of many deployed distributed systems as they enable a strictly consistent order for stateful client operations even in the face of failures and message delays/retransmissions. Most consensus algorithms [15, 44, 48] follow a common pattern: an elected leader proposes an order for arriving client requests, verifies with a set of replicas that it is still leader, and commits the request. It then performs any necessary application logic (e.g., to update state), replies back to the client, and informs the other replicas, so that they can also perform the application logic in the same order.

A common type of application built on top of consensus is a key-value (KV) store. To achieve higher throughput, the key space is often sharded with a leader and replica set for each slice. However, even with sharding, consistent reads can be expensive, because the leader must validate, each time, that it is still the leader before replying with the value stored with the key. As a result, it is common in practice to configure the system to return stale reads, allowing the leader to reply immediately [20, 32]. This places a burden on the client developer to handle the (rare) case where a failover can lead to inconsistent client data.

In our evaluation, we show that a consensus accelerator can help reduce the cost of consistency. Our accelerator operates as a witness, that is, it only validates the leader and tracks the operation order; it does not execute client operations. Single node fault tolerance can be achieved with one leader, one witness, and one replica. To add further fault tolerance, we add additional witnesses and replicas. For example, two-node fault tolerance can be achieved with one leader, two witnesses, and two replicas. To validate a read

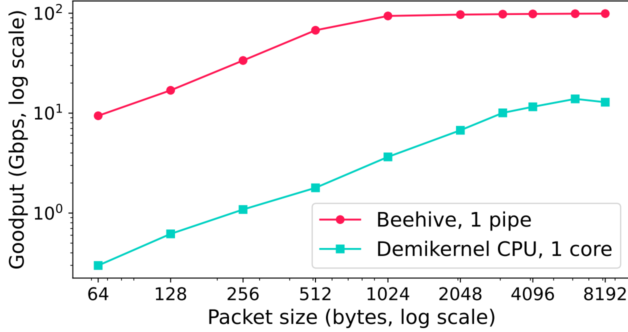

Figure 6. Packet size vs. goodput for a UDP echo application. Beehive with one instance of every protocol tile provides higher bandwidth than Demikernel at all packet sizes.

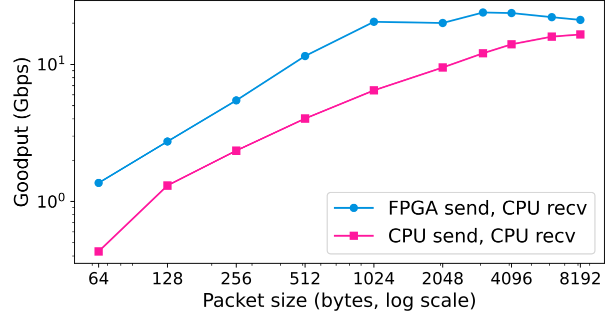

Figure 7. Packet size vs. goodput for Beehive and Linux TCP send. The (CPU send/FPGA receive) is omitted, as it is approximately the same as (CPU send/CPU receive) due to the CPU send path being the bottleneck.

or write operation, the leader only needs to receive a verification from the witnesses before replying to the client. The witness can be designed in hardware to reply with low and reliable latency.

Our witness protocol is based on a modified version of the Viewstamped Replication (VR) used in previous studies of high-performance consensus [51]. VR witnesses are integrated into Beehive as UDP applications. To handle multiple shards, we use one VR witness tile per shard. Unlike the RS encoder, the VR witness is not stateless and requests for a shard must always go to the same tile. We distribute work to the VR tiles by matching on the destination port number.

5.3 TCP live migration

Live TCP migration is useful for load balancing and server consolidation [30]. It allows TCP connections to remain active while virtual machines are migrated in response to offered load, such as during scale down where connections are consolidated onto fewer nodes. To support this, we integrate a network address translation (NAT) service into Beehive, as described in Section 4.5. Specifically, we integrate the NAT tiles between the TCP and IP tiles both on the receive and transmit paths to perform the translation. We also add a new tile to the control plane to handle the management interface. Then, we set up the routing tables of the IP tiles on the receive path to send to the NAT tile, and the NAT tile to send to the TCP tile. We also modify the routing tables of the TCP tiles on the transmit tile to send packets to the NAT tile instead of the IP tile. The internal controller is integrated as an application with the TCP stack.

We can insert this new functionality without any modifications to either the IP or TCP protocol implementations, because the protocols are separated into their own tiles.

Table 1. Lines of code per new tile Beehive for end-to-end applications. XML configuration numbers are given as LoC for declaring the tile plus the LoC to add it as a destination.

