

Flexible Network Bandwidth and Latency Provisioning in the Datacenter

Vimalkumar Jeyakumar¹, Abdul Kabbani², Jeffrey C. Mogul², Amin Vahdat^{2,3}

¹Stanford University, ²Google, ³UCSD

Abstract

Predictably sharing the network is critical to achieving high utilization in the datacenter. Past work has focused on providing bandwidth to endpoints, but often we want to allocate resources among multi-node *services*. In this paper, we present Parley, which provides service-centric minimum bandwidth guarantees, which can be composed hierarchically. Parley also supports service-centric weighted sharing of bandwidth in excess of these guarantees. Further, we show how to configure these policies so services can get low latencies even at high network load. We evaluate Parley on a multi-tiered oversubscribed network connecting 90 machines, each with a 10Gb/s network interface, and demonstrate that Parley is able to meet its goals.

1. Introduction

Multi-tenancy is inevitable at large scale. In Google’s datacenters, we see on average tens of jobs per server. A typical server is shared by applications that have diverse performance requirements from the network, such as:

- **Bandwidth Intensive:** A MapReduce job during its read and shuffle stages; or, large file copies.
- **Latency Sensitive:** The front-end for user-facing services, such as web search.
- **Bandwidth and Latency Sensitive:** Infrastructure services, such a distributed file system (DFS), often have a mix of bandwidth and latency sensitive flows.

These applications can also be **adversarial**. For example, tenant virtual machines (VMs) sharing the same network, in public clouds such as Amazon EC2, Google Compute Engine, HP Cloud, and Windows Azure, are not likely to be cooperative.

Prior work (see [18] for a detailed summary) in this space has already made the case for sharing bandwidth across entities that are coarser-grained than a TCP flow. For example, Oktopus [2], Gatekeeper [23], EyeQ [14] and ElasticSwitch [20] all provide, for a tenant’s collection of VMs, the abstraction of a dedicated physical network, with a specified guaranteed bandwidth for each endpoint VM. This abstraction of a shared network can be useful for providers wishing to support predictable behavior for tenant applications provisioned at a VM granularity.

However, these prior systems that support bandwidth guarantees raise the following questions:

- How should we realize notions of sharing that are (a) more flexible than guarantees (e.g. *weighted* sharing); and (b) service-centric, where services are a collection of endpoints?
- How should the service provision bandwidth to achieve service-level objectives, such as a bounded 99th percentile (tail) latency?

In this paper, we present Parley, a system we built to understand answers to the above questions. Parley has several features which make it attractive for practical deployment: (1) it supports a mixture of policies, such as bandwidth guarantees to service end-points, and hierarchical, weighted sharing, constructed by nesting service endpoints; (2) it uses a simple model of traffic characteristics to predict a service’s tail latency based on its own, and collocated services’ provisioned bandwidth, that is accurate at high loads.

We built Parley by systematically leveraging prior bandwidth sharing systems to achieve our goals (we modify the open-sourced system EyeQ [14]). There are two parts to Parley: (a) static provisioning, which the provider specifies when services are instantiated, and (b) runtime provisioning, which Parley determines based on bandwidth usage. While bandwidth guarantees are easy to statically provision using admission control, Parley differs from EyeQ in that runtime provisioning for hierarchical and weighted sharing requires information about services’ bandwidth usage.

For example, Figure 1 shows an example of a static sharing policy, in which the DFS service is allocated at least 6Gb/s, and at most 8Gb/s; and all (non-paying) VMs are allocated a maximum of 1Gb/s. That aggregate of 1Gb/s is then shared in a ‘weighted max-min’ [4, Sec. 6.5.2] fashion among VMs in the rack. Each VM’s allocation depends on their own, and the total bandwidth consumption of other VMs in the rack. This requires global service-level visibility of bandwidth usage that EyeQ does not have.

Parley adds this missing piece of global visibility on top of EyeQ, while retaining its strengths, as follows: At the core of Parley’s runtime system is a “bandwidth

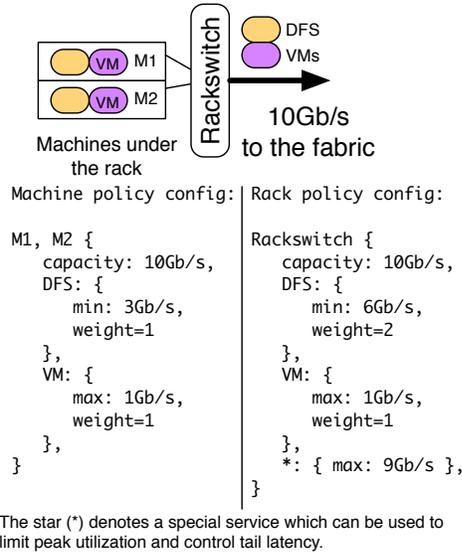


Figure 1. Parley supports a hierarchical sharing policy. The top figure shows the service placement on machines under the rack of a given capacity, and the bottom half illustrates sharing policies at the machine and rack.

broker” that allocates capacities to services in a hierarchical fashion. At the lowest level, each VM (or job, or service endpoint) on a machine is allocated an aggregate transmit and receive *hose capacity*. These capacities are periodically tuned by the bandwidth broker to conform to the global policy, based on (i) service endpoint’s machine-level utilization (§3.2.1), (ii) service’s rackswitch-level utilization (§3.2.2), and (iii) service’s overall fabric utilization (§3.2.3). Hose capacities are enforced directly in the *dataplane* using distributed, end-to-end congestion control. This decomposition of the global sharing objective into rack and machine sub-problems enables Parley to scale to a large datacenter.

Specifically, the main contributions of our paper are:

- The design and implementation of Parley that manages datacenter network bandwidth among services, in a flexible, and scalable fashion. Parley adds support for hierarchical allocations and distributed rate limiting, generalizing the notion of bandwidth guarantees provided by systems such as EyeQ and ElasticSwitch.
- Demonstrating that a service’s soft-realtime latency requirements can be met by appropriately controlling peak network load.

2. Background and Requirements

Datacenter bandwidth is not free. Modern datacenter network topologies are typically multi-staged Clos topologies [10] with few over-subscription points (typ-

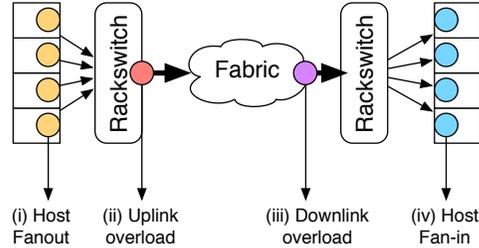


Figure 2. The contention points for intra-cluster bandwidth that Parley addresses.

ically only at the top-of-rack switch). Measurements from a highly-utilized datacenter show that packet drops occur predominantly at such over-subscription points in the topology, as shown in Figure 2. The contention can be broadly classified into three cases:

- **Host level:** multiple jobs sharing a host share the host (NIC) transmit and receive bandwidth. On the transmit side, the host network stack is back-pressured, so we do not see any drops. However, a large fraction of dropped packets happen at the receiver, due to incast-like traffic patterns. We call this *host fanin*.
- **First rackswitch:** Due to rackswitch over-subscription, hosts under a rack can overwhelm the limited-capacity uplinks. We call this the *uplink overload*.
- **Last rackswitch:** Finally, some small fraction of drops happen *within the fabric*, because the fabric overwhelms the rackswitch’s limited capacity to its hosts. We call this *downlink overload*.

Very few drops occur elsewhere within the fabric. We observe similar trends in other highly-utilized clusters.

Our goal is to address contention-induced packet drops, because TCP is the dominant protocol in the datacenter, and its default response to dropped packets aims at fair per-flow bandwidth allocation—without taking service-level requirements into consideration.

2.1 What does ‘predictability’ mean?

Our primary goal is to share bandwidth in a programmable and well-specified fashion, and to have low, bounded tail latency. We argue that many seemingly different notions of predictability can be mapped to offering guarantees on aggregate bandwidth at a contention point.

