

Low latency stream processing: Apache Heron with Infiniband & Intel Omni-Path

Supun Kamburugamuve
School of Informatics, Computing, and Engineering
Indiana University
Bloomington, IN 47408
skamburu@indiana.edu

Martin Swany
School of Informatics, Computing, and Engineering
Indiana University
Bloomington, IN 47408
swany@indiana.edu

Karthik Ramasamy
Streaml.io
Palo Alto, CA
karthik@streaml.io

Geoffrey Fox
School of Informatics, Computing, and Engineering
Indiana University
Bloomington, IN 47408
gcf@indiana.edu

ABSTRACT

Worldwide data production is increasing both in volume and velocity, and with this acceleration, data needs to be processed in streaming settings as opposed to the traditional store and process model. Distributed streaming frameworks are designed to process such data in real time with reasonable time constraints. Apache Heron is a production-ready large-scale distributed stream processing framework. The network is of utmost importance to scale streaming applications to large numbers of nodes with a reasonable latency. High performance computing (HPC) clusters feature interconnects that can perform at higher levels than traditional Ethernet. In this paper the authors present their findings on integrating Apache Heron distributed stream processing system with two high performance interconnects; Infiniband and Intel Omni-Path and show that they can be utilized to improve performance of distributed streaming applications.

KEYWORDS

Streaming data, Infiniband, Omni-Path, Apache Heron

1 INTRODUCTION

With ever increasing data production by users and machines alike, the amount of data that needs to be processed has increased dramatically. This must be achieved both in real time and as batches to satisfy different use cases. Additionally, with the adoption of devices into Internet of Things setups, the amount of real time data are exploding, and must be processed with reasonable time constraints. In distributed stream analytics, the large data streams are partitioned and processed in distributed sets of machines to keep up with the high volume data rates. By definition of large-scale streaming data processing, networks are a crucial component in

transmitting messages between the processing units for achieving efficient data processing.

There are many hardware environments in which big data systems can be deployed including High performance computing (HPC) clusters. HPC clusters are designed to perform large computations with advanced processors, memory, IO systems and high performance interconnects. High performance interconnects in HPC clusters feature microsecond latencies and large bandwidths. Thanks to recent advancements in hardware, some of these high performance networks have become cheaper to set up than their Ethernet counterparts. With multi-core and many-core systems having large numbers of CPUs in a single node, the demand for high performance networking is increasing as well.

Advanced hardware features such as high performance interconnects are not fully utilized in the big data computing frameworks, mostly because they are accessible to low level programming languages and most big data systems are written on Java platform. In recent years we have seen efforts to utilize high performance interconnects into big data frameworks such as Spark [24] and Hadoop [22]. Big data frameworks such as Spark and Hadoop focus on large batch data processing and hence their communication requirements are different compared to streaming systems which are more latency sensitive.

There are many distributed streaming frameworks available today for processing large amounts of streaming data in real time. Such systems are largely designed and optimized for commodity hardware and clouds. Apache Storm [29] was one of the popular early systems developed for processing streaming data. Apache Heron¹ [19] is similar to Storm with a new architecture for streaming data processing. It features a hybrid design with some of the performance-critical parts written in C++ and others written in Java. This architecture allows the integration of high performance enhancements directly rather than going through native wrappers such as Java Native Interface(JNI). When these systems are deployed on clusters that include high performance interconnects, they need to use an TCP interface to high performance interconnect which doesn't perform as well as a native implementation.

To utilize these hardware features, we have integrated Infiniband and Intel Omni-Path interconnects to Apache Heron to accelerate

¹<http://incubator.apache.org/projects/heron.html>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UCC'17, December 5–8, 2017, Austin, Texas, USA

© 2017 ACM. ISBN 978-1-4503-5149-2/17/12...\$15.00

DOI: <http://dx.doi.org/10.1145/3147213.3147232>

its communications. Infiniband [4] is an open standard protocol for high performance interconnects that is widely used in today’s high performance clusters. Omni-Path [6] is a proprietary interconnect developed by Intel and is available with the latest Knights Landing architecture-based (KNL) [27] many-core processors. With this implementation, we have observed significantly lower latencies and improved throughput in Heron. The main contribution in this work is to showcase the benefits of using high performance interconnects for distributed stream processing. There are many differences in hardware available for communications with different bandwidths, latencies and processing models. Ethernet has comparable hardware available to some of the high performance interconnects; it is not our goal to show that one particular technology is superior to others, as different environments may have alternate sets of these technologies.

The remainder of the paper is organized as follows. Section 2 presents the background information on Infiniband and Omni-Path. Section 3 describes the Heron architecture in detail and section 4 the implementation details. Next the experiments conducted are described in sections 5 and results are presented and discussed in section 6. Section 7 presents related work. The paper concludes with a look at future work.

2 BACKGROUND

2.1 Infiniband

Infiniband is one of the most widely used high performance fabrics. It provides a variety of capabilities including message channel semantics, remote memory access and remote atomic memory operations, supporting both connection-oriented and connectionless endpoints. Infiniband is programmed using the Verbs API, which is available in all major platforms. The current hardware is capable of achieving up to 100Gbps speeds with microsecond latencies. Infiniband does not require the OS Kernel intervention to transfer packets from user space to the hardware. Unlike in TCP, its protocol aspects are handled by the hardware. These features mean less CPU time spent on the network compared to TCP for transferring the same amount of data. Because the OS Kernel is bypassed by the communications, the memory for transferring data has to be registered in the hardware.

2.2 Intel Omni-Path

Omni-Path is a high performance fabric developed by Intel. Omni-Path fabric is relatively new compared to Infiniband and there are fundamental differences between the two. Omni-Path does not offload the protocol handling to network hardware and it doesn’t have the connection oriented channels as in Infiniband. Unlike in Infiniband the Omni-Path network chip can be built into the latest Intel Knights Landing (KNL) processors. Omni-Path supports tagged messaging with a 96 bit tag in each message. A Tag can carry any type of data and this information can be used at the application to distinguish between different messages. Omni-Path is designed and optimized for small high frequency messaging.