	Lines of Code	
	XML Configuration	Verilog Top Level
Reed-Solomon	25 + 6	13
Viewstamped Replication	18 + (6 × # of UDP tiles)	17
TCP Migration	2 × (34 + 6)	2 × 15

6 Evaluation

Our evaluation tests Beehive's ability to support scalability, low latency, and flexibility in a range of network stack configurations. We begin by evaluating Beehive with UDP and TCP microbenchmarks designed to test RPC performance and then evaluate three different application use cases: Reed-Solomon encoding acceleration, and Viewstamped Replication acceleration, and TCP connection live migration.

6.1 Flexibility

As a quantitative proxy for flexibility, we count the lines of code (LoC) required to insert an additional instance of an implemented service (network function or application) into the design for our three designs. Results are shown in Table 1. We count both NAT tiles for the TCP migration service.

6.2 Setup

We use Vivado 2021.2 for building our FPGA images. Beehive is configured on an Alveo U200 at 250 MHz. The FPGA and the clients are connected to an Edgecore Arista DCS-7060CX-32S-R 100G switch with jumbo frames enabled. We use five machines during evaluation with Turboboost disabled. All of them have Mellanox ConnectX-5 100G NICs and are running Ubuntu 20.04. Two machines have Intel Xeon Gold 6226R CPUs; the other three machines have Intel Xeon Gold 5218 CPUs.

Table 2. Energy consumption and goodput for Reed-Solomon encoding using Beehive versus CPU for 1, 2, 3 and 4 application instances.

Apps	1	2	3	4
CPU Energy (mJ/op)	1.1	0.59	0.41	0.32
Beehive Energy (mJ/op)	0.05	0.03	0.02	0.02
Energy efficiency	22x	20x	20x	16x
CPU Goodput (Gbps)	2.0	4.0	6.0	8.0
Beehive Goodput (Gbps)	15	31	45	62
Speedup	7.5x	7.8x	7.5x	7.8x

In experiments where energy is measured, we use the RAPL counters on the CPUs and the Alveo CMS registers on the FPGA. For CPU energy experiments we use a two-socket machine, so we run all the application and network processing code on one socket and poll the counters from the other socket. We only use RAPL's CPU counters, which is an underestimate as we do not include DRAM energy or network interface energy. On the FPGA, we use the Corundum framework to read the CMS registers that report instantaneous power and current usage [59]. We poll these counters every second to calculate energy over the benchmarking period.

6.3 UDP Echo

Throughput: We first compare Beehive's UDP echo goodput to Demikernel, an optimized DPDK network stack, on different packet sizes. Three Intel Gold 5128 machines act as clients; for the CPU experiments, the server runs on the Intel Gold 6226R machine. To stay comparable, we configure one echo application tile on the FPGA and one CPU core in Demikernel. While the CPU server uses DPDK, clients use the standard Linux network stack. We spawn the number of threads that yields highest server bandwidth for that packet size; threads send in an open-loop manner. The results of this experiment are shown in Figure 6.

Beehive achieves line rate with 1024-byte packets while the optimized CPU stack remains far below maximum bandwidth even with jumbo frames. The difference in performance is especially pronounced at small packet sizes where Beehive is able to sustain echoing 9 Gbps of 64-byte packets (18392 KReq/s) whereas single core Demikernel provides 0.3 Gbps (584 KReq/s), a 31× speedup.

Latency: For our latency experiment, we use Beehive and a single client thread to ping-pong a single 1-byte UDP packet. We record the latency by tagging the packet with a timestamp when it enters the network stack at the Ethernet parsing layer, taking another timestamp when it finally exits the Ethernet layer on transmit, and recording both timestamps into a log which we read back over the network. The latency through Beehive is 368 ns (92 cycles).

6.4 TCP throughput

To characterize the throughput performance of our TCP engine, we run a single-connection experiment and measure unidirectional sending and receiving performance across a range of packet payload sizes. Because Demikernel's TCP implementation is optimized for latency, it performs worse than Linux on this experiment, so we configure Demikernel to use Linux TCP as its backend. The sending application sits in a tight loop, submitting data into the network stack as fast as possible; the receiver pulls data out of the network stack without doing further processing on it.

We vary whether the sender or the receiver is the FPGA or the CPU. The results are shown in Figure 7. We omit the (CPU send/FPGA receive) results, because they are almost the same as the all-CPU configuration; in both situations, the CPU sender is the bottleneck. The CPU is more efficient at streaming TCP data than UDP data because it allows batching data into jumbo frames. By contrast, Beehive's TCP stack is slower than its UDP stack, because of the complexity of stateful packet handling in hardware. In particular, our TCP engine is designed to only achieve full bandwidth across multiple simultaneous connections. Even so, Beehive outperforms Linux TCP across all request sizes. The speedup is most pronounced at small packet sizes, where Beehive achieves 2666 KReq/s versus the CPU's 843 KReq/s, a 3.2× speedup.

