HONE: Joint Host-Network Traffic Management in Software-Defined Networks

Peng Sun, Minlan Yu, Michael J. Freedman,
Jennifer Rexford, David Walker

Abstract Applications running in modern datacenters interact with the underlying network in complex ways, making administrators run multiple traffic-management tasks to tune the system. However, today's traffic-management solutions are limited by an artificial division between the hosts and the network. While network devices only have knowledge of the network layer, the hosts can provide more visibility into how applications interact with the network. This paper proposes to extend the scope of traffic management to the end-host network stack. We present a Software-Defined Networking (SDN) platform for joint HOst-NEtwork (HONE) traffic management. HONE presents a uniform view of a diverse collection of measurement data, minimizes measurement overhead by performing lazy materialization of fine-grained statistics, and scales the analysis by processing data locally on the end hosts. HONE offers a simple and expressive programming framework for network and service administrators. We evaluate HONE by implementing several canonical traffic-management applications, measuring its efficiency with micro-benchmarks, and demonstrating its scalability with larger experiments on Amazon EC2.

1 Introduction

Modern datacenters run a wide variety of applications that generate large amount of traffic. These applications have a complex relationship with the underlying network, including varying traffic patterns [1], TCP incast [2], suboptimal TCP parameters [3], and elephant flows

Peng Sun, Michael J. Freedman, Jennifer Rexford, David Walker
Princeton University
35 Olden Street, Princeton, NJ 08540
E-mail: pengsun, mfreed, jrex, dpw@cs.princeton.edu

Minlan Yu
University of Southern California
941 Bloom Walk, SAL 228, Los Angeles, CA 90089
E-mail: minlanyu@usc.edu

that overload certain paths [4]. To optimize application performance, datacenter administrators perform various traffic-management tasks, such as performance monitoring, server load balancing [5], rate limiting [6], and traffic engineering [7, 8].

However, today's traffic-management solutions are constrained by an artificial division between the end hosts and the network. Most existing solutions only rely on network devices. Without access to the application and transport layers, network devices cannot easily associate traffic statistics with applications. For instance, it is difficult to infer the root causes of application performance problems or infer the backlog of application traffic in the send socket buffer, when just analyzing network-level statistics. Furthermore, traffic-management solutions are limited by the CPU and memory resources of network devices.

Compared to network devices, hosts have better visibility into application behavior, greater computational resources, and more flexibility to adopt new functionality. Increasingly, datacenter traffic management capitalizes on the opportunity to move functionality from network devices to end hosts [3, 9, 10, 11, 12, 13, 14, 15, 16]. But, in these works, the hosts are just used as software network devices (e.g., soft switches). The unique advantages of hosts in traffic management are not fully harnessed.

Instead, the scope of traffic management should expand further, into the host network stack, to measure and analyze fine-grained traffic information. The host network stack can provide detailed data about how applications interact with the underlying network, while remaining application-agnostic. By combining host and network data, traffic-management solutions can understand the applications better in order to improve application performance and efficiency of network resource usage.

Bringing the host network stack into traffic management, we face multiple challenges: First, today's end-hosts have numerous interfaces for performing network-related functions. Datacenter administrators use heterogeneous tools to collect TCP logs and kernel network statistics from hosts (e.g., Windows ETW [17], Web10G [18], and vCenter [19]), and they have multiple tools for controlling network behavior (e.g., Linux `tc` [20], `iptables` [21], and Open vSwitch [22]). To simplify traffic management, we need to provide a *uniform* interface for datacenter administrators to collect measurement data from hosts and switches.

Furthermore, datacenter administrators do not settle on a fixed set of data to measure in advance. Multiple traffic-management tasks need to run at the same time, and they need to measure different data from the hosts and switches. Additionally, new solutions will adopt new measurement data as the application mix and the network design evolve. However, the overhead of measuring all data blindly is prohibitive, especially when the detailed statistics of host stacks are included in traffic management. Thus, we need a *flexible* way to selectively collect measurements, tailored to the demands of the management tasks.

A final challenge is the efficiency and scalability of traffic management. Although hosts provide us with detailed statistics, the sheer volume of data poses a scalability challenge for the real-time analysis of the measured data. At the same time, the computational resources of hosts provide an opportunity to execute the analysis logic locally on the same hosts that collect the measurements. To utilize the hosts' resources, the analysis logic of management tasks should be *partitioned* into two parts: one that runs locally with monitoring on the hosts, and the other that runs on a controller with a global view of all the hosts and switches.

In this paper, we present a scalable and programmable platform for joint HOst-NEtwork (HONE) traffic management. As shown in Figure 1, the architecture includes a logically centralized controller, HONE agents running on each host, and a module interacting with network devices, following the recent trend of Software Defined Networking (SDN). HONE performs monitoring and analysis on streams of measurement data in a programmable fashion. A management program can easily define the measurement of various data on hosts and

Fig. 1: Overview of HONE System

network with minimum collection overhead. The data-analysis logic can be divided into local parts that execute together with measurement on hosts, and global parts running on the controller. To be more specific, two key technical contributions underlie HONE:

Lazy measurement of various data on hosts and network: HONE first abstracts the heterogeneous data on hosts and switches into a uniform representation of database-like tables. Using our query language on top of the uniform interface, administrators can easily specify what data to measure from various sources. Collecting many statistics for many connections at arbitrary time granularity would be too expensive to support directly. Instead, HONE enables lazy materialization of the measurement data, i.e., the controller and the host agents analyze the queries to collect only the necessary statistics for multiple management tasks at the appropriate frequencies.

Partitioning of analysis logic for local execution on hosts: We design a set of data-parallel streaming operators for programming the data-analysis logic. With these operators, programmers can easily link the measurement queries with the analysis logic in a streaming fashion. Our design of the operators further enables partitioning of the analysis logic between the host agents and the controller. To further boost scalability, we also design operators for programmers to hierarchically aggregate/process analysis results among multiple hosts.

HONE combines and extends techniques in stream programming and distributed systems in a unique way to offer programmable and scalable platform for joint host-network traffic management in SDN. To demonstrate the effectiveness of our design, we have built a prototype of HONE. Our prototype host agent collects socket activities by intercepting Linux system calls, and measures TCP statistics with Web10G [18, 23]. Our controller interacts with switches using OpenFlow [24]. The micro-benchmark experiments show that HONE can measure and calculate an application's throughput, and aggregate the results across 128 EC2 machines within a 90^{th}-percentile latency of 58ms.

To further demonstrate the power of HONE programming framework, we build a collection of canonical management applications, such as flow scheduling [4, 25], distributed rate limiting [6, 26], network diagnosis [3], etc. These examples demonstrate the expressiveness of our language, as well as the scalability of our data collection and analysis techniques.