2.3 Channel & Memory Semantics

High performance interconnects generally supports two modes of operations called channel and memory semantics [13]. With

channel semantics queues are used for communication. In memory semantics a process can read from or write directly to the memory of a remote machine. In channel mode, two queue pairs for transmission and receive operations are used. To transfer a message, a descriptor is posted to the transfer queue, which includes the address of the memory buffer to transfer. For receiving a message, a descriptor needs to be submitted along with a pre-allocated receive buffer. The user program queries the completion queue associated with a transmission or a receiving queue to determine the success or failure of a work request. Once a message arrives, the hardware puts the message into the posted receive buffer and the user program can determine this event through the completion queue. Note that this mode requires the receiving buffers to be pre-posted before the transmission can happen successfully.

With memory semantics, Remote Direct Memory Access (RDMA) operations are used. Two processes preparing to communicate register memory and share the details with each other. Read and write operations are used instead of send and receive operations. These are one-sided and do not need any software intervention from the other side. If a process wishes to write to remote memory, it can post a write operation with the local addresses of the data. The completion of the write operation can be detected using the completion queue associated. The receiving side is not notified about the write operation and has to use out-of-band mechanisms to figure out the write. The same is true for remote reads as well. RDMA is more suitable for large message transfers while channel mode is suitable for small messages. In general RDMA has 1 – 2 μ s latency advantage over channel semantics for Infiniband and this is not significant for our work.

2.4 Openfabrics API

Openfabrics² provides a library called libfabric [12] that hides the details of common high performance fabric APIs behind a uniform API. Because of the advantage of such an API, we chose to use libfabric as our programming library for implementing the high performance communications for Heron. Libfabric is a thin wrapper API and it supports different providers including Verbs, Aries interconnect from Cray through GNI, Intel Omni-Path, and Sockets.

2.5 TCP & High performance Interconnects

TCP is one of the most successful protocols developed. It provides a simple yet powerful API for transferring data reliably across the Internet using unreliable links and protocols underneath. One of the biggest advantages of TCP is its wide adoption and simplicity to use. Virtually every computer has access to a TCP-capable adapter and the API is well supported across different platforms. TCP provides a streaming API for messaging where the fabric does not maintain message boundaries. The messages are written as a stream of bytes to the TCP and the application has to define mechanisms such as placing markers in between messages to indicate the boundaries. On the other hand, Infiniband and Omni-Path both support message boundaries.

High performance interconnects have drivers that make them available through the TCP protocol stack. The biggest advantage of such implementation is that an existing application written using

²<https://www.openfabrics.org/>

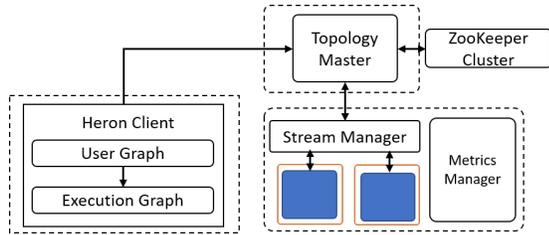


Figure 1: High level architecture of Heron. Each outer box shows a resource container allocated by a resource scheduler like Mesos or Slurm. The arrows show the communication links between different components.

the TCP stack can use high performance interconnect without any modifications to the code. It is worth noting that the native use of the interconnect through its API always yields better performance than using it through TCP/IP stack. A typical TCP application allocates memory in user space and the TCP stack needs to copy data between user space and Kernel space. Also each TCP invocation involves a system call which does a context switch of the application. Recent advancements like Netmap [26] and DPDK [11] removes these costs and increases the maximum number of packets that can be processed per second. Also it is possible to achieve direct packet copy from user space to hardware. High performance fabrics usually do not go through OS kernel for network operations and the hardware is capable of copying data directly to user space buffers. Infiniband offloads the protocol processing aspects to hardware while Omni-Path still involves the CPU for protocol processing.

3 APACHE HERON

Heron is a distributed stream processing framework developed at Twitter and now available as an open source project in Apache Incubator³. Heron is similar to Apache Storm in its API with many differences in underlying engine architecture. It retains the same Storm API, allowing applications written in Storm to be deployed with no or minimal code changes.

3.1 Heron Data Model

A stream is an unbounded sequence of high level objects named events or messages. The streaming computation in Heron is referred to as a Topology. A topology is a graph of nodes and edges. The nodes represent the processing units executing the user defined code and the edges between the nodes indicate how the data (or stream) flows between them. There are two types of nodes: spouts and bolts. Spouts are the sources of streams. For example, a Kafka [18] spout can read from a Kafka queue and emit it as a stream. Bolts consume messages from their input stream(s), apply its processing logic and emit new messages in their outgoing streams.

Heron has the concept of a user defined graph and an execution graph. The user defined graph specifies how the processing units are connected together in terms of message distributions. On the other hand, the execution graph is the layout of this graph in actual nodes with network connections and computing resources allocated

to the topology to execute. Nodes in the execution graph can have multiple parallel instances (or tasks) running to scale the computations. The user defined graph and the execution graph are referred to as logical plan and physical plan respectively.

3.2 Heron Architecture

The components of the Heron architecture are shown in Fig. 1. Each Heron topology is a standalone long-running job that never terminates due to the unbounded nature of streams. Each topology is self contained and executes in a distributed sandbox environment in isolation without any interference from other topologies. A Heron topology consists of multiple containers allocated by the scheduler. These can be Linux containers, physical nodes or sandboxes created by the scheduler. The first container, referred to as the master, always runs the Topology Master that manages the topology. Each of the subsequent containers have the following processes: a set of processes executing the spout/bolt tasks of the topology called Heron instances, a process called a stream manager that manages the data routing and the connections to the outside containers, and a metrics manager to collect information about the instances running in that container.

Each Heron instance executes a single task of the topology. The instances are connected to the stream manager running inside the container through TCP loop-back connection. It is worth noting that Heron instances always connect to other instances through the stream manager and they do not communicate with each other directly even if they are on the same container. The stream manager acts as a bridge between Heron instances. It forwards the messages to the correct instances by consulting the routing tables it maintains. A message between two instances in different containers goes through two stream managers. Containers can have many Heron instances running in them and they all communicate through the stream manager. Because of this design, it is important to have highly efficient data transfers at the stream manager to support the communication requirements of the instances.