6.5 Reed-Solomon Encoding Acceleration

To evaluate Beehive's scaling architecture, we evaluate a duplicated Reed-Solomon (RS) encoding accelerator on Beehive versus a CPU implementation of the same algorithm. The client sends blocks of 4 KB to the encoder using UDP; the accelerator replies with 1K of erasure data. This could be organized into an (8,2) stripe for double fault tolerance. One instance of the Reed-Solomon encoder can consume data at 15 Gbps; our FPGA has room for four encoder instances, which consume data at 62 Gbps as shown in Table 2. For comparison, we use the open-source Reed-Solomon encoding implementation from BackBlaze [8] running on CPUs which we then duplicate across cores.

We also compare the energy efficiency of the two approaches in Table 2. The FPGA is about 20× more efficient per operation than the CPU implementation.

6.6 Viewstamped Replication Witness Acceleration

We next turn to a latency-sensitive application, evaluating Beehive hosting a viewstamped replication (VR) witness appliance. We first evaluate the witness on a single shard. We then take advantage of Beehive's ability to duplicate both internal components and applications to host a 4-shard witness appliance. We also duplicate protocol tiles to prevent them from becoming a bottleneck

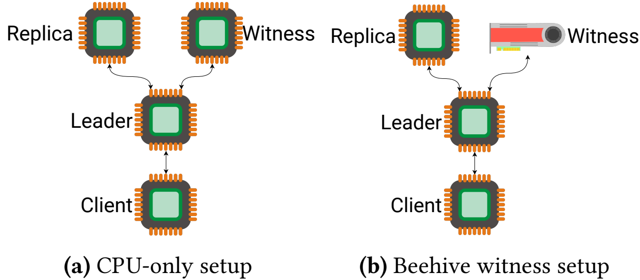

Figure 8. Experimental setups for VR evaluation

For all experiments, we evaluate a three-node VR configuration as shown in Figure 8 with either the FPGA or CPU serving as a witness. The other nodes run on CPU. The CPU VR replicas run on Intel Xeon Gold 5218 CPUs. Client threads run on Intel Xeon Gold 6226R CPUs. Each client thread sends one request and waits for a reply before sending another. We use UDP as our transport protocol, because VR does not assume reliable message delivery.

We evaluate our VR accelerator with a replicated key-value store application with 64-byte keys and 64-byte values. We use a read-write mix of 90% reads and 10% writes and a uniform key distribution. We set the IRQ affinity to the CPU socket where the replicas are running.

We test both single-shard and multishard configurations. Latency is measured at the clients as the time between the initial request and the eventual response. The shard leaders are distributed evenly between two CPU machines. The CPU witness(es) run on a separate server to allow us to measure the energy used by a CPU witness appliance.

Latency & throughput: We plot latency versus throughput for differing numbers of shards in Figure 9. We increase offered load by increasing the number of client threads sending requests to the leader. The results are shown in Figure 9. The system using the FPGA witness can provide up to 1.14× more per-core throughput and up to 1.13× lower median latency.

Energy: For each shard, we take the median energy measurement, throughput, median latency, and 99th-percentile (p99) latency at each circled point in Figure 9. The results are shown in Table 3. The FPGA is between 2.07× and 2.32× more energy efficient per operation compared to the CPU while providing better overall throughput and latency to key-value store clients.

6.7 TCP Connection Live Migration

To demonstrate Beehive's flexibility, we add support for live migration of a remote TCP endpoint to our hardware TCP stack by adding tiles and reusing other components.

For the remote TCP endpoint, we use Demikernel [61], because their TCP implementation is able to pause a connection, serialize the current state, and reinstall the state in

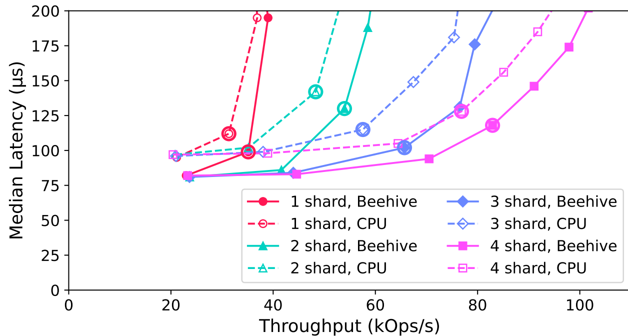

Figure 9. Latency vs. throughput for the key-value store workload varying the number of shards and client threads. The FPGA witness consistently outperforms the equivalent CPU cores in both latency and throughput.

Table 3. Energy per operation (measured at the witness) and performance metrics (measured at the clients) at the circled points in Figure 9

Shards	1	2	3	4
CPU Energy (mJ/op)	1.51	1.03	0.90	0.70
Beehive Energy (mJ/op)	0.73	0.48	0.39	0.31
Energy efficiency	2.07×	2.16×	2.32×	2.27×
CPU Throughput (kOps/s)	31	48	58	77
Beehive Throughput (kOps/s)	35	54	66	83
Speedup	1.12×	1.12×	1.14×	1.08×
CPU Median Latency (µs)	112	142	115	128
Beehive Median Latency (µs)	99	130	102	118
Improvement	1.13×	1.09×	1.13×	1.08×
CPU p99 Latency (µs)	273	372	339	412
Beehive p99 Latency (µs)	281	334	304	394
Improvement	0.97×	1.11×	1.12×	1.05×

the Demikernel network stack on a different node. Our test application is an echo RPC server. During the experiment, we run one closed-loop client on a CPU generating 64-byte payloads every 100 μs. The throughput around the migration event is shown in Figure 10. The migration latency, measured on the FPGA as the time between the last request from the first client and the first request from the second client, is 500 μs.