HONE focuses on different contexts and problems than prior works in streaming database [27, 28] and MapReduce [29, 30, 31]. Compared to streaming databases which operate

Measurement	Analysis	Control

Server Load Balancing

Measure servers' utilization and request incoming rates	→	Compute total request rate and a target distribution of requests	→	Reconfigure load balancing policies to enforce the target

Elephant Flow Scheduling

Measure sockets' backlog traffic demand	→	Detect elephant flows. Compute routes for them	→	Install routing rules in network

Fig. 2: Three Stages of Traffic Management

on existing data, HONE focuses on how to lazily monitor a minimum necessary set of statistics on hosts, and dynamically generate streams of measurement data for later analysis. MapReduce's key-value programming model fits well with naturally parallelizable data with multiple keys, and it uses data shuffling by keys for reduce operations. In contrast, HONE focuses on the measurement data which are inherently associated with hosts that collected them. HONE partitions and places the analysis logic together with monitoring to consume data locally and reduce data-transmission overhead.

2 HONE Programming Framework

HONE's programming framework is designed to enable a wide range of traffic-management tasks. Traffic management is usually oriented around a three-stage "control loop" of measurement, data analysis, and control. Figure 2 presents two representative applications that serve as running examples throughout the paper:

Server load balancing: The first application distributes incoming requests across multiple server replicas. After measuring the request rate and the server load (e.g., CPU, memory, and bandwidth usage) at each host, the application estimates the total request rate, computes a new target division of requests over the hosts, and configures switches accordingly.

Elephant flow scheduling: The second application is inspired by how Hedera [4] and Mahout [25] schedule large flows. After measuring the backlog in the socket buffer for each TCP connection, the application identifies the elephant flows and directs them over paths that minimize network congestion.

2.1 Measurement: Query on Global Tables

HONE's data model unifies the representation of statistics across a range of formats, locations, types of devices, and modes of access. The HONE controller offers a simple abstraction of a central set of database tables. Programmers can launch sophisticated queries, and rely on HONE to distribute the monitoring to the devices, materialize the necessary tables, transform the data to fit the schema, perform local computations and data reduction, and aggregate the data. The data model reduces a complex and error-prone distributed programming task to a set of simple, tabular queries that can usually be crafted in just tens of lines of code.

The HONE data model is organized around the protocol layers and the available data sources. Table 1 shows the tables that our current prototype supports. On the hosts, HONE

Table Name	Row for each	Columns
Connections	Connection	App, TCP/UDP five-tuple, TCP-stack statistics.
Applications	Process	Host ID, PID, app's name, CPU/memory usage.
Machines	Host	Host ID, total CPU usage, total memory usage, IP.
Links	Link	IDs/ports of two ends, capacity.
SwitchStats	Switch interface	Switch ID, port, timestamp, per-port counters.

Table 1: Global Tables Supported in HONE Prototype

Query	:=	*Select*(Stats) * *From*(Table) * *Where*(Criteria) * *Groupby*(Stat) * *Every*(Interval)
Table	:=	Connections | Applications | Machines | Links | SwitchStats
Stats	:=	*Columns* of Table
Interval	:=	*Integer* in Seconds or Milliseconds
Criteria	:=	Stat Sign *value*
Sign	:=	> | < | ≥ | ≤ | = | ≠

Table 2: Measurement Query Language Syntax

collects socket logs and TCP connection statistics, in order to capture the relationship between applications and the network stack while remaining application-agnostic. On the switches, HONE collects the topology, the routing configurations, and per-port counters using OpenFlow. However, we can easily extend our prototype to support more interfaces (e.g., NetFlow, SNMP) by adding new tables, along with implementations for collecting the data.

HONE offers programmers a familiar, SQL-like query language for collecting the data, as summarized in Table 2 [1]. The query language gives programmers a way to state declaratively *what* data to measure, rather than *how*. More sophisticated analysis, transformation, filtering, and aggregation of the data take place in the analysis phase. To illustrate how to create a HONE program, consider the three example queries needed for elephant-flow scheduling:

Backlog in socket buffer: This query generates the data for computing the backlog in the socket buffers:

```
def ElephantQuery():
  return (
    Select([SrcIp, DstIp, SrcPort, DstPort, BytesWritten, BytesSent]) *
    From(Connections) *
    Every(Seconds 1) )
```

The query produces a stream of tables, with one table every second. In each table, each row corresponds to a single connection and contains the endpoint IP addresses and port numbers, as well as the amount of data written into the socket buffer and sent into the network. Later, the analysis phase can use the per-connection `BytesWritten` and `BytesSent` to compute the backlog in the socket buffer to detect elephant flows.

---

[1] In the syntax, the star operator (*) glues together the various query components. Each query term generates a bit of abstract syntax that our runtime system interprets.

Connection-level traffic by host pair: This query collects the data for computing the traffic matrix:

```
def TrafficMatrixQuery():
  return(
    Select([SrcIp, DstIp, BytesSent, Timestamp]) *
    From(Connections) *
    Groupby([SrcIp,DstIp]) *
    Every(Seconds 1) )
```

The query uses the `Groupby` operator to convert each table (at each second) into a *list* of tables, each containing information about all connections for a single pair of end-points. Later, the analysis phase can sum the `BytesSent` across all connections in each table in the list, and compute the difference from one time period to the next to produce the traffic matrix.

Active links and their capacities: This query generates a stream of tables with all unidirectional links in the network:

```
def LinkQuery():
  return(
    Select([BeginDevice, EndDevice, Capacity]) *
    From(Links) *
    Every(Seconds 1) )
```

Together, these queries provide the information needed for the elephant-flow application. They also illustrate the variety of different statistics that HONE can collect from both hosts and switches—all within a simple, unified programming framework. Under the hood, the HONE controller analyzes these queries to merge overlapping parts. Then the host agents or the network module will lazily collect only the queried statistics at appropriate frequencies to minimize the measurement overhead.

2.2 Analysis: Data-Parallel Operators

HONE enables programmers to analyze data across multiple hosts, without worrying about the low-level details of communicating with the hosts or tracking their failures. HONE's functional data-parallel operators allow programmers to say *what* analysis to perform, rather than *how*. Programmers can associate their own functions with the operators to apply these functions across sets of hosts, as if the streams of tabular measurement data were all available at the controller. Yet, HONE gives the programmers a way to express whether their functions can be (safely) applied in parallel across data from different hosts, to enable the runtime system to reduce the bandwidth and processing load on the controller by executing these functions at the hosts. HONE's data-parallel operators include the following:

– `MapSet(f)`: Apply function f to every element of a stream in the set of streams, producing a new set of streams.
– `FilterSet(f)`: Create a new set of streams that omits stream elements e for which $f(e)$ is false.
– `ReduceSet(f,i)`: "Fold" function f across each element for each stream in the set, using i as an initializer. In other words, generate a new set of streams where $f(\dots f(f(i, e_1), e_2) \dots, e_n)$ is the n^{th} element of each stream when e_1, e_2, ..., e_n were the first n elements of the original stream.
– `MergeHosts()`: Merge a set of streams on the hosts into one single global stream. (Currently in HONE, the collection of switches already generate a single global stream of

measurement data, given that our prototype integrates with an SDN controller to access data from switches.)