Heron is designed from the ground up to be extensible, and important parts of the core engine are written in C++ rather than Java, the default language of choice for Big Data frameworks. The rationale for the use of C++ is to leverage the advanced features offered by the OS and hardware. Heron instances and schedulers are written in Java while stream manager and topology master are written in C++.

3.2.1 Acknowledgements. Heron uses an acknowledgement mechanism to provide at least once message processing semantics that ensures the message is always processed in the presence of process/machine failures. In order to achieve at least once, the stream manager tracks the tuple tree generated while a message progresses through the topology. When a bolt emits a message, it anchors the new message to the parent message and this information is sent to originating stream manager (in the same container) as a separate message. When every new message finishes its processing, a separate message is again sent to the originating stream manager. Upon receiving such control messages for every emit in the message tree, the stream manager marks the message as fully processed.

³<https://github.com/twitter/heron>

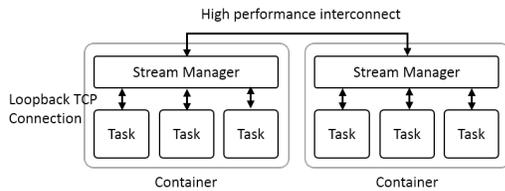


Figure 2: Heron high performance interconnects are between the stream managers

3.2.2 Processing pipeline. Heron has a concept called max messages pending with spouts. When a spout emits messages to a topology, this number dictates the amount of in-flight messages that are not fully processed yet. The spout is called to emit messages only when the current in-flight message count is less than the max spout pending messages.

3.3 Heron Stream Manager

Stream manager is responsible for routing messages between instances inside a container and across containers. It employs a single thread that use event-driven programming using non-blocking socket API. A stream manager receives messages from instances running in the same container and other stream managers. These messages are Google protocol buffer [30] serialized messages packed into binary form and transmitted through the wire. If a stream manager receives a message from a spout, it keeps track of the details of the message until all the acknowledgements are received from the message tree. Stream manager features a in-memory, store and forward architecture for forwarding messages and can batch multiple messages into single message for efficient transfers. Because messages are temporarily stored, there is a draining function that drains the store at a user defined rate.

4 IMPLEMENTATION

Even though high performance interconnects are widely used by HPC applications and frameworks, they are not often utilized in big data systems. Furthermore, experiences in using these interconnects in big data systems are lacking in the public domain. In this implementation, Infiniband and Omni-Path interfaces with Heron through its stream manager, as shown in Fig. 2. Infiniband or Omni-Path message channels are created between each stream manager in the topology. These then carry the data messages going between the stream managers. The control messages that are sent between stream manager and topology master still use the TCP connections. They are not frequent and do not affect the performance of the data flow. The TCP loop-back connections from the instances to the stream manager are not altered in this implementation as loop back connection is much more efficient than the network. Both Infiniband and Omni-Path implementations use channel semantics for communication. A separate thread is used for polling the completion queues associated with the channels. A credit based flow control mechanism is used for each channel along with a configurable buffer pool.

4.1 Bootstrapping

Infiniband and Omni-Path require information about the communication parties to be sent out-of-band through other mechanisms like TCP. Infiniband uses the RDMA(Remote direct memory access) Connection manager to transfer the required information and establish the connections. RDMA connection manager provides a socket-like API for connection management, which is exposed to the user in a similar fashion through Libfabric API. The connection manager also uses the IP over Infiniband network adapter to discover and transfer the bootstrap information. Omni-Path has a built-in TCP server for discovering the endpoints. Because Omni-Path does not involve connection management, only the destination address is needed for communications. This information can be sent using an out-of-band TCP connection.

4.2 Buffer management

Each side of the communication uses buffer pools with equal size buffers to communicate. Two such pools are used for sending and receiving data for each channel. For receiving operations, all the buffers are posted at the beginning to the fabric. For transmitting messages, the buffers are filled with messages and posted to the fabric for transmission. After the transmission is complete the buffer is added back to the pool. The message receive and transmission completions are discovered using the completion queues. Individual buffer sizes are kept relatively large to accommodate the largest messages expected. If the buffer size is not enough for a single message, the message is divided in to pieces and put into multiple buffers. Every network message carries the length of the total message and this information can be used to assemble the pieces.

The stream manager de-serializes the protocol buffer message in order to determine the routing for the message and handling the acknowledgements. The TCP implementation first copies the incoming data into a buffer and then use this buffer to build the protocol buffer structures. This implementation can directly use the buffer allocated for receiving to build the protocol message.

4.3 Flow control at communication level

Neither Infiniband nor Omni-Path implement flow control between the communication parties, and it is up to the application developer to implement the much higher level functions [33]. This implementation uses a standard credit-based approach for flow control. The credit available for sender to communicate is equal to the number of buffers posted into the fabric by the receiver. Credit information is passed to the other side as part of data transmissions, or by using separate messages in case there are no data transmissions to send it. Each data message carries the current credit of the communication party as a 4-byte integer value. The credit messages do not take into account the credit available to avoid deadlocks, otherwise there may be situations where there is no credit available to send credit messages.

4.4 Interconnects

Infiniband implementation uses connection-oriented endpoints with channel semantics to transfer the messages. The messages are transferred reliably by the fabric and the message ordering is

guaranteed. The completions are also in order of the work request submissions.

Intel Omni-Path does not support connection-oriented message transfers employed in the Infiniband implementation. The application uses reliable datagram message transfer with tag-based messaging. Communication channels between stream managers are overlaid on a single receive queue and a single send queue. Messages coming from different stream managers are distinguished based on the tag information they carry. The tag used in the implementation is a 64-bit integer which carries the source stream manager ID and the destination stream manager ID. Even though all the stream managers connecting to a single stream manager are sharing a single queue, they carry their own flow control by assigning a fixed amount of buffers to each channel. Unlike in Infiniband, the work request completions are not in any order of their submission to the work queue. Because of this, the application keeps track of the submitted buffers and their completion order explicitly.