6.8 Hardware Resource Utilization

The hardware utilization of the Beehive infrastructure is shown in Table 4. For the UDP stack used in Section 6.3, Beehive components use 4% of the LUTs available on the Alveo U200 and 2% of the BRAMs. In a tile, a router uses around 6000 LUTs, twice the size of the UDP processing. For comparison with a more complex module, we include the utilization of the TCP receive path.

We also evaluate the scalability of the hardware implementation by building a UDP network stack as before and

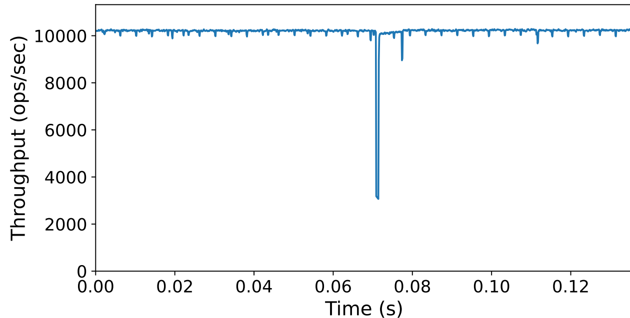

Figure 10. Request throughput over time. A single client sends requests every 100 μs to the server. Migration occurs around $t = 0.07$ s. From $t = 0$ s until migration, the TCP connection is on the first client. Throughput dips during migration, then recovers after 500 μs when the connection is brought up on the second client.

Table 4. FPGA resource utilization of selected modules.

	LUTs	BRAM
Beehive UDP full	57272	40
UDP RX Tile	9780	9.5
Router	5961	0
NoC Message Processing	515	0
UDP RX Processing	2984	9.5
TCP RX Engine	11672	16.5

then duplicating the echo application tile. On the Alveo U200, we can place 22 application tiles and 28 tiles total. We are limited by timing rather than resource utilization; the critical path is between NoC routers. Each router is fairly expensive, because the 512-bit width of the bus results in a number of high-fanout wires. This is exacerbated by the fact that the FPGA part in the Alveo U200 is made up of several chiplets, and chiplet crossings add significant delay. Several FPGAs [5, 60] now support hardened NoC resources and could improve the quality of results.

7 Related Work

7.1 Packet processing

PANIC [43] is a framework that supports integration of arbitrary packet processing elements, including general purpose cores. It uses a similar model to Beehive of chaining message-passing elements over a NoC, but it relies on a crossbar, limiting scalability. While the paper does not directly address deadlock, their central scheduler drops packets when it runs out of buffer space, preventing deadlock. This is not suitable for reliable network processing above layer 3. Because it aims to be a NIC framework, PANIC does not need to support higher-level protocols. For example, it does not support multi-packet application requests.

ClickNP [41] also supports the integration of arbitrary processing elements. However, it does not use a NoC. Instead components are directly connected via FIFOs, which makes it harder to replicate elements. It also assumes a PCIe connection to a CPU, which it relies on for control plane configuration.

Rosebud [36] is an FPGA framework for middleboxes. It uses an interconnect to connect custom processing elements they call reconfigurable processing units (RPUs) that can include accelerators. Because it targets middleboxes, they do not evaluate a network stack with full reliable transport protocol support. While it does provide support to chain RPUs, they acknowledge it was not designed to do so, and inter-RPU traffic has a fairly significant latency penalty.

A more restrictive approach leverages reconfigurable match-action tables. An action (e.g. strip a header, rewrite a field, drop a packet) is taken based on some header fields in the header of the packet. Typically, there is a pipeline of these processing elements [12, 27, 34]. However, match-action style processing is not well-suited for highly stateful processing [50] typical for application-level offloads. Other models have been proposed for stateful packet processing. Flowblaze uses an FSM-based model [50]. However, they specifically say that workloads above the transport layer are out of scope. hXDP proposed a processor for eBPF bytecode [13] designed for offloading kernel-level eBPF programs. Because of its sequential execution model, hXDP performs best on small programs and is a poor fit for more complex processing such as Reed-Solomon encoding.

7.2 Transport protocol offloads

Another related vein of work are transport protocol offloads. Most of these are TCP offload engines available as custom chips [16] or encrypted IP cores for FPGAs [2–4]. They generally do not support the full range of functions found in datacenter network stacks.