`MapSet`, `FilterSet`, and `ReduceSet` are parallel operators. Using them indicates that the analysis functions associated with the operators can run in a distributed fashion. It offers HONE knowledge of how to partition the analysis logic into the local part that can run in parallel on each host and the global part that sits on the controller. HONE also enables analysis on a single global stream with corresponding operators, such as `MapStream`, `FilterStream`, and `ReduceStream`. To combine queries and analysis into a single program, the programmer simply associates his functions with the operators, and "pipes" the result from one query or operation to the next (using the >> operator).

Consider again the elephant-flow scheduling application, which has three main parts in the analysis stage:

Identifying elephant flows: Following the approach suggested by Curtis *et al.* [25], the function `IsElephant` defines elephant flows as the connections with a socket backlog (i.e., the difference between bytes `bw` written by the application and the bytes `bs` acknowledged by the recipient) in excess of 100KB:

```
def IsElephant(row):
  [sip,dip,sp,dp,bw,bs] = row
  return (bw-bs > 100)

def DetectElephant(table):
  return (FilterList(IsElephant, table))

EStream = ElephantQuery() >>
        MapSet(DetectElephant) >>
        MergeHosts()
```

`DetectElephant` uses `FilterList` (the same as `filter` in Python) to apply `IsElephant` to select only the rows of the connection table that satisfy this condition. Finally, `DetectElephant` is applied to the outputs of `ElephantQuery`, and the results are merged across all hosts to produce a single stream `EStream` of elephant flows at the controller.

Computing the traffic matrix: The next analysis task computes the traffic matrix, starting from aggregating the per-connection traffic volumes by source-destination pair, and then computing the difference across consecutive time intervals:

```
TMStream = TrafficMatrixQuery() >>
        MapSet(MapList(SumBytesSent) >>
        ReduceSet(CalcThroughput, {}) >>
        MergeHosts() >>
        MapStream(AggTM)
```

The query produces a stream of lists of tables, where each table contains the per-connection traffic volumes for a single source-destination pair at a point in time. `MapList` (i.e., the built-in `map` in Python) allows us to apply a custom function `SumBytesSent` that aggregates the traffic volumes across connections in the same table, and `MapSet` applies this function over time. The result is a stream of tables, which each contains the cumulative traffic volumes for every source-destination pair at a point in time. Next, the `ReduceSet` applies a custom function `CalcThroughput` to compute the differences in the total bytes sent from one time to the next. The last two lines of the analysis merge the streams from different hosts and apply a custom function `AggTM` to create a global traffic matrix for each time period at the controller.

Policy	:= [Rule]+
Rule	:= *if* Criteria *then* Action
Criteria	:= Predicate [(and | or) Predicate]*
Predicate	:= Field = *value*
Field	:= AppName | SrcHost | DstHost | Headers
Headers	:= SrcIP | DstIP | SrcPort | DstPort | · · ·
Action	:= rate-limit *value* | forward-on-path *path*

Table 3: Control Policy in HONE Prototype

Constructing the topology: A last part of our analysis builds a network topology from the link tables produced by `LinkQuery`, which is abstracted as a single data stream collected from the network:

```
TopoStream = LinkQuery() >>
            MapStream(BuildTopo)
```

The auxiliary `BuildTopo` function (not shown) converts a single table of links into a graph data structure useful for computing paths between two hosts. The `MapStream` operator applies `BuildTopo` to the stream of link tables to generate a stream of graph data structures.

2.3 Control: Uniform and Dynamic Policy

In controlling hosts and switches, the datacenter administrators have to use various interfaces. For example, administrators use `tc`, `iptables`, or Open vSwitch on hosts to manage traffic, and they use SNMP or OpenFlow to manage the switches. For the purpose of managing traffic, these different control interfaces can be unified because they share the same pattern of generating control policies: for a group of connections satisfying criteria, define what actions to take. Therefore, HONE offers administrators a uniform way of specifying control policies as *criteria + action* clauses, and HONE takes care of choosing the right control implementations, e.g., we implement *rate-limit* using `tc` and `iptables` in the host agent.

The criteria can be network identifiers (e.g., IP addresses, port numbers, etc.). But this would force the programmer to map his higher-level policies into lower-level identifiers, and identify changes in which connections satisfy the higher-level policies. Instead, we allow programmers to identify connections of interest based on higher-level attributes, and HONE automatically tracks which traffic satisfies these attributes as connections come and go. Our predicates are more general than network-based rule-matching mechanisms in the sense that we can match connections by applications with the help of hosts. Table 3 shows the syntax of control policies, each of which our current prototype supports.

Continuing the elephant-flow application, we define a function `Schedule` that takes inputs of the detected elephant flows, the network topology, and the current traffic matrix. It assigns a routing path for each elephant flow with a greedy *Global First Fit* [4] strategy, and creates a HONE policy for forwarding the flow along the picked path. Other non-elephant flows are randomly assigned to an available path. The outputs of policies by `Schedule` will be piped into `RegisterPolicy` to register them with HONE.

```
def Schedule(elephant, topo, traffic):
  routes = FindRoutesForHostPair(topo)
  policies = []
```

```
for four_tuples in elephant:
  path = GreedilyFindAvailablePath(four_tuples, routes, traffic)
  criteria = four_tuples
  action = forward-on-path path
  policies.append([criteria, action])
return policies
```

2.4 All Three Stages Together

Combining the measurement, analysis, and control phases, the complete program merges the data streams, feeds the data to the `Schedule` function, and registers the output of policies. With this concrete example of an elephant-flow detection and scheduling application, we have demonstrated the simple and straightforward way of designing traffic-management tasks in HONE programming framework.

```
def ElephantFlowDetectionScheduling():
  MergeStreams([EStream, TopoStream, TMStream]) >>
  MapStream(Schedule) >>
  RegisterPolicy()
```

3 Efficient & Scalable Execution

Monitoring and controlling many connections for many applications on many hosts could easily overwhelm a centralized controller. HONE overcomes this scalability challenge in four main ways. First, a distributed directory service dynamically tracks the mapping of management tasks to hosts, applications, and connections. Second, the HONE agents lazily materialize virtual tables based on the current queries. Third, the controller automatically partitions each management task into global and local portions, and distributes the local part over the host agents. Fourth, the hosts automatically form a tree to aggregate measurement data based on user-defined aggregation functions to limit the bandwidth and computational overhead on the controller.

3.1 Distributed Directory Service

HONE determines which hosts should run each management task, based on which applications and connections match the queries and control policies. HONE has a directory service that tracks changes in the active hosts, applications, and connections. To ensure scalability, the directory has a two-tiered structure where the first tier (tracking the relatively stable set of active hosts and applications) runs on the controller, and the second tier (tracking the large and dynamic collection of connections) runs locally on each host. This allows the controller to decide which hosts to inform about a query or control policy, while relying on each local agent to determine which connections to monitor or control.