IP over Fabric or IP over Infiniband (IPoIB) is a mechanism to allow a regular TCP application to access the underlying high performance interconnects through the TCP Socket API. For using IPoIB heron stream manager TCP sockets are bound to the IPoIB network interface explicitly without changing the existing TCP processing logic.

5 EXPERIMENTS

An Intel Haswell HPC cluster was used for the Infiniband experiments. The CPUs are Intel Xeon E5-2670 running at 2.30GHz. Each node has 24 cores (2 sockets x 12 cores each) with 128GB of main memory, 56Gbps Infiniband interconnect and 1Gbps dedicated Ethernet connection to other nodes. Intel Knights Landing (KNL) cluster was used for Omni-Path tests. Each node in KNL cluster has 72 cores (Intel Xeon Phi CPU 7250F, 1.40GHz) and is connected to a 100Gbps Omni-Path fabric and 1Gbps Ethernet connection. There are many variations of Ethernet, Infiniband and Omni-Path performing at different message rates and latencies. We conducted the experiments in the best available resources to us.

We conducted several micro-benchmarks to measure the latency and throughput of the system. In these experiments the primary focus was given to communications and no computation was conducted in the bolts. The tasks in each experiment were configured with 4GB of memory. A single Heron stream manager was run in each node. We measured the IP over Fabric (IPoIB) latency to showcase the latency possible in case no direct implementation of Infiniband or Omni-Path.

5.1 Experiment Topologies

To measure the behavior of the system, two topologies shown in Fig. 4 and Fig. 3 is used. Topology A in Fig. 3 is a deep topology with multiple bolts arranged in a chain. The parallelism of the topology determines the number of parallel task for bolts and spout in the topology. Each adjacent component pair is connected by a shuffle grouping. Topology B in Fig. 4 is a two-component topology with a spout and a bolt. Spouts and bolts are arranged in a shuffle grouping so that the spouts load balance the messages among the bolts.

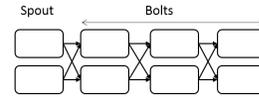


Figure 3: Topology A: Deep topology with a spout and multiple bolts arranged in a chain. The spout and bolt run multiple parallel instances.

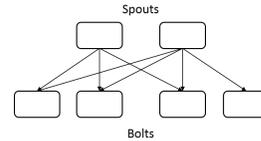


Figure 4: Topology B: Shallow topology with a spout and bolt connected in a shuffle grouping. The spout and bolts run multiple parallel instances.

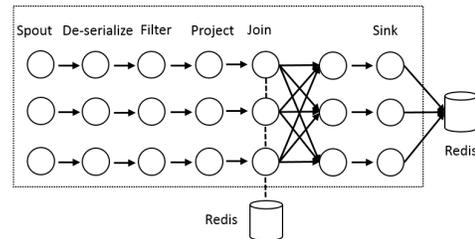


Figure 5: Yahoo Streaming benchmark with 7 stages. The join bolts and sink bolts communicate with a Redis server.

In both topologies the spouts generated messages at the highest sustainable speed with acknowledgements. The acknowledgements acted as a flow control mechanism for the topology. The latency is measured as the time it takes for a tuple to go through the topology and its corresponding acknowledgement to reach the spout. Since tuples generate more tuples when they go through the topology, it takes multiple acknowledgements to complete a tuple. In topology A, it needs control tuples equal to the number of Bolts in the chain to complete a single tuple. For example if the length of the topology is 8, it takes 7 control tuples to complete the original tuple. Tests were run with 10 in flight messages through the topology.

5.2 Yahoo streaming benchmark

For evaluating the behavior with a more practical streaming application, we used a benchmark developed at Yahoo!⁴ to test streaming frameworks. We modified the original streaming benchmark to support Heron and added additional features to support our case. The modified benchmark is available open source in Github⁵. It focuses on an advertisement application where ad events are processed using a streaming topology as shown in Fig. 10. We changed the original benchmark to use a self-message-generating spout instead

⁴<https://yahoeng.tumblr.com/post/135321837876/benchmarking-streaming-computation-engines-at>

⁵<https://github.com/iotcloud/streaming-benchmarks.git>

of reading messages through Kafka [18]. This was done to remove any bottlenecks and variations imposed by Kafka.

The benchmark employs a multiple stage topology with different processing units. The data is generated as a JSON object and sent through the processing units, which do de-serialization, filter, projection and join operations on each tuple. At the joining stage, it uses a Redis [8] database to query data about the tuples. Additionally, the last bolt saves information to the Redis database. At the filter step about 66% of the tuples are dropped. We use an acking enabled topology and measured the latency as the time it takes for a tuple to go through the topology and its ack to return back to the spout. For our tests we used 16 nodes with 8 parallelism at each step, totaling 56 tasks. Each spout was sending 100,000 messages per second, which gave 800,000 messages per second in total. Note that this topology accesses the Redis database system for 33% messages it receives.

6 RESULTS & DISCUSSION

Fig. 6 shows the latency of the Topology A with varying message sizes and parallelism for IB, TCP and IPoIB. Infiniband performed the best while Ethernet showed the highest latency. Fig. 7 shows the latency of the Topology B with varying message sizes and parallelism as in Fig. 6. For small messages, we see both IPoIB and Ethernet performing at a similar level while Infiniband was performing better. When increasing the parallelism, the latency increased as expected and Infiniband showed a smaller increase than IPoIB and Ethernet. As expected Topology B showed less improvements compared to Topology A. For most practical applications, the minimum latency possible will be within the results observed in these two topologies as these represent the minimum possible topology and a deep topology for most practical applications.