Service-level guarantees. Consider large-scale data-analysis jobs with multiple MapReduce stages, which care most about total completion time. The total job completion time is a complex function of the number of CPUs (number of workers), the available network bandwidth, and the number of barriers (stages) in the job [8]. Job completion time generally does not depend on short-term network latency. Parley allows the job scheduler to request bandwidth while accounting for contention at

the machine and rack level. This explicit guarantee allows the job scheduler to make better decisions [8] when placing jobs and orchestrating transfers, *enabling* predictable job completion times.

Latency. Infrastructure services also have stringent requirements on tail latency (e.g., <20ms response time for 1MB IO operations at the 99th percentile). Achieving tail latency guarantees is hard, as this fundamentally depends on the stochastic nature of flow arrivals. This fundamental dependence cannot be avoided. However, given a model of flow arrivals, a bandwidth guarantee can *bound* tail completion times.

For example, if the inter-flow arrival times are exponentially distributed, and their processing times (i.e. flow sizes) are also exponentially distributed, then the distribution of flow completion times can be explicitly written down as a function of the service instance’s allocated bandwidth. That is, in a steady state $M/M/1/FIFO$ queue, the flow completion time t (i.e. waiting time + service time) has the probability density function [12]

$$f(t) = \begin{cases} \mu(1-\rho) \exp(-\mu(1-\rho)t) & t > 0 \\ 0 & \text{otherwise.} \end{cases}$$

where μ is the average service *rate* (a function of the available capacity), and $\rho \in (0, 1)$ is the average load on queue. For example, if the average flow size 1MB, μ is 1.25 per ms at 10Gb/s; we can then deduce that 99% of flow completion times are <18.4ms.

By empirically measuring the distribution of flow sizes, service times, and inter-arrival times (all of which change relatively infrequently), we can do an offline simulation to derive a bandwidth guarantee requirement that will support robust bounds on a flow’s tail latency. Though the formula makes an assumption about the distribution of flow sizes and their inter-arrival times, we discuss in §4 how we can bound latency without modeling or making assumptions about the arrival process, using a simple counting argument.

Thus, to deal with latency requirements, Parley can explicitly guarantee aggregate capacity to service endpoints, and control the peak load on the network.

Summary: The main takeaway is that for many use cases, explicit bandwidth guarantees at the level of service endpoints (e.g. VMs), and control of the peak load on the network, helps application-specific schedulers to maintain predictability for their own metrics. Hence, Parley provides knobs to share network bandwidth in a flexible fashion, and strives to meet these requirements as quickly as possible.

2.2 Related Work

Many works address bandwidth management within a cluster. Since we build on top of EyeQ [14], we inherit

its benefits over prior work such as Seawall [24], Gatekeeper [23], Oktopus [2], and FairCloud [21]. We therefore focus on more recent work.

Both EyeQ [14] and ElasticSwitch [20] only offer work-conserving (and non work-conserving) bandwidth guarantees to individual service endpoints. CloudMirror [16] argues for asymmetric transmit and receive guarantees. However, they are limited in their ability to handle network congestion in a flexible fashion. If the network is congested, bandwidth is shared in a max-min fashion either across receiving endpoints (EyeQ), or across contending source/destination host pairs (ElasticSwitch), which need not necessarily match service-specific goals. In §6.3 we show how this flexibility helps in controlling tail latency at high loads.

The bandwidth enforcer in Google’s B4 SDN WAN [13] allocates bandwidth at the granularity of (QoS,src-datacenter,dst-datacenter) tunnels. This approach does not scale to allocations within a data-center, due to the large number of (source,destination) pairs. Instead, Parley’s optimization decomposition decouples mechanisms for pairwise allocations between source/destination hosts from service-level allocations, which helps scalability.

Hadrian [1] describes a hierarchical, work-conserving hose model to guarantee a tenant a minimum inter-tenant bandwidth. However, Hadrian requires switch changes to count the number of flows on a per-tenant basis, which changes quite frequently. It is unclear if this is practical, and whether Hadrian’s allocations are stable over time. However, the policy framework we describe in this work is rich enough to capture inter-tenant isolation as well.

DRL [22] focuses on distributed rate limits by having service endpoints exchange demands using a gossip-protocol. However, distributed rate limits alone is insufficient in sharing bandwidth in a predictable fashion, as a service can target its entire bandwidth share at a single machine. Our approach is inspired by DRL, but uses a hierarchical decomposition to explicitly handle the major oversubscription points. For instance, by aggregating service endpoints at a machine and rack level, we flexibly share limited machine and rack bandwidth.

Commercial switches [6] do support hierarchical sharing policies. However, these policies are limited: First, switches can only drop packets, and we need traffic admission control to deal with malicious traffic. Second, switches only offer a link-centric sharing policy. In practice, we desire an *end-to-end* view for allocating bandwidth, regardless of the physical topology.

The line of work on Dominant Resource Fairness (DRF) [9, 5] defines a fairness metric when sharing *across* multiple resource types, such as CPU, memory, and network, but rely on other mechanisms to *enforce*

these allocations. Parley can be used as an enforcer for network bandwidth.

3. Design

We now detail the design of Parley, starting with its approach to specifying bandwidth guarantees. We then describe the high-level system architecture, and each component in more detail.

3.1 Specifying Sharing Policies

In Parley, a policy is specified with respect to hierarchies of services at one contention point. Services are traffic bundles that are uniquely identifiable using a packet filter; a service can be a single endpoint, a collection of endpoints, or sub-flows within an endpoint. The service hierarchy must be a tree (i.e. no loops). Each service has both a machine-level and rack-level policy, to cover the primary contention points in a datacenter. A service can have different policies for transmit and receive.

For instance, at the machine level, all flows terminating at the Distributed File System (DFS) service port are part of the DFS service, even if they involve multiple source servers. At the rack level, the DFS endpoints can be aggregated into one service S1, and VM endpoints can be aggregated into another service S2, but a DFS endpoint cannot be part of both S1 and S2.

A flow between a source and destination can be part of multiple services; e.g., traffic from a MapReduce client to a DFS server consumes bandwidth both at MapReduce and DFS, and is therefore charged to both services. The guarantee for traffic between two services is limited to the minimum of the two service-specific guarantees. E.g., a flow between a MapReduce client guaranteed 2Gb/s and a DFS server guaranteed 1Gb/s will be guaranteed only 1Gb/s.

A policy is configured using several parameters:

- **Min bandwidth:** the guaranteed bandwidth for the service; the default is 0, implying no guarantee.
- **Max bandwidth:** the limit on the service's bandwidth; the default is "unlimited."
- **Weight:** excess bandwidth (what remains after all guarantees have been met) is shared among contenders based on their weights (subject to their max limits). Weights encode a service's relative importance. The default weight is 1.

Policies can be specified statically, or can be changed any time to support dynamic reservations [27]. Note that in the case of guarantees, admission control and job placement must ensure that the guarantees can be satisfied in the worst case (i.e., when all services demand their full guaranteed bandwidth). When aggregating services, the guarantee for the parent service must at least be the sum of guarantees of its child services.

The most constrained policy determines the service allocation. For instance, consider a situation where there

are 10 MapReduce jobs in a rack, and each job has machine-level policy (weight = 1, max = 1Gb/s). If MapReduce jobs in aggregate have a rack-level policy (max = 5Gb/s), and all jobs are active, each job can send at 0.5Gb/s. However, if only one job is active, it cannot grab the entire 5Gb/s, since its machine-level policy is the most constrained, and takes effect. Thus, besides the static policy, there is a dynamically computed *runtime policy* which is actually enforced. Note that service policies are implicitly constrained by machine and rack capacities.

Why have hierarchical policies? A hierarchical policy is one way to quantify the over-provisioning *risk*, which makes sharing *unambiguous* at every contention point, instead of letting services battle it out between themselves. Without hierarchy, we would have to provision for the worst case simultaneously at each contention point, which usually does not happen often, and can waste capacity. Hierarchies are also useful from an operational perspective; since Parley exposes both guarantees and runtime policies for every service, it is easier to debug when a service does not get its bandwidth.