Some TCP offload engines could potentially support modification. Limago [53] is an open-source TCP and RoCEv2 offload engine written in Vivado HLS. However, it does not provide any specific APIs or hooks for adding other protocols, so introducing a new network function or new protocol would require fairly extensive modifications to the stack itself. Tonic [7] is an open-source implementation of the TCP send path and supports customization of the transport protocol, but does not address any lower-level packet processing layers; it also lacks a complete receive path implementation. FlexTOE [54] is a software implementation of TCP offload engine using the Netronome DPU, a processor designed specifically for network processing that is programmable using C or eBPF. While they do support network functions, their work targets TCP offload for CPUs while our work shows that a direct-attached hardware accelerator does not need a CPU core to support software stack functionality.

Microsoft Catapult's FPGAs use a custom transport protocol called LTL [14], which is a reliable transport protocol over UDP. Similar to most TCP engines, it is presented as a fixed IP core with no interface for extension. Catapult also supports a single-layer RMT, used for network virtualization [27]. However, it is unknown if these are ever combined and if so, how it would support new protocols or network functions.

8 Conclusion

Modern datacenter networking relies on a variety of network functions and protocols, but current hardware network stacks fall short on these features. As datacenters continue to offload computation to accelerators, it is becoming increasingly important to enable direct-attached accelerators to reduce network overhead. In this paper, we presented the design and implementation of Beehive, a NoC-based network stack for direct-attached accelerators designed to be customizable and to support the variety of protocols and management functions in datacenter networking. We demonstrated that Beehive can combine replicated protocol elements and replicated applications for higher bandwidth, consistent low latency, and minimal overhead. We have open-sourced Beehive for reuse.

References

[1] Kernel TLS operation. https://docs.kernel.org/networking/tls-offload.html. Accessed: 2022-6-28.

[2] TCP Offload engine – Offloading TCP into hardware. https://www.easics.com/tcp-offload-engine/. Accessed: 2022-6-28.

[3] TCP/IP Offload Engine 10/25G. https://www.xilinx.com/products/intellectual-property/1-y7rb2p.html#productspecs. Accessed: 2022-6-28.

[4] TCP/UDP/IP Network Protocol Accelerator Platform (NPAP). https://www.missinglinkelectronics.com/index.php/menu-products/menu-network-protocol-accelerator. Accessed: 2022-6-28.

[5] Achronix. Revolutionary New 2D Network-on-Chip. https://www.achronix.com/revolutionary-new-2d-network-chip. Accessed: 2022-7-4.

[6] Aditya Akella, Amin Vahdat, Arjun Singhvi, Behnam Montazeri, Dan Gibson, Hassan Wassel, Joel Scherpelz, Milo M. K. Martin, Monica C Wong-Chan, Moray Mclaren, Prashant Chandra, Rob Cauble, Sean Clark, Simon Sabato, and Thomas F. Wenisch. 1RMA: Re-Envisioning Remote Memory Access for Multi-Tenant Datacenters. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, page 708–721, New York, NY, USA, 2020.

[7] Mina Tahmasbi Arashloo, Alexey Lavrov, Manya Ghobadi, Jennifer Rexford, David Walker, and David Wentzlaff. Enabling Programmable Transport Protocols in High-Speed NICs. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, February 2020.

[8] Backblaze. JavaReedSolomon. https://github.com/Backblaze/JavaReedSolomon. accessed: 2023-11-28.

[9] Jonathan Balkind, Michael McKeown, Yaosheng Fu, Tri Nguyen, Yanqi Zhou, Alexey Lavrov, Mohammad Shahrad, Adi Fuchs, Samuel Payne, Xiaohua Liang, Matthew Matl, and David Wentzlaff. OpenPiton: An Open Source Manycore Research Framework. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '16, page 217–232, New York, NY, USA, 2016. Association for Computing Machinery.

[10] Jeffrey R Ballard, Ian Rae, and Aditya Akella. Extensible and scalable network monitoring using opensafe. *INM/WREN*, 10, 2010.

[11] Ethan Blanton, Dr. Vern Paxson, and Mark Allman. TCP Congestion Control. RFC 5681, September 2009.

[12] Pat Bosshart, Glen Gibb, Hun-Seok Kim, George Varghese, Nick McKeown, Martin Izzard, Fernando Mujica, and Mark Horowitz. Forwarding Metamorphosis: Fast Programmable Match-Action Processing in Hardware for SDN. *SIGCOMM Comput. Commun. Rev.*, 43(4):99–110, aug 2013.

[13] Marco Spaziani Brunella, Giacomo Belocchi, Marco Bonola, Salvatore Pontarelli, Giuseppe Siracusano, Giuseppe Bianchi, Aniello Cammarano, Alessandro Palumbo, Luca Petrucci, and Roberto Bifulco. hXDP: Efficient Software Packet Processing on FPGA NICs. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 973–990. USENIX Association, November 2020.