Fig. 8 shows the latency results of Omni-Path implementation with Topology A, in the KNL cluster for large and small messages. IPoFabric is the IP driver for Omni-Path. The KNL machine has less powerful CPU cores and hence the latency was higher compared to Infiniband in the tests. Fig 10 shows the latency distribution of the Topology A with 100 and 10 inflight messages with 128k message size for TCP and IB. The graphs show that the network latencies are according to a normal distribution with high latencies at the 99th percentile. Fig. 11 shows the latency distribution seen by the Yahoo stream benchmark with Infiniband fabric. For all three networking modes, we have seen high spikes at the 99th percentile. This was primarily due to Java garbage collections at the tasks which is unavoidable in JVM-based streaming applications. The store and forward functionality of the stream manager contributes to the distribution of latencies as a single message can be delayed up to 1ms randomly at stream manager.

Fig. 9 present the message rates observed with Topology B. The experiment was conducted with 32 parallel bolt instances and varying numbers of spout instances and message sizes. The graph shows that Infiniband had the best throughput, while IPoIB achieved second-best and Ethernet came in last. Fig. 9 c) and d) show the throughput for 128K and 128 bytes messages with varying number of spouts. When the number of parallel spouts increases, the IPoIB and Ethernet maxed out at 16 parallel spouts for 128k messages, while Infiniband kept on increasing. Fig. 12 shows the total time

required to serialize protocol buffers and time it took to complete the same number of messages with TCP and IB implementations. It is evident that IB implementation overhead is much less compared to TCP.

The results showed good overall results for Infiniband and Omni-Path compared to the TCP and IPoIB communication modes. The throughput of the system is bounded by the CPU usage of stream managers. For TCP connections, the CPU is used for message processing as well as the network protocol. Infiniband, on the other hand, uses the CPU only for message processing at the stream manager, yielding better performance.

The results showed much higher difference between TCP and Infiniband for large message sizes. For small messages the bandwidth utilization is much lower than large messages. This is primarily due to the fact that CPU is needed to process every message. For larger messages, because of the low number of messages transferred per second, the CPU usage is low for that aspect. For smaller messages, because of the large number of messages per second, the CPU usage is much higher. Because of this, for small messages, the stream managers saturate the CPU without saturating the communication channel. For practical applications that require large throughput, Heron can bundle small messages into a large message in-order to avoid some of the processing overheads and transfer overheads. This makes it essentially a large message for the stream manager and large message results can be observed with elevated latencies for individual messages. We used 10 buffers with 1 megabyte allocated to each for sending and receiving messages. Protocol buffers are optimized for small messages and it is unlikely to get more than 1 megabyte messages for streaming applications targeting Heron.

The KNL system used for testing Omni-Path has a large number of processes with less CPU frequencies compared to Haswell cluster. In order to fully utilize such a system, multiple threads need to be used. For our implementation we did not explore such features specific to KNL and tried to first optimize for the Omni-Path interconnect. The results show that Omni-Path performed considerably better than the other two options. In this work the authors did not try to pick between Omni-Path or Infiniband as a better interconnect as they are tested in two completely different systems under varying circumstances. The objective of the work is to show the potential benefits of using interconnects to accelerate stream processing.

7 RELATED WORK

In large part, HPC and Big Data systems have evolved independently over the years. Despite this, there are common requirements that raise similar issues in both type of systems. Some of these issues are solved in HPC and some in Big Data frameworks. As such, there have been efforts to converge both HPC and Big Data technologies to create better systems that can work in different environments efficiently. SPIDAL [9] and HiBD are two such efforts to enhance the Big Data frameworks with ideas and tools from HPC. This work is part of an ongoing effort by the authors to improve stream engine performance using HPC techniques. Previously we showed [15] how to leverage shared memory and collective communication algorithms to increase the performance of Apache Storm.

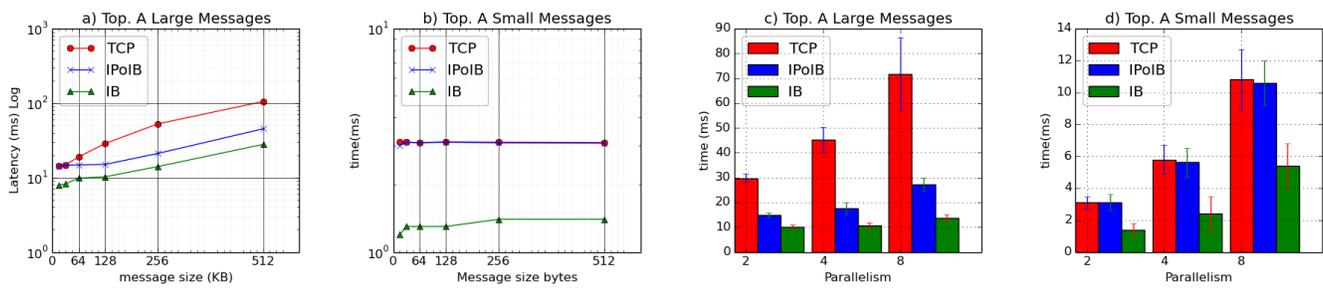


Figure 6: Latency of the Topology A with 1 spout and 7 bolt instances arranged in a chain with varying parallelism and message sizes. a) and b) are with 2 parallelism and c) and d) are with 128k and 128bytes messages. The results are on Haswell cluster with IB.

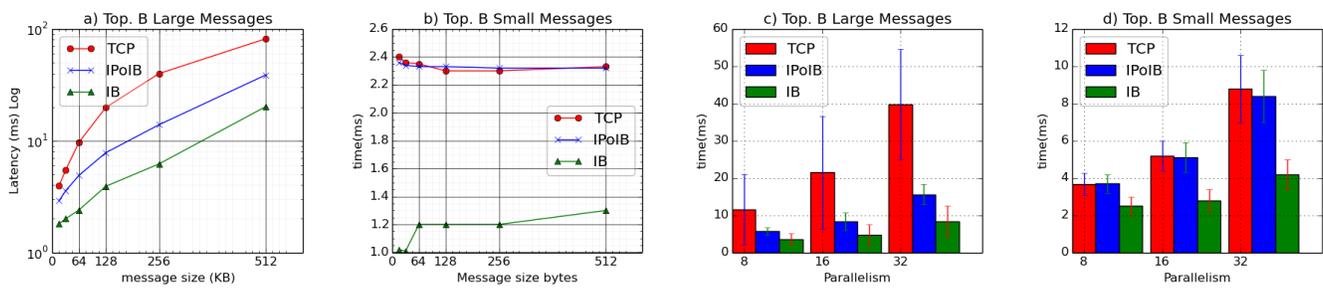


Figure 7: Latency of the Topology B with 32 parallel bolt instances and varying number of spouts and message sizes. a) and b) are with 16 spouts and c) and d) are with 128k and 128bytes messages. The results are on Haswell cluster with IB.