Tracking hosts and applications: Rather than build the first tier of the directory service as a special-purpose component, we leverage the HONE programming framework to run a standing query:

```
def DirectoryService():
  (Select([HostID, App]) *
   From(Applications) *
   Every(Seconds 1) ) >>
  ReduceSet(GetChangeOfAppAndHealth,[]) >>
  MergeHosts() >>
  MapStream(NotifyRuntime)
```

which returns the set of active hosts and their applications. `GetChangeOfAppAndHealth` identifies changes in the set of applications running on each host, and the results are aggregated at the controller. The controller uses its connectivity to each host agent as the host's health state, and the host agent uses `ps` to find active applications.

Tracking connections: To track the active connections, each host runs a Linux kernel module we build that intercepts the socket system calls (i.e., `connect`, `accept`, `send`, `receive`, and `close`). Using the kernel module, the HONE agent associates each application with the TCP/UDP connections it opens in an event-driven fashion. This avoids the inevitable delay of poll-based alternatives using Linux `lsof` and `/proc`.

3.2 Lazily Materialized Tables

HONE gives programmers the abstraction of access to diverse statistics at any time granularity. To minimize measurement overhead, HONE lazily materializes the statistics tables by measuring only certain statistics, for certain connections, at certain times, as needed to satisfy the queries. The HONE controller analyzes the queries from the management programs, and identifies what queries should run on hosts or switches. For queries to run on hosts, the host agents merge the collection of overlapping statistics to share among management programs. The agents collect only the statistics as specified in the queries with appropriate measurement techniques, instead of measuring all statistics in the virtual tables. The network module also merges the collection of shared statistics among queries, and collects the requested statistics from switches using OpenFlow.

Returning to the elephant-flow application, the controller analyzes the `ElephantQuery` and decides to run the query on the hosts. Since the query does not constrain the set of hosts and applications, the controller instructs *all* local agents to run the query. Each HONE agent periodically measures the values of `SrcIP`, `DstIP`, `SrcPort`, `DstPort`, and `BytesSent` from the network stack (via Web10G [18]), and collects the `BytesWritten` from the kernel module discussed earlier in §3.1. Similarly, HONE queries the switches for the `LinkQuery` data; in our prototype, we interact with network devices using the Open-Flow protocol. HONE does not collect or record any unnecessary data. Lazy materialization supports a simple and uniform data model while keeping measurement overhead low.

3.3 Host-Controller Partitioning

In addition to selectively collecting traffic statistics, the hosts can significantly reduce the resulting data volume by filtering or aggregating the data. For example, the hosts could identify connections with a small congestion window, sum throughputs over all connections, or find the top *k* flows by traffic volume.

However, parallelizing an arbitrary controller program would be difficult. Instead, HONE provides a `MergeHosts` operator that explicitly divides a task into its local and global parts. Analysis functions before `MergeHosts` run locally on each host, whereas functions after

Hosts' Execution Plans

Fig. 3: Partitioned Execution Plan of Elephant-Flow Scheduling Program

`MergeHosts` run on the controller. HONE hides the details of distributing the computation, communicating with end devices, and merging the results. Having an explicit `MergeHosts` operator obviates the need for complex code analysis for automatic parallelization.

HONE coordinates the parallel execution of tasks across a large group of hosts.[2] We first carry out industry-standard clock synchronization with NTP [32] on all hosts and the controller. Then the HONE runtime stamps each management task with its creation time t_c. The host agent dynamically adjusts when to start executing the task to time $t_c + nT + \epsilon$, where n is an integer, ϵ is set to 10ms, and T is the period of task (as specified by the `Every` statement). Furthermore, the host agent labels the local execution results with a logical sequence number (i.e., n), in order to tolerate the clock drift among hosts. The controller buffers and merges the data bearing the same sequence number into a single collection, releasing data to the global portion of task when either receiving from all expected hosts or timing out after T.

Using our elephant-flow-scheduling application, Figure 3 shows the partitioned execution plan of the management program. Recall that we merge `EStream`, `TMStream`, and `TopoStream` to construct the program. The measurement queries are interpreted as parallel *Measure* operations on the host agents, and the query of switch statistics from the network module. HONE executes the `EStream` and `TMStream` tasks on each host in parallel (to detect elephant flows and calculate throughputs, respectively), and streams these local results to the controller (i.e., *ToController*). The merged local results of `TMStream` pass through a throughput aggregation function (*AggTM*), and finally merge together with the flow-detection data and the topology data from `TopoStream` to feed the *Schedule* function.

3.4 Hierarchical Data Aggregation

Rather than transmit (filtered and aggregated) data directly to the controller, the hosts construct a hierarchy to combine the results using user-specified functions. HONE automatically

---

[2] The HONE controller ships the source code of the local portion of management tasks to the host agent. Since HONE programs are written in Python, the agent can execute them with its local Python interpreter, and thus avoids the difficulties of making the programs compatible with diverse environments on the hosts.

Fig. 4: Aggregation Tree: 8 Hosts with Branching of 2

constructs a *k*-ary tree rooted at the controller[3] and applies a `TreeMerge` operator at each level. All hosts running the task are leaves of the tree. For each group of *b* hosts, HONE chooses one to act as their parent in the tree. These parents are grouped again to recursively build the tree towards the controller. User-defined functions associated with `TreeMerge` are applied to all non-leaf nodes of the tree to aggregate data from their children. HONE is unique among research efforts on tree-based aggregation [33, 34], since prior works focus on aggregating data with *a priori* knowledge of the data structure, and don't allow users to specify their own aggregation functions.

Many aggregation functions used in traffic management are both commutative and associative; such functions can be applied hierarchically without compromising correctness. For example, determining the top *k* values for heavy-hitter analysis is amenable to either direct processing across all data or to breaking the data into subsets for intermediate analysis and combining the results downstream. Calculating the total throughput of connections across all hosts can also be calculated in such a distributed manner, as the arithmetic sum is also a commutative and associative function.

Making the user-defined aggregation functions be associative and commutative ensures that HONE can apply them correctly in a hierarchical manner. Using `TreeMerge`, HONE assumes that the associated functions have the required properties, avoiding the semantics analysis. `TreeMerge` is similar to `MergeHosts` in the sense that they both combine local data streams from multiple hosts into one data stream on the controller, and intermediate hosts similarly buffer data until they receive data from all their children or a timeout occurs. But with `TreeMerge`, HONE also applies a user-defined aggregation function, while `MergeHosts` simply merges all hosts' data at the controller without reduction.

The algorithm of constructing the aggregation tree is an interesting extensible part of HONE. We can group hosts based on their network locality, or we can dynamically monitor the resource usage on hosts to pick the one with most available resource to act as the intermediate aggregator. In our prototype, we leave those interesting algorithms to future works, but offer a basic one of incrementally building the tree by when hosts join the HONE system. Subject to the branching factor *b*, the newly joined leaf greedily finds a node in one level up with less than *b* children, and links with the node if found. If not found, the leaf promotes itself to one level up, and repeats the search. When the new node reaches the highest level and still cannot find a place, the controller node moves up one level, which increases the height of the aggregation tree. Figure 4 illustrates an aggregation tree under the basic algorithm when 8 hosts have joined and *b* is 2.