[14] Adrian Caulfield, Eric 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. A Cloud-Scale Acceleration Architecture. In *Proceedings of the 49th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 1–13. IEEE Computer Society, October 2016.

[15] Tushar Deepak Chandra, Robert Griesemer, and Joshua Redstone. Paxos made live - an engineering perspective (2006 invited talk). In *Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing*, 2007.

[16] Chelsio. Terminator 6 ASIC. https://www.chelsio.com/terminator-6-asic/. Accessed: 2019-10-24.

[17] Inho Cho, Ahmed Saeed, Joshua Fried, Seo Jin Park, Mohammad Alizadeh, and Adam Belay. Overload Control for u-scale RPCs with Breakwater. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 299–314. USENIX Association, November 2020.

[18] Cilium. L7 load balancing and url re-writing. https://docs.cilium.io/en/latest/gettingstarted/servicemesh/envoy-traffic-management/. Accessed: 2022-6-28.

[19] Cloudflare. What is mutual TLS (mTLS)? https://www.cloudflare.com/learning/access-management/what-is-mutual-tls/. Accessed: 2022-6-28.

[20] Natacha Crooks, Youer Pu, Lorenzo Alvisi, and Allen Clement. Seeing is believing: A client-centric specification of database isolation. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, page 73–82, New York, NY, USA, 2017. Association for Computing Machinery.

[21] Dally and Seitz. Deadlock-free message routing in multiprocessor interconnection networks. *IEEE Transactions on Computers*, C-36(5):547–553, 1987.

[22] William J. Dally and Brian Towles. Route Packets, Not Wires: On-Chip Inteconnection Networks. In *Proceedings of the 38th Annual Design Automation Conference*, DAC '01, page 684–689, New York, NY, USA, 2001. Association for Computing Machinery.

[23] Michael Dalton, David Schultz, Jacob Adriaens, Ahsan Arefin, Anshuman Gupta, Brian Fahs, Dima Rubinstein, Enrique Cauich Zermeno, Erik Rubow, James Alexander Docauer, Jesse Alpert, Jing Ai, Jon Olson, Kevin DeCabooter, Marc de Kruijf, Nan Hua, Nathan Lewis, Nikhil Kasinadhuni, Riccardo Crepaldi, Srinivas Krishnan, Subbaiah Venkata, Yossi Richter, Uday Naik, and Amin Vahdat. Andromeda: Performance, Isolation, and Velocity at Scale in Cloud Network Virtualization. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 373–387, Renton, WA, April 2018. USENIX

Association.

[24] Daniel E. Eisenbud, Cheng Yi, Carlo Contavalli, Cody Smith, Roman Kononov, Eric Mann-Hielscher, Ardas Cilingiroglu, Bin Cheyney, Wentao Shang, and Jinnah Dylan Hosein. Maglev: A Fast and Reliable Software Network Load Balancer. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 523–535, Santa Clara, CA, 2016.

[25] Envoy Proxy. https://www.envoyproxy.io/, 2022.

[26] Daniel Firestone. VFP: A virtual switch platform for host SDN in the public cloud. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 315–328, Boston, MA, March 2017. USENIX Association.

[27] Daniel Firestone, Andrew Putnam, Hari Angepat, Derek Chiou, Adrian Caulfield, Eric Chung, Matt Humphrey, Kalin Ovtcharov, Jitu Padhye, Doug Burger, Dave Maltz, Albert Greenberg, Sambhrama Mundkur, Alireza Dabagh, Mike Andrewartha, Vivek Bhanu, Harish Kumar Chandrappa, Somesh Chaturmohta, Jack Lavier, Norman Lam, Fengfen Liu, Gautham Popuri, Shachar Raindel, Tejas Sapre, Mark Shaw, Gabriel Silva, Madhan Sivakumar, Nisheeth Srivastava, Anshuman Verma, Qasim Zuhair, Deepak Bansal, Kushagra Vaid, and David A. Maltz. Azure Accelerated Networking: SmartNICs in the Public Cloud. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2018.

[28] Alex Forencich, Alex C. Snoeren, George Porter, and George Papen. Corundum: An Open-Source 100-Gbps NIC. In *28th IEEE International Symposium on Field-Programmable Custom Computing Machines*, 2020.

[29] Prateesh Goyal, Preey Shah, Kevin Zhao, Georgios Nikolaidis, Mohammad Alizadeh, and Thomas E Anderson. Backpressure Flow Control. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 779–805, 2022.

[30] Yutaro Hayakawa, Michio Honda, Douglas Santry, and Lars Eggert. Prism: Proxies without the Pain. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 535–549. USENIX Association, April 2021.

[31] Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, and Sergey Yekhanin. Erasure coding in windows azure storage. In *2012 USENIX Annual Technical Conference (USENIX ATC 12)*, pages 15–26, Boston, MA, June 2012. USENIX Association.