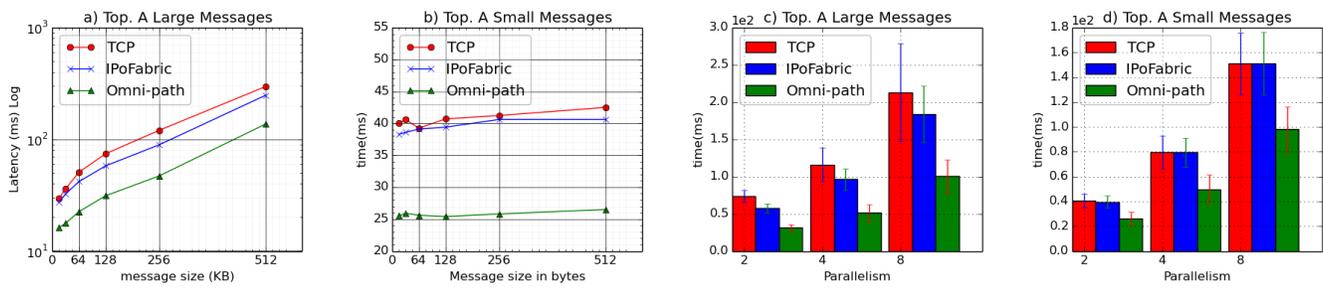


Figure 8: Latency of the Topology A with 1 spout and 7 bolt instances arranged in a chain with varying parallelism and message sizes. a) and b) are with 2 parallelism and c) and d) are with 128k and 128bytes messages. The results are on KNL cluster.

Also the authors have looked at various available network protocols for Big Data applications [28].

There are many distributed stream engines available today including Apache Spark Streaming [32], Apache Flink [7], Apache Apex [1] and Google Cloud Data flow [3]. All these systems follow the data flow model with comparable features to each other. Stream bench [21] is a benchmark developed to evaluate the performance of these systems in detail. These systems are primarily optimized for commodity hardware and clouds. There has been much research done around Apache Storm and Apache Heron to improve its capabilities. [10] described architectural and design improvements to Heron that improved its performance much further.

In recent years there have been multiple efforts to integrate high performance interconnects with machine learning and big data frameworks. Infiniband has been integrated into Spark [24] where the focus is on the Spark batch processing aspects rather than the streaming aspects. Spark is not considered a native stream processing engine and only implements streaming as an extension to its batch engine, making its latency inadequate for low latency applications. Recently Infiniband has been integrated into Hadoop [22], along with HDFS [13] as well. Hadoop, HDFS and Spark all use Java runtime for their implementations, hence the RDMA was integrated using JNI wrappers to C/C++ codes that invoke the underlying RDMA implementations. Recent work has

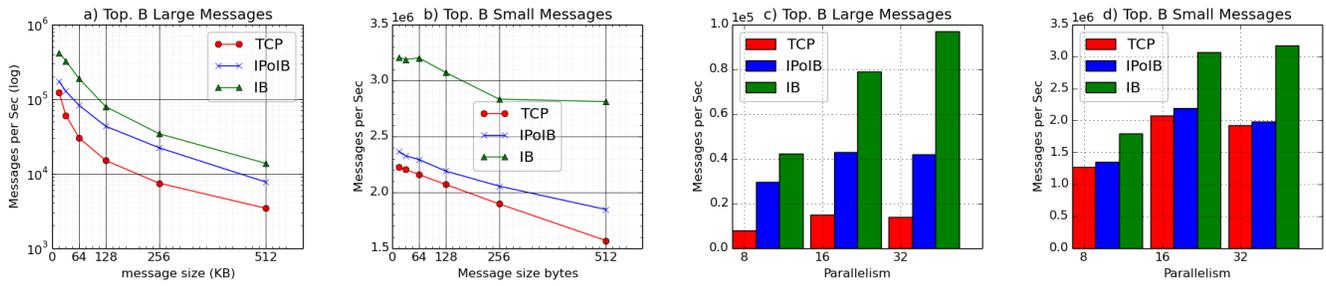


Figure 9: Throughput of the Topology B with 32 bolt instances and varying message sizes and spout instances. The message size varies from 16K to 512K bytes. The spouts are changed from 8 to 32.

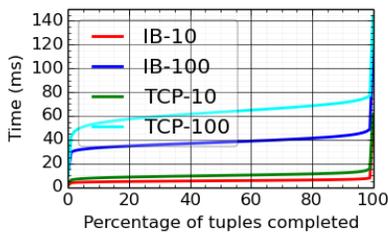


Figure 10: Percent of messages completed within a given latency for the Topology A with in-flight messages at 100 and 10 with 128K messages and 8 parallelism

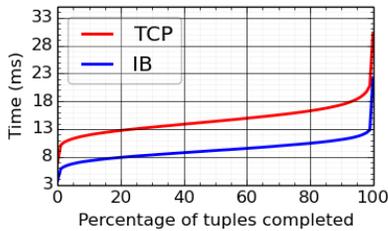


Figure 11: Percent of messages completed within a given latency for the Yahoo! streaming benchmark with Infiniband network. The experiment was conducted in Haswell cluster with 8 nodes and each stage of topology have 8 parallel tasks.

added RDMA support for the Tensorflow [2] machine learning framework.