3.2 Design components

Parley consists of three components, each responsible for computing and enforcing both transmit and receive bandwidth allocations, at different contention points. Figure 3 shows the high-level architecture. These components work together to enforce the static policy.

- The **Machine shaper** handles host fanin and fanout. Its job is to enforce service shares only at a machine granularity, according to the static and runtime machine policy. It allocates bandwidth at the granularity of (source,destination) pairs of communicating hosts [14].
- The **Rack broker** handles the downlink and uplink overload on every rack. Its job is to dynamically compute per-service per-machine policy, so that the static rack policy is enforced by the machine shaper.
- The **Fabric broker** handles distributed rate limits, to enforce bandwidth caps on various services at the global scale. It dynamically computes a per-rack, per-service policy to enforce the static fabric policy.

Figure 1 in §1 shows an example with two services: VMs and DFS. To simplify discussion, we assume there are no fabric-level shares for both services. There are two machines with 10Gb/s capacities under the rackswitch, which has a 10Gb/s bandwidth to the fabric. The rack sharing policy is to divide bandwidth such that VMs gets at most 1Gb/s, and DFS gets at least 6Gb/s. Thus, when every service is active, (M1,VM), (M2,VM) will get 0.5Gb/s each, and (M1,DFS), (M2,DFS) are allocated 4Gb/s each. When (M2,DFS) is idle, (M1,DFS) is allocated 8Gb/s, and

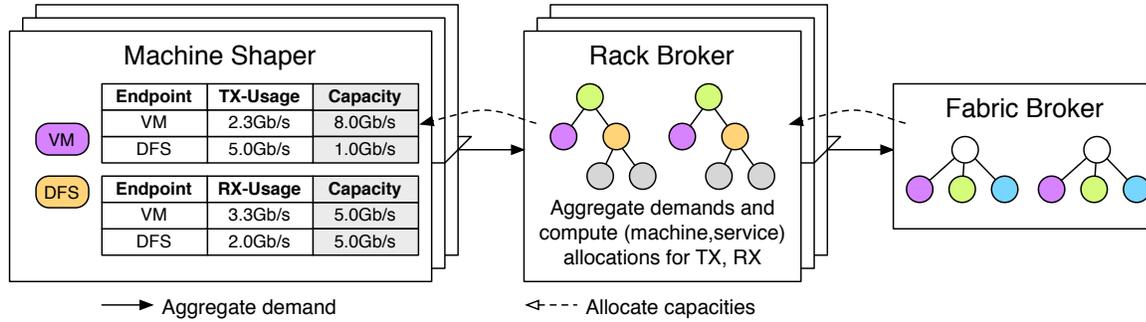


Figure 3. Parley consists of distributed machine shapers, rack-local brokers and a fabric broker that work together to enforce policies at different contention points.

when all VMs are idle, (M1,DFS) is allocated the entire 9Gb/s.

Note that the rack broker only computes (machine,service) shares, and is not concerned with traffic patterns that change frequently. The machine shaper works at fast timescales, to ensure machine runtime policy is enforced regardless of fine-grained communication patterns. This timescale requirement has a direct impact on the design of the individual components, which we detail next.

3.2.1 Machine shaper

A machine may run a number of services, each perhaps in its own virtual machine. The problem of enforcing machine-level bandwidth shares is the same problem solved by EyeQ [14]: On the transmit side, the machine shaper consists of a single root rate limiter that enforces an aggregate rate limit on the service’s traffic leaving the machine. In addition, per-destination rate limiters enforce receive-bandwidth shares. The receive-bandwidth shares are computed at the receiver, and signalled back to the sources using a feedback packet.

On the receive side, each service is attached to a rate meter, which is allocated some capacity C . The rate meter periodically measures the aggregate bandwidth utilization $y(t)$ of the service, and continuously computes *one* rate $R(t)$ such that the aggregate utilization $y(t)$ matches the service capacity C . This is done iteratively using a control equation:

$$R(t+T) = R(t) \times \left(1 - \alpha \frac{y(t) - C}{C} - \mathbb{1}_{\text{marked}(t,t+T)} \frac{\beta}{2} \right)$$

where T is a configurable time period, α is a parameter chosen for stability. The value β is the fraction of Explicit Congestion Notification (ECN) marked packets, and is used only if there are marked packets in the interval t and $t+T$ (denoted as $\mathbb{1}_{\text{marked}(t,t+T)}$). The control equation water-fills $R(t)$ such that the utilization $y(t)$ matches capacity C .

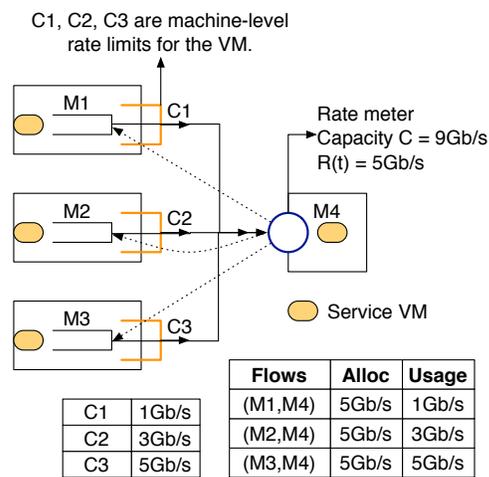


Figure 4. An example of a max-min fair allocation (across senders at a receiver) computed by the machine shaper. The allocation is computed at the receiver, and communicated to senders using feedback messages. The end-to-end allocation is computed in a distributed fashion using only local information.

The receiver samples incoming packets (1 every 10kB) and sends feedback to the source IP address. Sampling ensures that heavy hitters are rate limited. The feedback message encodes $R(t)$ serialized in a custom IP packet. When a sender service receives feedback from a particular destination, it creates/updates a rate limiter to the destination, under the sender’s root rate limiter, to enforce the bandwidth allocation (as shown in Figure 4).

Weighted allocation across senders. When two services communicate with each other, their receivers can divide allocated capacity in a weighted fashion across senders: The sender simply scales its feedback to $w_{\text{sender}} \times R(t)$. The weight ensures that the rates of senders 1 and 2 are always in the ratio $w_1 : w_2$.

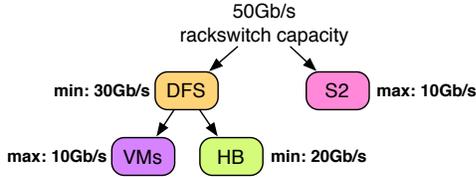


Figure 5. Inter-tenant sharing is a special case of hierarchical sharing.

Inter-tenant sharing. For the heavy bandwidth consumers such as DFS, a common requirement is to share bandwidth among its own users, according to some specified policy (typically weighted shares). Since the notion of a ‘service’ is flexible, it can further be *deaggregated* into ‘service:user.’ For instance, if HBase (HB) and VMs both use DFS, we can create the service hierarchy shown in Figure 5 (S2 being another service). Deaggregation enables delegation, where services can use Parley to manage their bandwidth. Each packet is associated with a leaf node identified using packet classification rules installed on the machine, both on the transmit and the receive path.

Parameter guidelines: We use the same parameter settings as in EyeQ: $T = 200\mu s$ and $\alpha = 0.5$. Note that the receiver *does not keep track of the number of senders*. This is by design: datacenter measurements [3] show high flow-arrival rates at a host (hundreds/second), which makes per-flow tracking expensive and complex.

3.2.2 Rack broker

The machine shaper realizes an allocation using only information locally available at a machine, so it can only satisfy the machine policy. The rack broker aggregates service-level bandwidth usages across machines in a rack, uses these values as ‘demands,’ and computes a machine-level transmit/receive runtime policy for each service. This policy determines the service’s transmit and receive capacity. The transmit capacity is set on the (machine,service) root rate limiter, and the receive capacity is set on the (machine,service) root rate meter. Once these capacities are set, the machine shaper enforces these capacities continuously. *Note that the rack broker does not track (source-ip,destination-ip) communicating pairs*. This is the key difference between our mechanism and a similar control mechanism in Oktopus [2, §4.3].