[32] Patrick Hunt, Mahadev Konar, Flavio P. Junqueira, and Benjamin Reed. ZooKeeper: Wait-free coordination for internet-scale systems. In *2010 USENIX Annual Technical Conference (USENIX ATC 10)*. USENIX Association, June 2010.

[33] Istio. Traffic management. https://istio.io/latest/docs/concepts/traffic-management/. Accessed: 2022-6-28.

[34] Antoine Kaufmann, SImon Peter, Naveen Kr. Sharma, Thomas Anderson, and Arvind Krishnamurthy. High Performance Packet Processing with FlexNIC. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '16, page 67–81, New York, NY, USA, 2016. Association for Computing Machinery.

[35] Osama Khan, Randal C Burns, James S Plank, William Pierce, and Cheng Huang. Rethinking erasure codes for cloud file systems: Minimizing I/O for recovery and degraded reads. In *10th USENIX Conference on File and Storage Technologies (FAST 12)*, San Jose, CA, February 2012. USENIX Association.

[36] Moein Khazraee, Alex Forencich, George C. Papen, Alex C. Snoeren, and Aaron Schulman. Rosebud: Making fpga-accelerated middlebox development more pleasant. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, ASPLOS 2023, page 586–605, New York, NY, USA, 2023. Association for Computing Machinery.

[37] Hyojoon Kim and Nick Feamster. Improving network management with software defined networking. *IEEE Communications Magazine*, 51(2):114–119, 2013.

[38] Teemu Koponen, Keith Amidon, Peter Balland, Martin Casado, Anupam Chanda, Bryan Fulton, Igor Ganichev, Jesse Gross, Paul Ingram, Ethan Jackson, Andrew Lambeth, Romain Lenglet, Shih-Hao Li, Amar Padmanabhan, Justin Pettit, Ben Pfaff, Rajiv Ramanathan, Scott Shenker, Alan Shieh, Jeremy Stribling, Pankaj Thakkar, Dan Wendlandt, Alexander Yip, and Ronghua Zhang. Network Virtualization in Multi-tenant Datacenters. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 203–216, Seattle, WA, April 2014. USENIX Association.

[39] Gautam Kumar, Nandita Dukkipati, Keon Jang, Hassan MG Wassel, Xian Wu, Behnam Montazeri, Yaogong Wang, Kevin Springborn, Christopher Alfeld, Michael Ryan, et al. Swift: Delay is Simple and Effective for Congestion Control in the Datacenter. In *Proceedings of the ACM SIGCOMM 2020 Conference*, pages 514–528, 2020.

[40] Andreas Lankes, Thomas Wild, Andreas Herkersdorf, Soeren Sonntag, and Helmut Reinig. Comparison of Deadlock Recovery and Avoidance Mechanisms to Approach Message Dependent Deadlocks in On-chip Networks. In *2010 Fourth ACM/IEEE International Symposium on Networks-on-Chip*, pages 17–24, 2010.

[41] Bojie Li, Kun Tan, Layong (Larry) Luo, Yanqing Peng, Renqian Luo, Ningyi Xu, Yongqiang Xiong, Peng Cheng, and Enhong Chen. ClickNP: Highly Flexible and High Performance Network Processing with Reconfigurable Hardware. In *Proceedings of the 2016 ACM SIGCOMM Conference*, SIGCOMM '16, page 1–14, New York, NY, USA, 2016. Association for Computing Machinery.

[42] Yuliang Li, Rui Miao, Hongqiang Harry Liu, Yan Zhuang, Fei Feng, Lingbo Tang, Zheng Cao, Ming Zhang, Frank Kelly, Mohammad Alizadeh, and Minlan Yu. HPCC: High Precision Congestion Control. In *Proceedings of the ACM SIGCOMM 2019 Conference*, page 44–58, 2019.

[43] Jiaxin Lin, Kiran Patel, Brent E. Stephens, Anirudh Sivaraman, and Aditya Akella. PANIC: A High-Performance Programmable NIC for Multi-tenant Networks. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 243–259. USENIX Association, November 2020.

[44] Barbara Liskov and James Cowling. Viewstamped replication revisited. Technical Report MIT-CSAIL-TR-2012-021, MIT, July 2012.

[45] Michael Marty, Marc de Kruijf, Jacob Adriaens, Christopher Alfeld, Sean Bauer, Carlo Contavalli, Mike Dalton, Nandita Dukkipati, William C. Evans, Steve Gribble, Nicholas Kidd, Roman Kononov, Gautam Kumar, Carl Mauer, Emily Musick, Lena Olson, Mike Ryan, Erik Rubow, Kevin Springborn, Paul Turner, Valas Valancius, Xi Wang, and Amin Vahdat. Snap: a Microkernel Approach to Host Networking. In *In ACM SIGOPS 27th Symposium on Operating Systems Principles*, New York, NY, USA, 2019.