High performance interconnects have been widely used by the HPC community, and most MPI(Message passing Interface) implementations have support for a large number of interconnects that are available today. Some early work that describes in detail about RDMA for MPI can be found in [20]. There has even been some work to build Hadoop-like systems using the existing MPI capabilities [23]. Photon [16] is a higher level RDMA library that can be used as a replacement to libfabric. RoCE [5] and iWARP [25] are protocols designed to use RDMA over Ethernet to increase the maximum packets per second processed at a node while decreasing

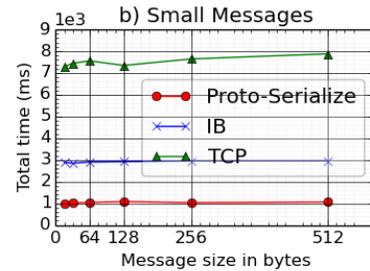
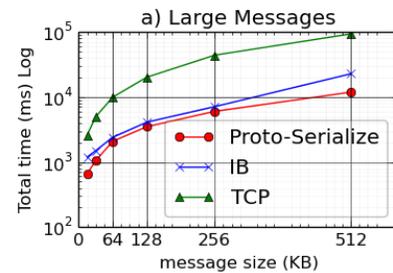


Figure 12: The total time to finish the messages vs the total time to serialize the messages using protobuf and Kryo. Top. B with 8 parallelism for bolts and spouts used. Times are for 20000 large messages and 200000 small messages. The experiment is conducted on Haswell cluster.

latency. Heron doesn't take into account the network when scheduling tasks into available nodes and it would be interesting to consider network latencies as described in [14, 17] for task scheduling.

8 CONCLUSIONS & FUTURE WORK

Unlike other Big Data systems which are purely JVM-based, Heron has a hybrid architecture where it uses both low level and high level languages appropriately. This architecture allows the addition of high performance enhancement such as Infiniband natively rather than going through additional layers of JNI as done in high performance Spark and Hadoop. With Infiniband and Omni-Path integration to Heron, we have seen good performance gains both in latency and throughput. With Infiniband integration we have

seen the communication overhead is moving closer to object serialization cost. The architecture and the implementations can be improved further to reduce the latency and increase the throughput of the system. Even though the authors use Apache Heron in this paper, it is possible to implement high performance interconnects to other stream processing engines such as Apache Flink [7], Apache Spark [31] and Apache Apex [1] using JNI wrappers.

Past research has shown that the remote memory access operations of Infiniband are more efficient than using channel semantics for transferring large messages. A hybrid approach can be adopted to transfer messages using both channel semantics and memory semantics. It is evident that the CPU is a bottleneck at the stream managers to achieve better performance. The protocol buffer processing along with message serialization are the dominant CPU consumers. A more efficient binary protocol that does not require protocol buffer processing at the stream manager can avoid these overheads. Instances to stream manager communication can be improved with a shared memory approach to avoid the TCP stack. With such approach the Infiniband can be improved to directly use the shared memory for the buffers without relying on data copying.

Because of the single-process single-threaded approach used by Heron processes, many core systems such as Knights Landing cannot get optimum performance out of Heron. Having a hybrid architecture where multiple threads are used for both communication and computation utilizing the hardware threads of the many core systems can increase the performance of Heron in such environments. Since we are using a fabric abstraction to program Infiniband and Omni-Path with Libfabric, the same code can be used with other potential high performance interconnects, though it has to be evaluated in such environments to identify possible changes. We would like to continue this work to include more high level streaming operations such as streaming joins, reductions on top of high performance interconnects.

ACKNOWLEDGMENT

This work was partially supported by NSF CIF21 DIBBS 1443054 and NSF RaPyDLI 1415459. We thank Intel for their support of the Haswell system, and extend our gratitude to the FutureSystems team for their support with the infrastructure. We would like to thank Heron team for their support of this work.