The design of the rack broker is best explained by an example. Consider the desired sharing policy shown in Figure 6 (assume for now that the demands were computed before the rack broker kicks in). The problem of determining the per-machine service policy that satisfies the rack policy can be done in two passes:

1. A bottom-up pass from child to root that computes the aggregate demand at each intermediate node in the sharing hierarchy.
2. A top-down pass that computes the new runtime policies at each node in the sharing hierarchy, according to the parent’s policy.

We make two assumptions: First, knowing the rack’s capacity to the fabric core (i.e., the 50Gb/s in Figure 6) greatly simplifies bandwidth allocation. This is feasible as datacenter operators have full control over their network infrastructure, and so the rack broker can periodically query (say) an OpenFlow controller about the rack uplink capacities. Second, we assume that the uplinks are evenly utilized, so we can treat them as one single link. But recall that the machine shaper’s allocations fall back gracefully even if there is in-network congestion (which is often only transient).

Allocation Algorithm: The runtime weighted max-min allocations and min/max guarantees can be computed by the classical water-filling algorithm [4, § 6.5.2]. Each iteration of the algorithm satiates the demand for one service. We do not rate limit endpoints whose demand is less than its capacity determined by the water-fill algorithm, as rate limiting shows a flow’s progress, and increases its flow completion time (§4,§7).

3.2.3 Fabric broker

The rack broker allocates bandwidth based on a capacity assigned to the rack, and based on (rack,service) limits. These limits can in turn be adjusted based on the global fabric bandwidth consumption of a particular service. The rack broker at each rack communicates the (rack,service) demands to a global fabric broker, which uses these demands and in turn computes new (rack,service) allocations using the same max-min algorithm shown above. The fabric broker operates at a slower timescale, running every $T_{\text{fabric}} = 10$ seconds.

3.3 Scalability

The EyeQ [14] paper discusses scalability of the machine shaper in detail. We therefore focus on the rack and fabric brokers, both of which have the same design.

Using measurements from production clusters, we find that the approach that aggregates (machine,service) tuples across machines under a rack and fabric scales well. To make the right allocations, the rack broker only needs to know the bandwidth utilization of the heavy bandwidth consuming services. If we define a “heavy rack-consumer” as a service using at least 100Mb/s of rack bandwidth (averaged over 5-minute intervals) we observe 10–100 of such heavy consumers per rack. If we set a higher threshold of 1Gb/s for a “heavy fabric-consumer” the number rapidly diminishes. It is therefore practical to collect and process service-level usages.

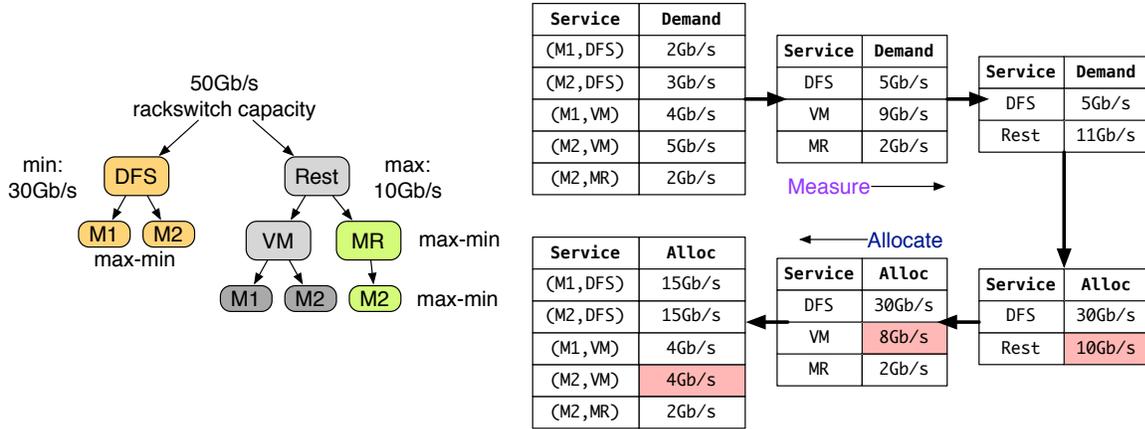


Figure 6. An example illustrating the hierarchical max-min allocation, computed at the rack level in two steps: (i) aggregate measurements, and (ii) enforce allocations. The red boxes denote the only entities (service instances) that are rate-limited.

3.4 Optimization Decomposition

The key idea behind splitting the problem into multiple sub-problems is the principle of optimization decomposition. A sharing objective is decomposed into multiple sub-problems that execute in a distributed fashion. We refer the reader to [19] for a detailed discussion on the mathematical foundations of optimization decomposition.

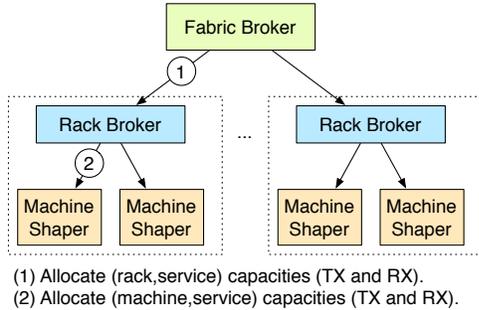


Figure 7. Optimization decomposition in Parley.

Parley uses this approach, which has also been adopted by other bandwidth-management systems [13, 25], and inherits several benefits:

- **Scalability:** Each shaper/broker is responsible only for a portion of the datacenter, and the optimization decomposition leads to a design that scales with the datacenter size. The machine shaper deals with no more than a few tens of services; the rack broker deals with no more than a few 100s of (machine,service) pairs in the rack; and the fabric broker deals with more no than tens to hundreds of racks.
- **Fault tolerance:** The scale-out design ensures graceful degradation, as un-failed parts of the system will continue normal operation. For instance, since the

machine shaper uses ECN marks to detect when the network is congested, the rack broker can fail without significantly affecting the system. The allocations degrade gracefully, and by design, the receivers will get a max-min share of the rackswitch bandwidth.

The main disadvantage of optimization decomposition is that the bandwidth allocation according to specified policy is not computed in one shot, but over a sequence of iterations. This can delay the speed at which a service can get its full allocation. While this is typically not a big concern at fabric scale, a service must be able to grab its fair bandwidth share quickly at the rack level. Our design choice to not rate limit services that are not consuming their full bandwidth share gives them the ability to ramp up quickly.

One possible concern is that services, when not rate limited, can burst to gain an unfair share of bandwidth and affect other service allocations. The effect of this burst is limited, as the machine-shaper kicks in quickly (within hundreds of micro-seconds) during transient network congestion, ensuring that services are not starved.

3.5 Are demands stable?

The bandwidth brokers operate at three different timescales: (i) the machine shaper operates at network round-trip timescales, (ii) the rack broker operates at timescales of seconds, and (iii) the fabric broker operates on multiple tens of seconds. One might wonder if the (machine,service) demands are stable for a long enough period that the rack broker can kick in, and the same for the global fabric broker. Figure 8 presents evidence that demands are indeed stable, so the individual brokers can operate effectively over the timescales they are designed for.

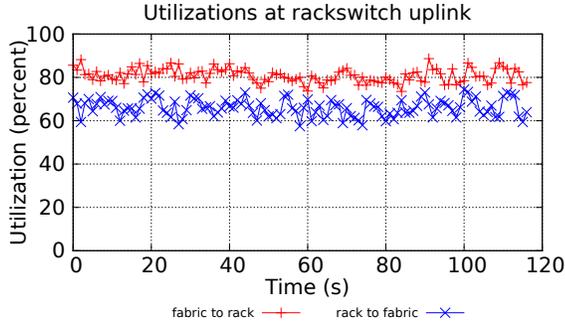


Figure 8. Time series of uplink and downlink utilizations (averaged over 1s) of a highly utilized rack, showing that the link utilizations are fairly stable over multiple seconds.

By design, the rack broker kicks in and enforces allocations only when the bandwidth usage is above limits specified by policy. Figure 8 shows the utilization (averaged over 1 sec.) of one highly-utilized rack in a production data center. We see that the aggregate transmit and receive utilizations remain fairly stable over multiple seconds, and thus the timescale of a few seconds is pragmatic. We don't have visibility at smaller timescales in production traffic.