[46] Microsoft. Create security policies with extended port access control lists. https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v-virtual-switch/create-security-policies-with-extended-port-access-control-lists. Accessed: 2022-6-28.

[47] Behnam Montazeri, Yilong Li, Mohammad Alizadeh, and John Ousterhout. Homa: A receiver-driven low-latency transport protocol using network priorities. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, page 221–235, New York, NY, USA, 2018. ACM.

[48] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, USENIX ATC'14, page 305–320, USA, 2014. USENIX Association.

[49] Parveen Patel, Deepak Bansal, Lihua Yuan, Ashwin Murthy, Albert Greenberg, David A. Maltz, Randy Kern, Hemant Kumar, Marios Zikos, Hongyu Wu, Changhoon Kim, and Naveen Karri. Ananta: Cloud Scale Load Balancing. *SIGCOMM Comput. Commun. Rev.*, 43(4):207–218, aug 2013.

[50] Salvatore Pontarelli, Roberto Bifulco, Marco Bonola, Carmelo Cascone, Marco Spaziani, Valerio Bruschi, Davide Sanvito, Giuseppe Siracusano, Antonio Capone, Michio Honda, Felipe Huici, and Giuseppe Siracusano. FlowBlaze: Stateful Packet Processing in Hardware. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 531–548, Boston, MA, February 2019. USENIX Association.

[51] Dan R. K. Ports, Jialin Li, Vincent Liu, Naveen Kr. Sharma, and Arvind Krishnamurthy. Designing distributed systems using approximate synchrony in data center networks. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 43–57, Oakland, CA, May 2015. USENIX Association.

[52] K. V. Rashmi, Nihar B. Shah, Dikang Gu, Hairong Kuang, Dhruba Borthakur, and Kannan Ramchandran. A solution to the network challenges of data recovery in erasure-coded distributed storage systems: A study on the facebook warehouse cluster. In *5th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 13)*, San Jose, CA, June 2013. USENIX Association.

[53] Mario Ruiz, David Sidler, Gustavo Sutter, Gustavo Alonso, and Sergio López-Buedo. Limago: An FPGA-Based Open-Source 100 GbE TCP/IP Stack. In *29th International Conference on Field Programmable Logic and Applications (FPL)*, pages 286–292, 2019.

[54] Rajath Shashidhara, Tim Stamler, Antoine Kaufmann, and Simon Peter. FlexTOE: Flexible TCP Offload with Fine-Grained Parallelism. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 87–102, Renton, WA, April 2022. USENIX Association.

[55] Nikita Shirokov and Ranjeeth Dasineni. Open-sourcing Katran, a scalable network load balancer. https://engineering.fb.com/2018/05/22/open-source/open-sourcing-katran-a-scalable-network-load-balancer/. Accessed: 2021-10-26.

[56] Yong Ho Song and T.M. Pinkston. A progressive approach to handling message-dependent deadlock in parallel computer systems. *IEEE Transactions on Parallel and Distributed Systems*, 14(3):259–275, 2003.

[57] Praveen Tammana, Rachit Agarwal, and Myungjin Lee. Simplifying datacenter network debugging with {PathDump}. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 233–248, 2016.

[58] David Wragg. Unimog - cloudflare's edge load balancer. https://blog.cloudflare.com/unimog-cloudflares-edge-load-balancer/. Accessed: 2021-10-26.

[59] Xilinx. Alveo card management solution subsystem product guide. https://docs.xilinx.com/r/en-US/pg348-cms-subsystem/Register-Space. Accessed: 2023-11-29.

[60] Xilinx. Versal premium series. https://www.xilinx.com/products/silicon-devices/acap/versal-premium.html. Accessed: 2022-7-4.

[61] Irene Zhang, Amanda Raybuck, Pratyush Patel, Kirk Olynyk, Jacob Nelson, Omar S. Navarro Leija, Ashlie Martinez, Jing Liu, Anna Kornfeld Simpson, Sujay Jayakar, Pedro Henrique Penna, Max Demoulin, Piali Choudhury, and Anirudh Badam. The Demikernel Datapath OS Architecture for Microsecond-Scale Datacenter Systems. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles*, SOSP '21, page 195–211, New York, NY, USA, 2021. Association for Computing Machinery.

[62] Yiwen Zhang, Gautam Kumar, Nandita Dukkipati, Xian Wu, Priyaranjan Jha, Mosharaf Chowdhury, and Amin Vahdat. Aequitas: Admission control for performance-critical rpcs in datacenters. In *Proceedings of the ACM SIGCOMM 2022 Conference*, SIGCOMM '22, page 1–18, New York, NY, USA, 2022. Association for Computing Machinery.

[63] Zhipeng Zhao, Hugo Sadok, Nirav Atre, James C. Hoe, Vyas Sekar, and Justine Sherry. Achieving 100Gbps Intrusion Prevention on a Single Server. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 1083–1100. USENIX Association, November 2020.