REFERENCES

- [1] 2017. Apache Apex: Enterprise-grade unified stream and batch processing engine. (2017). <https://apex.apache.org/>
- [2] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, and others. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* (2016).
- [3] Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael J Fernández-Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, and others. 2015. The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment* 8, 12 (2015), 1792–1803.
- [4] InfiniBand Trade Association and others. 2000. *InfiniBand Architecture Specification: Release 1.0*. InfiniBand Trade Association.
- [5] Motti Beck and Michael Kagan. 2011. Performance Evaluation of the RDMA over Ethernet (RoCE) Standard in Enterprise Data Centers Infrastructure. In *Proceedings of the 3rd Workshop on Data Center - Converged and Virtual Ethernet Switching (DC-CaVES '11)*. International Teletraffic Congress, 9–15. <http://dl.acm.org/citation.cfm?id=2043535.2043537>
- [6] Mark S Birrittella, Mark Debbage, Ram Huggahalli, James Kunz, Tom Lovett, Todd Rimmer, Keith D Underwood, and Robert C Zak. 2015. Intel® Omni-path Architecture: Enabling Scalable, High Performance Fabrics. In *High-Performance Interconnects (HOTI), 2015 IEEE 23rd Annual Symposium on*. IEEE, 1–9.
- [7] Paris Carbone, Stephan Ewen, Seif Haridi, Asterios Katsifodimos, Volker Markl, and Kostas Tzoumas. 2015. Apache flink: Stream and batch processing in a single engine. *Data Engineering* (2015), 28.
- [8] Josiah L Carlson. 2013. *Redis in Action*. Manning Publications Co.
- [9] Geoffrey Fox, Judy Qiu, Shantenu Jha, Saliya Ekanayake, and Supun Kamburugamuve. 2015. Big data, simulations and hpc convergence. In *Workshop on Big Data Benchmarks*. Springer, 3–17.
- [10] Maosong Fu, Ashvin Agrawal, Avriilia Floratou, Graham Bill, Andrew Jorgensen, Mark Li, Neng Lu, Karthik Ramasamy, Sriram Rao, and Cong Wang. 2017. Twitter Heron: Towards Extensible Streaming Engines. *2017 IEEE International Conference on Data Engineering* (Apr 2017). <http://icde2017.sdsc.edu/industry-track>
- [11] S. Gallenmiller, P. Emmerich, F. Wohlfart, D. Raumer, and G. Carle. 2015. Comparison of frameworks for high-performance packet IO. In *2015 ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*. 29–38. DOI: <http://dx.doi.org/10.1109/ANCS.2015.7110118>
- [12] P. Grun, S. Hefty, S. Sur, D. Goodell, R. D. Russell, H. Pritchard, and J. M. Squyres. 2015. A Brief Introduction to the OpenFabrics Interfaces - A New Network API for Maximizing High Performance Application Efficiency. In *2015 IEEE 23rd Annual Symposium on High-Performance Interconnects*. 34–39. DOI: <http://dx.doi.org/10.1109/HOTI.2015.19>
- [13] Nusrat S Islam, MW Rahman, Jithin Jose, Raghunath Rajachandrasekar, Hao Wang, Hari Subramoni, Chet Murthy, and Dhableswar K Panda. 2012. High performance RDMA-based design of HDFS over InfiniBand. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society Press, 35.
- [14] Virajith Jalaparti, Peter Bodik, Ishai Menache, Sriram Rao, Konstantin Makarychev, and Matthew Caesar. 2015. Network-Aware Scheduling for Data-Parallel Jobs: Plan When You Can. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM '15)*. ACM, New York, NY, USA, 407–420. DOI: <http://dx.doi.org/10.1145/2785956.2787488>
- [15] Supun Kamburugamuve, Saliya Ekanayake, Milinda Pathirage, and Geoffrey Fox. 2016. Towards High Performance Processing of Streaming Data in Large Data Centers. In *HPBDC 2016 IEEE International Workshop on High-Performance Big Data Computing in conjunction with The 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2016)*, Chicago, Illinois USA.
- [16] Ezra Kissel and Martin Swany. 2016. Photon: Remote memory access middleware for high-performance runtime systems. In *Parallel and Distributed Processing Symposium Workshops, 2016 IEEE International*. IEEE, 1736–1743.
- [17] Adrian Klein, Fuyuki Ishikawa, and Shinichi Honiden. 2012. Towards Network-aware Service Composition in the Cloud. In *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*. ACM, New York, NY, USA, 959–968. DOI: <http://dx.doi.org/10.1145/2187836.2187965>
- [18] Jay Kreps, Neha Narkhede, Jun Rao, and others. 2011. Kafka: A distributed messaging system for log processing. In *Proceedings of the NetDB*. 1–7.
- [19] Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg, Suresh Mittal, Jignesh M Patel, Karthik Ramasamy, and Siddarth Taneja. 2015. Twitter heron: Stream processing at scale. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM, 239–250.
- [20] Jiuxing Liu, Jiesheng Wu, Sushmitha P Kini, Pete Wyckoff, and Dhableswar K Panda. 2003. High performance RDMA-based MPI implementation over InfiniBand. In *Proceedings of the 17th annual international conference on Supercomputing*. ACM, 295–304.
- [21] Ruirui Lu, Gang Wu, Bin Xie, and Jingtong Hu. 2014. Stream bench: Towards benchmarking modern distributed stream computing frameworks. In *Utility and Cloud Computing (UCC), 2014 IEEE/ACM 7th International Conference on*. IEEE, 69–78.
- [22] Xiaoyi Lu, Nusrat S Islam, Md Wasi-Ur-Rahman, Jithin Jose, Hari Subramoni, Hao Wang, and Dhableswar K Panda. 2013. High-performance design of Hadoop RPC with RDMA over InfiniBand. In *2013 42nd International Conference on Parallel Processing*. IEEE, 641–650.
- [23] Xiaoyi Lu, Fan Liang, Bing Wang, Li Zha, and Zhiwei Xu. 2014. DataMPI: extending MPI to hadoop-like big data computing. In *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*. IEEE, 829–838.
- [24] Xiaoyi Lu, Md Wasi Ur Rahman, Nusrat Islam, Dipti Shankar, and Dhableswar K Panda. 2014. Accelerating spark with RDMA for big data processing: Early experiences. In *High-performance interconnects (HOTI), 2014 IEEE 22nd annual symposium on*. IEEE, 9–16.
- [25] M. J. Rashti and A. Afsahi. 2007. 10-Gigabit iWARP Ethernet: Comparative Performance Analysis with InfiniBand and Myrinet-10G. In *2007 IEEE International Parallel and Distributed Processing Symposium*. 1–8. DOI: <http://dx.doi.org/10.1109/IPDPS.2007.370480>
- [26] Luigi Rizzo. 2012. Netmap: a novel framework for fast packet I/O. In *21st USENIX Security Symposium (USENIX Security 12)*. 101–112.
- [27] A. Sodani. 2015. Knights landing (KNL): 2nd Generation Intel xAE; Xeon Phi processor. In *2015 IEEE Hot Chips 27 Symposium (HCS)*. 1–24. DOI: <http://dx.doi.org/>

org/10.1109/HOTCHIPS.2015.7477467

- [28] Brian Tierney, Ezra Kissel, Martin Swany, and Eric Pouyoul. 2012. Efficient data transfer protocols for big data. In *E-Science (e-Science), 2012 IEEE 8th International Conference on*. IEEE, 1–9.
- [29] Ankit Toshniwal, Siddarth Taneja, Amit Shukla, Karthik Ramasamy, Jignesh M Patel, Sanjeev Kulkarni, Jason Jackson, Krishna Gade, Maosong Fu, Jake Donham, and others. 2014. Storm@ twitter. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 147–156.
- [30] Kenton Varda. 2008. Protocol buffers: Googlefs data interchange format. *Google Open Source Blog*. Available at least as early as Jul (2008).
- [31] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association, 2–2.
- [32] Matei Zaharia, Tathagata Das, Haoyuan Li, Scott Shenker, and Ion Stoica. 2012. Discretized streams: an efficient and fault-tolerant model for stream processing on large clusters. In *Presented as part of the*.
- [33] Yibo Zhu, Haggai Eran, Daniel Firestone, Chuanxiong Guo, Marina Lipshteyn, Yehonatan Liron, Jitendra Padhye, Shachar Raindel, Mohamad Haj Yahia, and Ming Zhang. 2015. Congestion Control for Large-Scale RDMA Deployments. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM '15)*. ACM, New York, NY, USA, 523–536. DOI: <http://dx.doi.org/10.1145/2785956.2787484>