---

[3] The runtime uses information from the directory service to discover and organize hosts.

4 Performance Evaluation

In this section, we present micro-benchmarks on our HONE prototype to evaluate measurement overhead, the execution latency of management programs, and the scalability; §5 will demonstrate the expressiveness and ease-of-use of HONE using several canonical traffic-management applications.

We implement the HONE prototype in combination of Python and C. The HONE controller provides the programming framework and runtime system, which partitions the management programs, instructs the host agents for local execution, forms the aggregation hierarchy, and merges the data from hosts for the global portion of program execution. The host agent schedules the installed management tasks to run periodically, executes the local part of the program, and streams the serialized data to the controller or intermediate aggregators. We implement the network part of the prototype as a custom module in Floodlight [35] to query switch statistics and install routing rules.

Our evaluation of the prototype focuses on the following questions about our design decisions in §2 and §3.

1. How efficient is the host-based measurement in HONE?
2. How efficiently does HONE execute entire management tasks?
3. How much overhead does lazy materialization save?
4. How effectively does the controller merge data from multiple hosts using hierarchical aggregation?

We run the HONE prototype and carry out the experiments on Amazon EC2. All instances have 30GB memory and 8 virtual cores of 3.25 Compute Units each.[4]

4.1 Performance of Host-Based Measurement

The HONE host agent collects TCP connection statistics using the Web10G [18] kernel module. We evaluate the measurement overhead in terms of time, CPU, and memory usage as we vary the number of connections running on the host. To isolate the measurement overhead, we run a simple management task that queries a few randomly-chosen statistics of all connections running on the host every one second (we choose the four tuples, bytes of sent data, and the congestion window size). Our experiment consists of three EC2 instances—one for the controller, and two running the HONE agent.

To collect the statistics, the host agent must first identify what connections to measure. Then the agent queries the kernel via Web10G to retrieve the statistics. Finally, the agent organizes the statistics in the schema specified by the query and feeds the result to the management program. In Figure 5, we break down the latency in each portion. For each fixed number of connections, we run the management task for five minutes (i.e., about 300 iterations), and plot the average and standard deviation of time spent in each portion.

Figure 5 shows that the agent performs well, measuring 5k connections in an average of 532.6ms. The Web10G measurement takes the biggest portion–432.1ms, and the latency is linear in the number of active connections. The time spent in identifying connections to measure is relatively flat, since the agent tracks the relevant connections in an event-driven fashion via the kernel module of intercepting socket calls. The time spent in organizing the

---

[4] One EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or Xeon processor.

Fig. 5: Overhead of Collecting Connection Statistics

statistics rises slowly as the agent must go through more connections to format the results into the query's schema. The results set lower limit for the periods of management tasks that need measurement of different numbers of connections. The CPU and memory usage of the agent remain stable throughout the experiments, requiring an average of 4.55% CPU of one core and 1.08% memory of the EC2 instance.

4.2 Performance of Management Tasks

Next, we evaluate the end-to-end performance of several management tasks. To be more specific, we evaluate the latency of finishing one *round* of a task: from the agent scheduling a task to run, measuring the corresponding statistics, finishing the local analysis, sending the results to the controller, the controller receiving the data, till the controller finishing the remaining parts of the management program. We run three different kinds of management tasks which have a mix of leverages of hosts, switches, and the controller in HONE, in order to show the flexibility of HONE adapting to different traffic-management tasks. All experiments in this subsection run on a 8-host-10-switch fat-tree topology [36]. The switches are emulated by running Open vSwitch on an EC2 instance.

- **Task1** calculates the throughputs of all `iperf` connections on each host, sums them up, and aggregates the global `iperf` throughput at the controller. This task performs most of the analysis at the host agents, leaving relatively little work for the controller. Each host launches 100 `iperf` connections to another randomly chosen host.
- **Task2** queries the topology and switch statistics from the network, and uses the per-port counters on the switches to calculate the current link utilization. This task uses the network module in HONE a lot to measure data, and runs computation work on the controller. *Task2* is performed under the same setting of running `iperf` as *Task1*.
- **Task3** collects measurement data from the hosts to detect connections with a small congestion window (i.e., which perform badly). It also queries the network to determine

Fig. 6: Latency of One Round of Execution

the forwarding path for each host pair. The task then diagnoses the shared links among those problematic flows as possible causes of the bad network performance. *Task3* is a joint host-network job, which runs its computation across hosts, network, and the controller. *Task3* is still under the same setting, but we manually add rules on two links to drop 50% of packets for all flows traversing the links, emulating a lossy network.

Figure 6 illustrates the cumulative distribution function (CDF) of the latency for finishing one round of execution, as we run 300 iterations for each task. We further break down the latency into three parts: the execution time on the agent or the network, the data-transmission time from the host agent or network module to the controller, and the execution time on the controller. In Figure 7, we plot the average latency and standard deviation for each part of the three tasks. *Task1* finishes one round with a 90th-percentile latency of 27.8ms, in which the agent takes an average of 17.8ms for measurement and throughput calculation, the data transmission from 8 hosts to the controller takes another 7.7ms, and the controller takes the rest. Having a different pattern with *Task1*, *Task2*'s 140.0ms 90th-percentile latency is consisted of 87.5ms of querying the switches via Floodlight and 8.9ms of computation on the controller (the transmission time is near zero since Floodlight is running on the controller machine). *Task3*'s latency increases as it combines the data from both hosts and the network, and its CDF also has two stairs due to different responsiveness of the host agents and the network module.

Table 4 summarizes the average CPU and memory usage on the host agent and the controller when running the task. The CPU percentage is for one core of 8 cores of our testbed machines. The results show that HONE's resource usage are bind to the running management tasks: *Tasks3* is the most complex one with flow detection/rate calculation on the hosts, and having the controller join host and network data.

Fig. 7: Breakdown of Execution Latency

	CPU Agent	Memory Agent	CPU Controller	Memory Controller
Task1	3.71%	0.94%	0.67%	0.10%
Task2	N/A	N/A	0.76%	1.13%
Task3	7.84%	1.64%	1.03%	0.11%

Table 4: Average CPU and Memory Usage of Execution

4.3 Effects of Lazy Materialization

HONE lazily materializes the contents of the statistics tables. We evaluate how much overhead the feature can save for measurement efficiency in HONE.

We set up two applications (*A* and *B*) with 1k active connections each on a host. We run multiple management tasks with different queries over the statistics to evaluate the measurement overhead in terms of latency. Figure 8 illustrates the average and standard deviation of the latencies for different queries. The first program queries all 122 TCP-stack statistics available in Web10G of all 2k connections, and all applications' CPU and memory usage. The following ones query various statistics of `Connections` or `Applications` tables with details shown on Figure 8.