4. Provisioning for Latency

In §2, we showed how a bandwidth guarantee enables services to achieve a bound on the tail flow completion times. The formula, however, made an assumption that flow arrivals are Poisson, and that flow sizes are exponentially distributed with a given mean.

While this model serves as a useful guide to provision systems, it is often the case that flow arrivals are not Poisson. Prior datacenter measurement work shows that flow arrivals are *bursty* [3]. Therefore, we seek a guideline that does not make assumptions on flow sizes (except that they are finite), or flow service order (except that the system is work-conserving), or any particular flow arrival pattern. Since we make minimal assumptions, the bound will serve as a useful guideline to reason about the worst-case performance.

We model a latency-sensitive endpoint as a work-conserving queue of capacity C that is shared with other endpoints. It is impossible to have any guarantees unless we constrain the arrivals in some fashion. Suppose $B(t_1, t_2)$ refers to the number of bytes received by the queue between time t_1 and t_2 . Consider a constrained arrival process parameterized by (σ, ρ) [7] such that:

$$B(t_1, t_2) \leq \sigma + \rho C (t_2 - t_1) \quad (1)$$

for all $t_1 < t_2$. Here, σ quantifies the arrival burstiness, and $\rho \in (0, 1)$ quantifies the average long-term

rate. Then, we can show that the flow completion time $FCT(f)$ of any flow f of size $Z(f)$ satisfies:

$$FCT(f) \leq \frac{\sigma + Z(f)}{C(1 - \rho)} \quad (2)$$

The proof is by a counting argument: Let $S(f)$ and $F(f)$ denote the start time and finish time for a flow f . Now, the number of bytes arriving to the server between $S(f)$ and $F(f)$ is constrained by Equation 1. We have:

$$F(f) \leq \text{Start time} + \text{Waiting time} + \text{Service time}$$

$$F(f) \leq S(f) + \frac{B(S(f), F(f))}{C} + \frac{Z(f)}{C}$$

$$(1 - \rho) F(f) \leq (1 - \rho) S(f) + \frac{\sigma + Z(f)}{C}$$

Since $\rho \in (0, 1)$, the quantity $(1 - \rho) > 0$, and therefore:

$$F(f) - S(f) \leq \frac{\sigma + Z(f)}{C(1 - \rho)}$$

$$\implies FCT(f) \leq \frac{\sigma + Z(f)}{C(1 - \rho)}$$

Alternatively, this bound can be used as a guideline to limit the peak load ρ on the network to keep FCT under a desired value. The bound only assumes a (σ, ρ) constraint on arrivals to the work-conserving queue, and that latency-sensitive flows are not rate limited. A rate-based congestion control can control the long-term load ρ . The burst size σ is the maximum of two components: (i) the excess line-rate burst due to delayed convergence of the congestion control algorithm, and (ii) the burstiness of sender rate limiters. From experiments, we found that the burst due to convergence time is the dominant factor; i.e., if the congestion control algorithm converges in t seconds, then the queue sees a maximum burst of $C \times t$.

For example, if $C=100\text{Mb/s}$, and the congestion control algorithm takes 10ms to converge, then the maximum burst size is about 83 MTU sized packets. Once the congestion control algorithm converges, the only burstiness is due to rate limiters at senders. To quantify this burst, we set up an experiment using the Mininet network emulator as follows: 100 hosts talk to one receiver using long-lived TCP flows rate limited to to 1/100th of the desired load ρ at the receiver. The rate limiters were configured with a maximum burst size of 64kB, and we measure the queue sizes at the receiver. Figure 9 plots the distribution of queue sizes for various values of ρ . We find that even at 90% load, the 99th percentile queue size is less than 25 packets, which is smaller than the

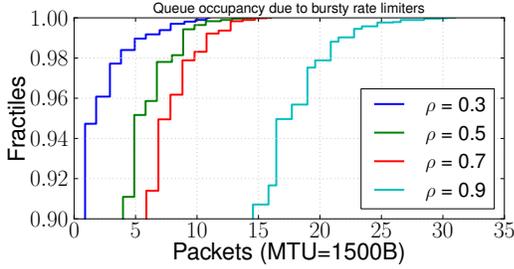


Figure 9. CDF of queue sizes as a function of load.

maximum burst due to delayed convergence (83 packets), indicating that the burst contribution due to the delayed convergence of congestion control algorithm plays a dominant role in determining the worst case flow completion time.

Though the bound is on the worst case FCT, flows can be reordered (e.g. through prioritization) to achieve better individual FCTs. How a service schedules its own flows in accordance to bandwidth guarantees is completely up to it. We echo the insight of recent work pFabric [alizadeh2013pfabric] that rate control is perhaps a poor choice for fine-grained prioritization across *mice flows* (e.g. flows less than a few 100kB). Parley’s rate control is decoupled from the flow’s priority. In practice, datacenter measurements [alizadeh2010data, 10] show that most of the bandwidth consumption is from the elephant flows that last long enough to be accurately rate limited (flows between a few MB to 100s of MB last a few msec. to 100s of msec.).

5. Implementation

Our prototype of Parley builds on top of EyeQ.¹ We first briefly review the implementation of the machine shaper in EyeQ, discuss our new improvements, and then describe the implementation of the rack broker and fabric broker. Though we prototyped Parley in software, its control logic (rate brokers) could interface with a hardware datapath (e.g., rate limiters on a NIC, or in switches within the fabric).

5.1 Machine shaper

The machine shaper consists of a transmit and receive side component. As in EyeQ, we use hierarchical rate limits to enforce sharing policies on the transmit side. At a service endpoint on a machine, the rate limiter hierarchy is shown in Figure 10. In the hierarchy, rack uplink traffic is controlled by rate limits set by the rack broker (on the service root rate limiter). This means that inter- and intra-rack traffic share the same rate limit if there is congestion at the rack uplinks; however, when we look at measurements from several datacenters, we see that most of the traffic is *inter-rack*, so we did not optimize

¹<http://jvimal.github.io/eyeq/>

for fate sharing. Nevertheless, we could separate inter- and intra-rack traffic below the service root rate limiter, and have the rack broker set limits only on inter-rack traffic. The rate limiters and rate meters are designed for high throughput using per-CPU datastructures.

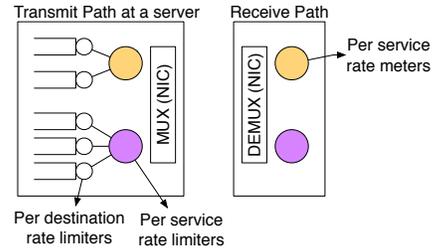


Figure 10. The rate limiter hierarchy at the sender, and rate meters at the receiver, shaded by service type.

5.2 Rack broker

We implemented the rack broker as a user-space program in C++. There are many ways to query the aggregate link utilization across services: the program can install counters on the rackswitch and use OpenFlow [17] to count the per-service utilization. However, for ease of deployment, we implemented a simple approach to aggregating the (machine,service) counters. The rack broker is distributed across machines, and each machine queries the local machine shaper to obtain a list of (machine,service) pairs. Each machine then broadcasts this list of tuples to other machines in the rack, and aggregates tuples from other machines. Thus, each machine will have its own copy of the list of (machine,service) counters for all machines and services in the rack. Using this information, each machine runs the water-fill algorithm to determine all (machine,service) allocations, and each machine locally enforces allocations for its own services. The machines broadcast the counters using UDP once every T_{rack} seconds.

Tolerance to failures: In the absence of rack shaper, the machine shapers fall back gracefully to a max-min fair allocation across receiving VMs/service instances (as in EyeQ). Thus, rack broker failures do not cause global failures. However, the rack broker is highly fault tolerant, as it runs on every machine. If a few machines die, the remaining rack brokers continue working as expected. If the broadcast packets are intermittently lost, the last updated value of (machine,service) is used to compute bandwidth shares. We prevent small rate allocations from sticking permanently by having machine shapers reset their allocations to their static configuration after a long timeout. This timeout value T_{rack}^l is chosen so that the machine shaper can be certain about a rack broker failure.