The lazy materialization of the tables lowers the measurement overhead by either measuring a subset of tables (*Query1* vs. others), rows (number of connections in *Query1* vs. *Query2* and *Query3*), and columns (number of statistics in *Query2* vs. *Query3*). The high overhead of *Query4* is due to the implementation of CPU measurement, which is, for each process, one of the ten worker threads on the agent keeps running for 50ms to get a valid CPU usage.

Fig. 8: Effects of Lazy Materialization

4.4 Evaluation of Scalability in HONE

We will evaluate the scalability of HONE from two perspectives. First, when HONE controller partitions the management program into local and global parts of execution, the controller will handle the details of merging the local results processed in the same time period from multiple hosts, before releasing the merged result to the global part of execution. Although the host clocks are synchronized via NTP as mentioned in §3.3, the clocks still drift slightly over time. It results in a buffering delay at the controller. Now we will evaluate how well the buffering works in terms of the time difference between when the controller receives the first piece of data and when the controller receives all the data bearing the same sequence number.

To focus on the merging performance, we use the *Task1* in §4.2. All hosts will directly send their local results to the controller without any hierarchical aggregation. Each run of the experiment lasts 7 minutes, containing about 400 iterations. We repeat the experiment, varying the number of hosts from 16 to 128.

Figure 9 shows the CDFs of the latencies for these experiments. The 90th-percentile of the controller's buffering delay is 4.3ms, 14.2ms, 9.9ms, and 10.7ms for 16, 32, 64, and 128 hosts respectively. The results show that the synchronization mechanism on host agents work well in coordinating their local execution of a management task, and the controller's buffering delay is not a problem in supporting traffic-management tasks whose execution periods are typically in seconds.

After evaluating how the controller merges distributed collection of data, we would evaluate another important feature of HONE for scalability, which is the hierarchical aggregation among the hosts. We continue using the same management task of summing the application's throughputs across hosts. But we change to use the `TreeMerge` operator to apply the aggregation function. In this way, the task will be executed by HONE through a k-ary tree consisted of the hosts.

In this experiment, we fix the branching factor k of the hierarchy to 4. We repeat the experiment with 16, 32, 64, and 128 hosts, in which case the height of the aggregation tree is 2, 3, 3, and 4 respectively. Figure 10 shows the CDFs of the latencies of one round of

Fig. 9: Controller's Buffering Delay of Merging Data from Multiple Host Agents

Num of Hosts	CPU Agent	Memory Agent	CPU Controller	Memory Controller
16	4.19%	0.96%	1.09%	0.05%
32	4.93%	0.96%	1.27%	0.05%
64	5.26%	0.97%	1.31%	0.06%
128	4.80%	0.97%	2.36%	0.07%

Table 5: Average CPU and Memory Usage in Hierarchical Aggregation Experiments

execution, which captures the time difference from the earliest agent starting its local part to the controller finishing the global part. The 90th-percentile execution latency increases from 32.2ms, 30.5ms, 37.1ms, to 58.1ms. Table 5 shows the average CPU and memory usage on the controller and the host agent. The host agent's CPU and memory usage come from the agent that multiplexes as local-data generator and the intermediate aggregators in all levels of the k-ary tree. It shows the maximum overhead that the host agent incurs when running in a hierarchy.

From the results above, we can conclude that HONE's own operations pose little overhead to the execution of management tasks. The performance of management tasks running in HONE will be mainly bound by their own program complexities, and the amount of data they need to process or transmit.

5 Case Studies

We have shown the micro-benchmark evaluation of HONE to demonstrate its efficiency and scalability. Now we will illustrate the expressiveness and ease-of-use of HONE by building a diversity of traffic-management tasks in data centers. Table 6 lists all the management tasks that we have built, ranging from conventional management operations in data centers (e.g.,

Fig. 10: End-to-end Latency of Execution under Hierarchical Aggregation

Management Task	Lines of Code
Summing application's throughputs	70
Monitoring CPU and memory usage	24
Collecting connection TCP statistics	19
Calculating traffic matrix	85
Calculating link utilizations	48
Discovering network topology	51
Network performance diagnosis	56
HONE's directory service	31
Elephant flow scheduling	140
Distributed rate limiting	74

Table 6: HONE-based Traffic-Management Tasks. Check http://hone.cs.princeton.edu/

calculating link utilizations) to recent proposals (e.g., network performance diagnosis [3]). Those conventional traffic-management tasks can actually serve as basic building blocks for more complex management tasks. The administrators can compose the code of those HONE programs to construct their own. HONE is an open-source project, and code for the management programs are also available at `http://hone.cs.princeton.edu/examples`.

In the following subsections, we pick two management tasks as case studies to illustrate more details, and evaluate the HONE-based solutions.

5.1 Elephant Flow Scheduling

In data centers, it is important to detect elephant flows with high traffic demands and properly route them to minimize network congestion. With HONE, we can easily build such an application to schedule the elephant flows. This example takes the scheduling strategies and the elephant-flow detection threshold from Hedera [4] and Mahout [25]. We implement Hedera's *Global-first-fit* routing strategy with 140 lines of code in HONE. The code of the management task has been already shown in previous sections as an example.

We deploy HONE on EC2 instances to emulate a data-center network with a 8-host-10-switch fat-tree topology (the switches are instances running Open vSwitch). We repeat an all-to-all data shuffle of 500MB (i.e., a 28GB shuffle) for 10 times. The HONE-based solution finishes the data shuffle with an average of 82.7s, compared to 103.1s of using ECMP. The improvement over shuffle time is consistent with Hedera's result.

5.2 Distributed Rate Limiting

Distributed rate limiting in data centers is used to control the aggregate network bandwidth used by an application, which runs on multiple hosts. It can help the application's owner to control the total cost of using a pay-per-use cloud provider.

Prior works [6, 26] proposed mechanisms to make distributed rate-limiters collaborate as a single, aggregate global limiter inside the network. HONE enables distributed rate limiting from the host side, which would introduce less overhead as the hosts have more computational power than the switches, and better visibility into the traffic demand of applications.

In HONE, the administrators do not need to worry about the complexity of collecting throughputs from multiple hosts in a synchronized way. Instead, they just need to write a simple program that sums up throughputs of an application's connections on each host, aggregates the throughputs across hosts, and then calculates their rate-limiting policies accordingly. The code written in HONE are shown below:

```
def DistributedRateLimiting():
  (Select([App, SrcIp, DstIp,
           BytesSent, Timestamp]) *
   From(Connections) *
   Where(App == X) *
   Every(Seconds 1) ) >>
  ReduceSet(CalculateThroughput, {}) >>
  MapSet(LocalAgg) >>
  ReduceSet(MovingAverage, initValue) >>
  MergeHosts() >>
  MapStream(GenerateRateLimitPolicy) >>
  RegisterPolicy()
```

We run the task to limit the aggregate throughput of application *X* to 100Mbps. The application *X* is set to send traffic on each host for 80 seconds with a default rate of 50Mbps. We launch *X* on 5 hosts, one by one every 10 seconds. Figure 11 shows the time series of the aggregate and individual traffic rates of *X*. The management task succeeds in limiting the total rate of *X* running on a distributed set of hosts to 100Mbps. Note that it takes one round of execution for the management task to discover new traffic and update the rate-limiting policies. That is why there are several 1-second spikes when *X* starts on new hosts.

Fig. 11: Time Series of Application Throughput. One Host Starts the App Every 10s.

6 Related Works

Recent projects have sought to incorporate end hosts into network management [9, 10]. But these solutions view the hosts only as software switches or as trusted execution environments for the network. Cooke *et al.* [37] use hosts to collect socket activities for a better understanding of network. However, their solution cannot selectively collect statistics, and the statistics at the transport layer are not covered. OpenTCP [38] dynamically adapts the configuration of TCP on the end hosts based on traffic conditions. Lee *et al.* [39] propose joint optimization across the application and network layers for better application throughput. In contrast, HONE supports diverse traffic-management tasks, and thus provides more measurement, data analysis, and control functionalities across hosts and switches.

There have also been industry efforts in simplifying cloud management, such as various commercial tools from vendors [19, 40, 41]. They aim at enabling better visualization and infrastructure control at hosts and switches in the cloud. HONE is complementary to these systems by focusing more on monitoring and analysis in traffic-management tasks and providing programmable interfaces for these tasks.

Prior works also adopt the stream abstraction for network traffic analysis [27, 28]. But they mainly focus on extending the SQL language, while we use functional language constructs to define traffic-management mechanisms more easily. Further, some of these works [27, 42] focus on a specific problem (e.g., intrusion detection) when designing their programming language, while HONE aims for a more generic programming interface for traffic management.

Finally, there are recent works proposing network programming languages [11, 43, 44, 45, 46, 47]. Their programming abstraction is the raw packets or traffic counters on a single switch. HONE mainly moves the programmability to the end hosts, provides an extensible platform for various types of measurement, and spans both the hosts and the switches.

7 Conclusion

HONE is a programmable and scalable platform for joint host-network traffic management. HONE expands the scope of traffic management to the host network stacks, in order to harness the detailed network-related statistics and the computational resources on the hosts. We design an integrated data model for diverse fine-grained measurement from hosts and network. The programming framework further offers data-parallel streaming operators to define measurement and analysis logic of traffic-management tasks. The system can selectively measure the data as needed by the management programs, and it divides the analysis logic to execute locally on hosts in real time. Thus the HONE system is efficient and scalable for traffic management with integration of host stacks. Micro-benchmarks and experiments with real management tasks demonstrate the performance and expressiveness of our system.

In our future work, we plan to build a wider range of management tasks, both to further demonstrate HONE's expressiveness and to continue optimizing our prototype's performance. We will specifically consider the tasks for datacenter infrastructure management, such as failure mitigation and switch power management. In addition, we plan to make HONE suitable for a multi-tenant cloud environment, where tenants may not want the cloud provider to collect measurement data within the virtual machine's guest operating system. Instead, we plan to collect measurement data from the hypervisor to infer the application's behaviors. Besides extending the functionalities, we also plan to further improve the performance of HONE, and seek deployment into real datacenter environment for larger-scale experiments. We believe that these future works, along with our existing support for programmable and scalable traffic management, can make HONE an invaluable platform for the research community and datacenter administrators.

Acknowledgment

References

1. Srikanth Kandula, Sudipta Sengupta, Albert Greenberg, Parveen Patel, and Ronnie Chaiken. The Nature of Datacenter Traffic: Measurements & Analysis. In *ACM IMC*, 2009.
2. Haitao Wu, Zhenqian Feng, Chuanxiong Guo, and Yongguang Zhang. ICTCP: Incast Congestion Control for TCP in Data Center Networks. In *ACM CoNEXT*, 2010.
3. Minlan Yu, Albert Greenberg, Dave Maltz, Jennifer Rexford, Lihua Yuan, Srikanth Kandula, and Changhoon Kim. Profiling Network Performance for Multi-tier Data Center Applications. In *USENIX NSDI*, 2011.
4. Mohammad Al-Fares, Sivasankar Radhakrishnan, Barath Raghavan, Nelson Huang, and Amin Vahdat. Hedera: Dynamic Flow Scheduling for Data Center Networks. In *USENIX NSDI*, 2010.
5. 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. In *ACM SIGCOMM*, 2013.

6. Barath Raghavan, Kashi Vishwanath, Sriram Ramabhadran, Kenneth Yocum, and Alex C. Snoeren. Cloud Control with Distributed Rate Limiting. In *ACM SIGCOMM*, 2007.

7. Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, Jon Zolla, Urs Hölzle, Stephen Stuart, and Amin Vahdat. B4: Experience with a Globally-Deployed Software Defined WAN. In *ACM SIGCOMM*, 2013.

8. Chi-Yao Hong, Srikanth Kandula, Ratul Mahajan, Ming Zhang, Vijay Gill, Mohan Nanduri, and Roger Wattenhofer. Achieving High Utilization with Software-Driven WAN. In *ACM SIGCOMM*, 2013.

9. Colin Dixon, Hardeep Uppal, Vjekoslav Brajkovic, Dane Brandon, Thomas Anderson, and Arvind Krishnamurthy. ETTM: A Scalable Fault Tolerant Network Manager. In *USENIX NSDI*, 2011.

10. Thomas Karagiannis, Richard Mortier, and Antony Rowstron. Network Exception Handlers: Host-network Control in Enterprise Networks. In *ACM SIGCOMM*, 2008.

11. Justine Sherry, Daniel C. Kim, Seshadri S. Mahalingam, Amy Tang, Steve Wang, and Sylvia Ratnasamy. Netcalls: End Host Function Calls to Network Traffic Processing Services. Technical Report UCB/EECS-2012-175, U.C. Berkeley, 2012.

12. Alan Shieh, Srikanth Kandula, Albert Greenberg, Changhoon Kim, and Bikas Saha. Sharing the Data Center Network. In *USENIX NSDI*, 2011.

13. David Zats, Tathagata Das, Prashanth Mohan, Dhruba Borthakur, and Randy Katz. DeTail: Reducing the Flow Completion Time Tail in Datacenter Networks. In *ACM SIGCOMM*, 2012.

14. Ben Pfaff, Justin Pettit, Keith Amidon, Martin Casado, Teemu Koponen, and Scott Shenker. Extending networking into the virtualization layer. In *ACM HotNets*, October 2009.

15. Chi-Yao Hong, Matthew Caesar, and P. Brighten Godfrey. Finishing Flows Quickly with Preemptive Scheduling. In *ACM SIGCOMM*, 2012.

16. Christo Wilson, Hitesh Ballani, Thomas Karagiannis, and Ant Rowtron. Better Never Than Late: Meeting Deadlines in Datacenter Networks. In *ACM SIGCOMM*, 2011.

17. Event Tracing for Windows. `http://support.microsoft.com/kb/2593157`, 2011.

18. Web10G Project. `http://web10g.org/`, 2012.

19. VMWare vCenter Suite. `http://www.vmware.com/products/datacenter-virtualization/vcenter-operations-management/overview.html`, 2013.