5.3 Fabric broker

The fabric broker uses the (rack,service) shares for a given rack, to determine new (rack,service) allocations on a per-rack basis. The fabric broker runs relatively infrequently on a few machines spread across the data-center. On each rack, one of the rack brokers is designated as the leader, and the leader sends the rack-level bandwidth consumption to the fabric broker’s service address. As with rack broker, the fabric broker need not be perfectly reliable.

The rack leader sends an RPC to the fabric broker once every 10 seconds. Even with 10000 racks within a cluster, sending 10kB data (which encodes utilizations for over 1000 services) to the fabric shaper consumes only 80Mb/s of traffic at the fabric broker.

Tolerance to failures: The strategy that rack brokers use to handle fabric broker failures is exactly the same as the strategy machine shapers use to handle rack broker failures. However, we set a larger timeout, $T'_{fabric} = 50$ seconds. If the rack broker doesn’t hear from the fabric broker in 50 seconds, it resets the runtime policy to the statically configured policy.

6. Evaluation

We demonstrate that Parley achieves our goals of predictable bandwidth sharing.

Since we built Parley on top of EyeQ, we refer the reader to measurements in the EyeQ paper [14, §5.1]. Parley inherits EyeQ’s results for: (i) low CPU overhead when compared to state-of-the-art software rate limiters; (ii) millisecond-timescale convergence times at the machine shapers.

The new aspects of Parley that we cover are:

- Rack brokers are highly efficient, and consumes less than 2ms wall-clock time to allocate bandwidth to as many as 100k services. The bandwidth overhead is less than 3Mb/s when rack broker interval is 1s. Moreover, they converge to the fair rates within a few seconds.
- In an emulation, the fabric broker converges within 10s to limit the maximum bandwidth consumption across 100 racks.
- The rack broker can control the maximum utilization on the rackswitch, ensuring predictable tail latency for services in the rack, even in the presence of a malicious service that tries to grab all bandwidth.
- Parley, unlike EyeQ, controls tail latency even if the network core is congested, while allowing work-conserving allocations to services.

Topology and parameters: Unless otherwise noted, we use a leaf-spine topology, where 9 rackswitches are connected to 2 spine switches in a bipartite graph as shown in Figure 11. There are 10 hosts under each of the 9 rackswitches. Each host has a 10Gb/s NIC. The

RCP aggressiveness	α	0.5
RCP time interval	T	$200\mu s$
Rack broker frequency	T_{rack}	1s
Fabric broker frequency	T_{fabric}	10s
Rack broker timeout at machine broker	T'_{rack}	5s
Fabric broker timeout at rack broker	T'_{fabric}	50s
ECN marking threshold		80kB

Table 1. Parameters used in evaluating Parley.

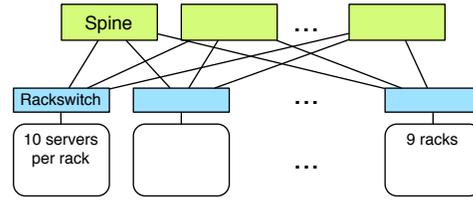


Figure 11. Topology of the testbed. The network is oversubscribed by a factor of 1.25 (100Gb/s to 80Gb/s) at the rackswitch.

network is oversubscribed 1:1.25 at the rackswitch. We use the Parley parameters shown in Table 1.

6.1 Machine shaper microbenchmarks

Graceful degradation during network congestion:

The machine shaper senses congestion when link utilization approaches capacity, which is not sufficient if the network is congested. Hence, the machine shaper also uses ECN feedback from the network to back off when links are congested. We evaluate this scenario by having two service instances under different racks send traffic to two their respective receivers under the same rack, but on two different machines. We induce network congestion by disabling all but one fabric link at the receiving rackswitch, creating a 2-to-1 over-subscription at the rackswitch.

Figure 12 highlights an important point that the machine shaper must operate at a fast enough timescale. At time $t=0$, only one service is active. At $t=40s$, both services are active, causing network congestion; therefore, bandwidth is shared almost equally (Jain’s Fairness Index is 0.99). At time $t=60s$, we increased T from $200\mu s$ to 1ms; notice that the machine shaper is now too slow to guarantee fair shares across receivers, as the transport layer (TCP) also backs off due to network congestion.

If we were allowed to change TCP to back off less aggressively, or to use UDP, a 1ms timescale for reaction time would be sufficient. We discuss the tradeoffs between accurate and coarse-grained rate limiting and its impact on flow completion times in §7.

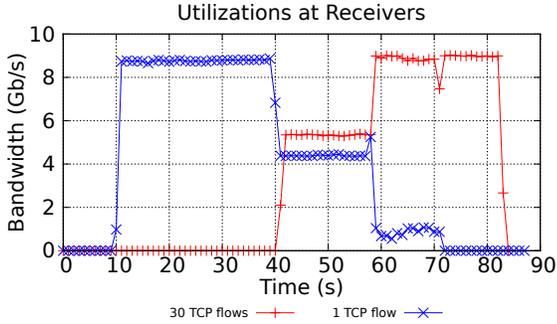


Figure 12. When there is network congestion at the rackswitch, Parley gracefully falls back to a sharing model where receivers get (almost) equal shares of the network bottleneck link.

6.2 Microbenchmarks: Rack and Fabric brokers

Bandwidth overhead: In our rack broker implementation, each machine broadcasts its own service utilizations to other machines under the rack. Each service can be represented as a 4B integer, and its utilization as another 4B floating point number. At 8B per service, even 1000 bandwidth-intensive services in a rack would generate only 8KB of data, plus some overhead for packet headers. If there are 40 machines in a rack, unicasting this data to each machine once/sec would cost less than 3Mb/s per machine.

Computation overhead: The rack broker repeatedly computes hierarchical max-min shares among N services. Table 2 shows the wall-clock time of computing these shares for a single level in a hierarchy, for various values of N (using one core of a 2.4GHz Linux desktop machine). We initialized demands randomly such that the load is close to the capacity. Although our implementation of the water-fill algorithm [4, Sec. 6.5.2] is $O(N^2)$, the wall-clock time to convergence scales well with N . This is because every iteration of the water-fill satisfies the demand for *at least* one service, and we track demands with a precision of 1Mb/s. Thus, it does not take more than ~ 80000 iterations to water-fill 80Gb/s capacity.

N	100	1k	10k	100k
Time	$2\mu\text{s}$	$12\mu\text{s}$	$320\mu\text{s}$	1.6ms

Table 2. Average wall-clock time per iteration of max-min share computation.

Convergence speed: The hierarchical max-min algorithm converges in *one invocation*, and therefore the convergence speed only depends on the speed of exchanging information and the computation time. The rack broker can also work at the rate of 10 times per second, consuming 30Mb/s bandwidth, converging to the allocations in 100ms. In practice, however, we found 1s

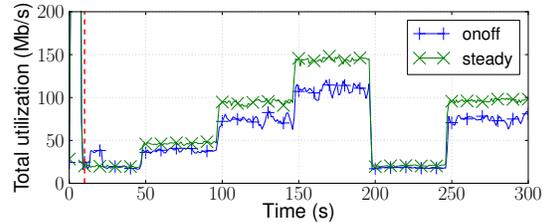


Figure 13. At $t=0$, we start the fabric broker, and it takes 10s for it to take effect. However, once the fabric broker starts to allocate bandwidth, it converges within a few iterations (at $t=30$ s) and limits the tenant’s usage to <20 Mb/s, despite its bursty traffic pattern.

time granularity to be sufficient in ensuring isolation, as the machine shaper works at faster timescales to ensure the network does not experience persistent congestion.

Fabric broker: To test the convergence of the fabric broker at large scale, we emulated a 100 rack cluster in Mininet [11] with one tenant that is rate limited to 20Mb/s. Each rack has a 100Mb/s link to the fabric; it picks another rack at random and sends a UDP data burst for 5s, and sleeps for another 2s until $t=300$ s (on-off). In another experiment, the traffic is less bursty (steady): it does not sleep for 2s, but instead, it immediately sends traffic to another random rack. The burst time is smaller than the 10s interval at which the racks communicate with the fabric broker. Figure 13 shows the time series of the aggregate utilization of the tenant: as we can see, the fabric broker converges after the first burst, and within a few iterations, the tenant is limited to its 20Mb/s, 50Mb/s, 100Mb/s, 150Mb/s, 20Mb/s and back to 100Mb/s every 50s (to illustrate convergence). In practice, the rack broker will ensure the initial burst does not adversely affect the performance of other services within the rack. This global service limit would not be feasible with proposals such as EyeQ/ElasticSwitch, as the service’s per-rack limit depends on bandwidth usages of other racks.

6.3 Macrobenchmarks

Policy: We demonstrate how Parley can protect the rackswitch links from persistent congestion on a real (i.e., non-emulated) network. On each machine in the topology, we instantiate two service endpoints A and B, each of which is given equal bandwidth at the machine level. At the rack uplink, service A is given at most 30Gb/s, and service B is given at least 30Gb/s, but the peak load on the rack is limited to 60Gb/s. Recall that Parley’s final allocation depends on the most constrained resource.

Throughput protection: One of the racks is designated as the receiver. The traffic pattern consists of long-lived transfers between every pair of sender and receiver (a full mesh). Figure 14 shows the utilization of each service at the receiving rackswitch. At the start of the

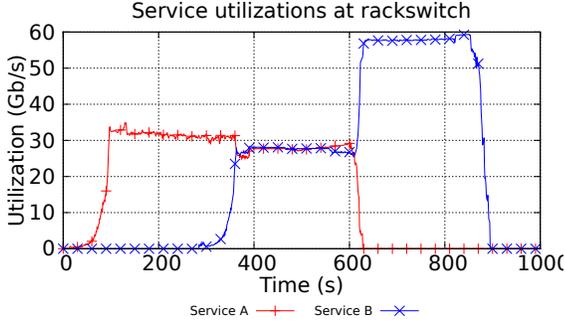


Figure 14. Two services A and B sharing limited rack-switch capacity. Service A can consume at most 30Gb/s, while service B can consume at least 30Gb/s but the total is limited to 60Gb/s. The utilization data is obtained every second, and the plot marks are every 30s.

experiment, only service A is active; notice that it uses only 30Gb/s. At about $T=300$ s, we start service B, which takes around 50s to ramp up. Parley converges faster than the time it takes for all flows of service B to ramp up, but by the time all flows are running at full line rate, service B is able to utilize its full 30Gb/s link capacity. When service A finishes, service B is able to utilize all the remaining capacity up to 60Gb/s.

Latency protection: Using the same policy as above, we show how Parley can protect the latency of a service against high-rate adversarial traffic. We run the same configuration as above, and both services mimic a client-server workload with all-to-all communication between senders spread across all but one rack, and receivers in the one remaining rack. There are 24 TCP connections between each (service,machine) pair.

Service A initiates 200kB transfers using RPCs such that its total ingress offered load is 14% of the receivers’ rackswitch capacity. Service B uses 1MB RPCs such that the total ingress offered load on the rackswitch is 15%, 50%, 70% and $>100\%$. Table 3 shows the 99th percentile FCT of 200kB RPCs, while we varied the total load at the rackswitch. The mean RPC inter-arrival time t_μ is chosen to match the load, and the inter-arrivals are sampled uniformly at random between 0 and $2t_\mu$. The 99th percentile FCT is computed after running for 20 minutes.

The results in Table 3 show, not surprisingly, that Parley does not significantly affect latency at offered loads below 100%. The table also shows that service A’s latency increases with the total offered load, even though A’s offered load is constant, since it is sharing contention points with service B. Therefore, one must assume that B is fully utilizing its maximum bandwidth allocation when deriving a tail-latency bound for A.

These results also demonstrate that without Parley, the tail latency of service A can become very large when

Service	Total offered load at receiver rackswitch			
	15%	50%	70%	$> 100\%$
<i>Without guarantees</i>				
A	2.34ms	3.50ms	6.17ms	939.79ms
B	7.89ms	7.64ms	14.30ms	-
<i>With EyeQ guarantees</i>				
A	2.33ms	3.44ms	6.03ms	701.26ms
B	9.12ms	7.71ms	14.19ms	-
<i>With Parley guarantees</i>				
A	2.34ms	3.42ms	6.08ms	38.74ms*
B	9.34ms	7.62ms	15.07ms	-
<i>Bounds (equation 2)</i>				
A	9.01ms	15.32ms	25.53ms	38.30ms
B	9.77ms	16.60ms	27.67ms	-

Table 3. The 99th percentile RPC completion time, as a function of the total load offered at the receiver rack-switch, with and without (w/o) Parley. The bounds on tail latency are computed using equation 2. As indicated in boldface, the latency without Parley suffers at high utilization.

B tries to increase its load past its allocated capacity, but Parley protects A’s tail latency. Naturally, B’s own tail latency becomes unbounded when its load exceeds its guaranteed bandwidth. Since the over-subscription is at the receiving rackswitch, and the load is uniformly spread across machines under that switch, Parley’s protection is primarily due to the rack broker’s allocations.

The total load at the rackswitch was limited to 60Gb/s (or about 80% of its capacity). For this experiment, the machine shaper iterates every 500μ s to recompute $R(t)$, and it converges (in the worst case) to within 30 iterations, to allocations within 0.01% of the ideal rate, regardless of the number of flows [14]. In practice, we found that fewer than 15 iterations suffice. Thus, the burst size σ can be bounded and is dominated by the convergence time of the congestion control algorithm. Table 3 also shows the bounds on the flow completion time computed using Equation 2. As we can see, the bounds are useful indicators of performance at high load. The bounds are a bit pessimistic at lower loads as the bounds are calculated making no assumptions on the arrival process; with more assumptions, they can be improved.

Parley improves over EyeQ because it can control tail latency even when services A and B are given work-conserving allocations, and therefore could congest the network fabric. In this scenario, these services will saturate the rackswitch uplink, causing overload at congestion point (iii) in Figure 2. Since EyeQ assumes a congestion-free core it cannot control tail latency in such scenarios; instead, EyeQ will default to a TCP-like behavior as highlighted in Table 3.

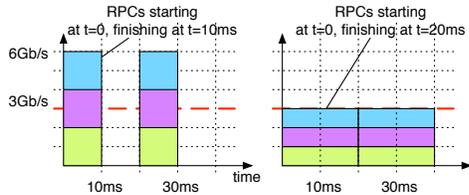


Figure 15. To rate limit or not: Rate limiting and accurate pacing improves sharing but increases RPC completion times.

In EyeQ, the only solution that would control tail latency would be to do admission control on service instances, ensuring that the rack uplink bandwidth is never over-committed. In other words, EyeQ could not support work-conserving allocations, and instead would have to cap (for example) each of the ten instances of both services to 3Gb/s, so as to limit the total allocation to 60Gb/s and prevent any congestion. Such static allocations, however, could waste a large fraction of the scarce uplink bandwidth if either service is inactive or lightly loaded, while the other is highly loaded. In contrast to EyeQ, Parley dynamically adjusts each service’s allocation, thus maximizing resource utilization under the constraint of limited tail latency.

Remark: Table 3 shows results averaged over 3 runs. The standard deviation was under 0.3ms, except (a) in the highlighted cases, where the means exceed 700ms, the std. dev. is above 100ms, and (b) in the case marked ‘*,’ where during one run, two (out of ten) rack brokers in the receiving rack crashed. In the latter case, the remaining eight rack brokers continued to function and under-estimated the total ingress utilization, which resulted in excessive 99th percentile latency.² We have not yet re-run this experiment with 3 bug-free runs.

7. Discussion

Timescales for traffic shaping: As the authors of EyeQ noted [14, §2.2], rate control must be fast enough to react before queues build up and increase latency for competing flows. However, we also found that such fast and accurate rate control can also hurt the completion times of RPC transfers. This insight is not specific to Parley, but applies to any bandwidth management scheme targeted at an environment with relatively high bandwidth and low latency.