20. Linux Advanced Routing & Traffic Control. `http://www.lartc.org/`, 2000.

21. Netfilter.org. `http://www.netfilter.org/`, 1999.

22. Open vSwitch. `http://openvswitch.org/`, 2011.

23. Matt Mathis, John Heffner, and Raghu Raghunarayan. TCP Extended Statistics MIB. RFC 4898, May 2007.

24. Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: Enabling Innovation in Campus Networks. *ACM CCR*, 38(2):69–74, 2008.

25. Andrew Curtis, Wonho Kim, and Praveen Yalagandula. Mahout: Low-Overhead Datacenter Traffic Management using End-Host-Based Elephant Detection. In *IEEE INFOCOM*,

2011.

26. Lucian Popa, Gautam Kumar, Mosharaf Chowdhury, Arvind Krishnamurthy, Sylvia Ratnasamy, and Ion Stoica. FairCloud: Sharing the Network in Cloud Computing. In *ACM SIGCOMM*, 2012.

27. Kevin Borders, Jonathan Springer, and Matthew Burnside. Chimera: A Declarative Language for Streaming Network Traffic Analysis. In *USENIX Security*, 2012.

28. Chuck Cranor, Theodore Johnson, Oliver Spataschek, and Vladislav Shkapenyuk. Gigascope: A Stream Database for Network Applications. In *ACM SIGMOD*, 2003.

29. Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *USENIX OSDI*, 2004.

30. Jeffrey Dean and Sanjay Ghemawat. MapReduce: A Flexible Data Processing Tool. *Commun. ACM*, 53(1):72–77, January 2010. ISSN 0001-0782. doi: 10.1145/1629175. 1629198. URL http://doi.acm.org/10.1145/1629175.1629198.

31. Andrew Pavlo, Erik Paulson, Alexander Rasin, Daniel J. Abadi, David J. DeWitt, Samuel Madden, and Michael Stonebraker. A Comparison of Approaches to Large-scale Data Analysis. In *ACM SIGMOD*, 2009.

32. NTP: The Network Time Protocol. http://www.ntp.org/, 2003.

33. Robbert van Renesse and Adrian Bozdog. Willow: DHT, Aggregation, and Publish/Subscribe in One Protocol. In *IPTPS*, 2004.

34. Praveen Yalagandula and Mike Dahlin. A Scalable Distributed Information Management System. In *ACM SIGCOMM*, 2004.

35. FloodLight. http://floodlight.openflowhub.org/, 2011.

36. Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A Scalable, Commodity Data Center Network Architecture. In *ACM SIGCOMM*, 2008.

37. Evan Cooke, Richard Mortier, Austin Donnelly, Paul Barham, and Rebecca Isaacs. Reclaiming network-wide visibility using ubiquitous end system monitors. In *USENIX ATC*, 2006.

38. Monia Ghobadi, Soheil Hassas Yeganeh, and Yashar Ganjali. Rethinking End-to-End Congestion Control in Software-Defined Networks. In *ACM HotNets*, 2012.

39. Young Lee, Greg Bernstein, Ning So, Tae Yeon Kim, Kohei Shiomoto, and Oscar Gonzalez de Dios. Research Proposal for Cross Stratum Optimization (CSO) between Data Centers and Networks. http://tools.ietf.org/html/draft-lee-cross-stratum-optimization-datacenter-00, March 2011.

40. OpenTSDB Project. http://www.opentsdb.net/, 2010.

41. Boundary. http://www.boundary.com/, 2010.

42. Xinming Ou, Sudhakar Govindavajhala, and Andrew W. Appel. MulVAL: A Logic-based Network Security Analyzer. In *USENIX Security*, 2005.

43. Nate Foster, Rob Harrison, Michael J. Freedman, Christopher Monsanto, Jennifer Rexford, Alec Story, and David Walker. Frenetic: A Network Programming Language. In *ACM ICFP*, 2011.

44. Christopher Monsanto, Joshua Reich, Nate Foster, Jennifer Rexford, and David Walker. Composing Software-Defined Networks. In *USENIX NSDI*, 2013.

45. Tim Nelson, Arjun Guha, Daniel J. Dougherty, Kathi Fisler, and Shriram Krishnamurthi. A Balance of Power: Expressive, Analyzable Controller Programming. In *ACM SIG-*

COMM HotSDN, 2013.

46. Lihua Yuan, Chen-Nee Chuah, and Prasant Mohapatra. ProgME: Towards Programmable Network Measurement. In *ACM SIGCOMM*, 2007.

47. Minlan Yu, Lavanya Jose, and Rui Miao. Software Defined Traffic Measurement with OpenSketch. In *USENIX NSDI*, 2013.

Author Biographies

Peng Sun is a Ph.D candidate in the computer science department of Princeton University. He received his B.E. in Electronic Engineering from Tsinghua University in 2010 and his M.A in Computer Science from Princeton University in 2012. His research interests are in large-scale network management and distributed systems for efficiency and reliability.

Minlan Yu is an assistant professor in the computer science department of University of Southern California. She received her B.A. in computer science and mathematics from Peking University in 2006 and her M.A. and Ph.D in computer science from Princeton University in 2008 and 2011. After that she was a postdoctoral scholar in UC Berkeley for one year. She has actively collaborated with companies such as AT&T, Microsoft, and Bell Labs. Her research interest includes data networking, distributed systems, enterprise and data center networks, network virtualization, and software-defined networking. She received ACM SIGCOMM doctoral dissertation award in 2012 and Google research award in 2013.

Michael J. Freedman is an Associate Professor in the Computer Science Department at Princeton University. His research broadly focuses on distributed systems, networking, and security, and has led to commercial products and deployed systems reaching millions of users daily. Honors include a Presidential Early Career Award (PECASE), Sloan Fellowship, NSF CAREER Award, ONR Young Investigator Award, DARPA CSSG membership, and multiple award publications.

Jennifer Rexford is the Gordon Y.S. Wu Professor of Engineering in the Computer Science department at Princeton University. Before joining Princeton in 2005, she worked for eight years at AT&T Labs–Research. Jennifer received her BSE degree in electrical engineering from Princeton University in 1991, and her PhD degree in electrical engineering and computer science from the University of Michigan in 1996. She is co-author of the book "Web Protocols and Practice" (Addison-Wesley, May 2001). She served as the chair of ACM SIGCOMM from 2003 to 2007. Jennifer was the 2004 winner of ACM's Grace Murray Hopper Award for outstanding young computer professional. She is an ACM Fellow (2008), and a member of the American Academy of Arts and Sciences (2013) and the National Academy of Engineering (2014).

David Walker is a Professor of Computer Science at Princeton University, where he studies programming language theory, design and implementation. He also currently serves as an associate editor for ACM Transactions on Programming Languages and Systems and is program chair for the ACM Symposium on Principles of Programming Languages in 2015.