To understand this, consider the scenario in Figure 15, where service receives three 2.5MB RPCs at the same time, and then waits for 10ms; thus, its demand oscillates between 0 and 6Gb/s (measured over 10ms) every 20ms. Suppose our policy limits the service to 3Gb/s; over a period of 20ms, its usage complies with

²The 99th percentile latency from the three runs were 15.54ms, 16.07ms, and 84.62ms.

the policy. However, over a 10ms interval (e.g., from 0–10ms), the service exceeds its allocation by 3Gb/s.

If we use a fast-acting rate-control mechanism, the service is always precisely rate limited, so demand from 0–10ms spills over to the next interval. This means that all RPCs take 20ms to finish. If, however, the service is not rate-limited at all, all of its RPCs will finish in 10ms, thus halving the maximum completion time. This highlights a fundamental tradeoff between accurate rate guarantees and RPC completion times. However, it is difficult to upper bound the rate-limiting burst size, as the example in Figure 15 still holds if RPC sizes (and timescales) are scaled up by a factor of 10.

Thus, at the machine shaper, the service rate limiters should have a configurable burst size for coarse-grained rate limiting, and rate meters should have a configurable time interval to measure ingress utilizations. In general, the burst size of a latency-sensitive service should be higher than a collocated throughput-sensitive service. This burst size approach is a special case of the more general *fair service curves* [26], an approach to decouple rate guarantees from delay guarantees. A rule of thumb is to set burst sizes larger than the size of RPCs that have low-latency requirements.

Tighter latency bounds: The latency bound shown in Equation 2 depends both on the capacity C and on the total load ρ on the network queues, and increases steeply as load increases (i.e. as $\rho \rightarrow 1$). Thus, if the resulting latency bounds are unacceptable, either (a) low-latency RPCs must not be exposed to load offered by other services, which can be achieved by prioritization, or (b) network capacity should be increased. Parley helps in two ways: (i) it can control the peak utilization at the rack uplinks/downlinks regardless of the number of service endpoints under the rack; (ii) Parley serves as a useful monitoring and protection mechanism for buggy services that might consume more bandwidth than what is specified according to a policy.

8. Conclusion

We presented Parley, a framework to flexibly and predictably share datacenter bandwidth across services at the machine, rack, and fabric levels. We showed how to systematically decompose a sharing objective into rack-local and machine-local objectives. On a cluster of 90 machines, we demonstrated that Parley is able to meet its goals. We also outlined how predictable bandwidth shares can help services maintain low tail latencies, despite potentially malicious traffic patterns of collocated jobs.

References

- [1] Hitesh Ballani, Keon Jang, Thomas Karagiannis, Changhoon Kim, Dinan Gunawardena, and Greg

- O'Shea. "Chatty Tenants and the Cloud Network Sharing Problem". In: *USENIX NSDI* (2013).
- [2] Hitesh Ballani, Paolo Costa, Thomas Karagiannis, and Antony Rowstron. "Towards Predictable Datacenter Networks". In: *SIGCOMM* (2011).
- [3] Theophilus Benson, Aditya Akella, and David A Maltz. "Network Traffic Characteristics of Data Centers in the Wild". In: *IMC* (2010).
- [4] Dimitri P Bertsekas, Robert G Gallager, and Pierre Humblet. *Data networks*. Prentice-Hall International, 1992.
- [5] Arka A Bhattacharya, David Culler, Eric Friedman, Ali Ghodsi, Scott Shenker, and Ion Stoica. "Hierarchical scheduling for diverse datacenter workloads". In: *SoCC* (2013).
- [6] Cisco: *Hierarchical Queueing Framework (HQF)*. http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/qos_frhqf_support.html. Retrieved October 4, 2013.
- [7] Rene L Cruz. "A calculus for network delay. I: Network elements in isolation". In: *IEEE Transactions on Information Theory* (1991).
- [8] Andrew D Ferguson, Peter Bodik, Srikanth Kandula, Eric Boutin, and Rodrigo Fonseca. "Jockey: Guaranteed Job Latency in Data Parallel Clusters". In: *EuroSys* (2012).
- [9] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. "Dominant Resource Fairness: Fair Allocation of Multiple Resource Types". In: *NSDI* (2011).
- [10] Albert Greenberg, James R Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A Maltz, Parveen Patel, and Sudipta Sengupta. "VL2: a scalable and flexible data center network". In: *SIGCOMM* (2009).
- [11] Nikhil Handigol, Brandon Heller, Vimalkumar Jeyakumar, Bob Lantz, and Nick McKeown. "Reproducible Network Experiments Using Container-based Emulation". In: *CoNEXT* (2012).
- [12] Peter G. Harrison. "Response time distributions in queueing network models". In: *Performance Evaluation of Computer and Communication Systems*. Ed. by Lorenzo Donatiello and Randolph Nelson. Vol. 729. LNCS. Springer Berlin Heidelberg.
- [13] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, et al. "B4: Experience with a Globally-Deployed Software Defined WAN". In: *SIGCOMM* (2013).
- [14] Vimalkumar Jeyakumar, Mohammad Alizadeh, David Mazieres, Balaji Prabhakar, Changhoon Kim, and Albert Greenberg. "EyeQ: Practical network performance isolation at the edge". In: *USENIX NSDI* (2013).
- [15] Vimalkumar Jeyakumar, Abdul Kabbani, Jeffrey C. Mogul, and Amin Vahdat. *Technical Report: Flexible Bandwidth and Latency Provisioning in a Datacenter*. <http://arxiv.org/abs/1405.0631>.
- [16] Jeongkeun Lee, Myungjin Lee, Lucian Popa, Yoshio Turner, Sujata Banerjee, Puneet Sharma, and Bryan Stephenson. "CloudMirror: Application-Aware Bandwidth Reservations in the Cloud". In: *USENIX HotCloud* (2013).
- [17] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. "OpenFlow: enabling innovation in campus networks". In: *SIGCOMM* (2008).
- [18] Jeffrey C Mogul and Lucian Popa. "What We Talk About When We Talk About Cloud Network Performance". In: *SIGCOMM* (2012).
- [19] Daniel Pérez Palomar and Mung Chiang. "A Tutorial on Decomposition Methods for Network Utility Maximization". In: *IEEE Journal on Selected Areas in Communications* (2006).
- [20] Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C Mogul, and Yoshio Turner Jose Renato Santos. "ElasticSwitch: Practical Work-Conserving Bandwidth Guarantees for Cloud Computing". In: *SIGCOMM* (2013).
- [21] Lucian Popa, Gautam Kumar, Mosharaf Chowdhury, Arvind Krishnamurthy, Sylvia Ratnasamy, and Ion Stoica. "FairCloud: Sharing The Network In Cloud Computing". In: *SIGCOMM* (2012).
- [22] Barath Raghavan, Kashi Vishwanath, Sriram Ramabhadran, Kenneth Yocum, and Alex C Snoeren. "Cloud Control with Distributed Rate Limiting". In: *SIGCOMM* (2007).
- [23] Henrique Rodrigues, Jose Renato Santos, Yoshio Turner, Paolo Soares, and Dorgival Guedes. "Gatekeeper: Supporting bandwidth guarantees for multi-tenant datacenter networks". In: *USENIX WIOV* (2011).
- [24] Alan Shieh, Srikanth Kandula, Albert Greenberg, Changhoon Kim, and Bikas Saha. "Sharing the Data Center Network". In: *USENIX NSDI* (2011).
- [25] David Shue, Michael J Freedman, and Anees Shaikh. "Performance Isolation and Fairness for Multi-Tenant Cloud Storage". In: *OSDI* (2012).
- [26] Ion Stoica, Hui Zhang, and T. S. Eugene Ng. "A Hierarchical Fair Service Curve Algorithm for Link-Sharing, Real-Time and Priority Services". In: *SIGCOMM* (1997).
- [27] Di Xie, Ning Ding, Y Charlie Hu, and Ramana Kompella. "The Only Constant is Change: Incorporating Time-Varying Network Reservations in Data Centers". In: *SIGCOMM* (2012).