

# Tapping into the Router's Unutilized Processing Power

Marat Radan and Isaac Keslassy

**Abstract**—The growing demand for network programmability has led to the introduction of complex packet processing features that are increasingly hard to provide at full line rates.

In this paper, we introduce a novel load-balancing approach that provides more processing power to congested linecards by tapping into the processing power of underutilized linecards. Using different switch-fabric models, we introduce algorithms that aim at minimizing the total average delay and maximizing the capacity region. Our simulations with real-life traces then confirm that our algorithms outperform current algorithms as well as simple alternative load-balancing algorithms. Finally, we discuss the implementation issues involved in this new way of sharing the router processing power.

## I. INTRODUCTION

### A. Motivation

The emerging and overwhelming trend towards a *software-based approach to networks* (through network function virtualization, software-defined networking, and expanded use of virtual machines) increasingly lets the network managers configure the packet processing functions needed in their network. For instance, one network manager may want to alter video packets by inserting ads or changing the streaming rate, while another manager may want to filter packets based on a specific key within the packet payload.

This growing demand for *network programmability* has already forced several vendors to integrate full 7-layer processing within switches and routers. For example, EZChip has recently introduced NPS (Network Processor for Smart networks), a 400Gbps C-programmable NPU (Network Processing Unit) [1]. Cisco also promotes its DevNet developer effort for ISR (Integrated Services Router) [2], and already put Akamai's software onto its ISR-AX branch routers.

The processing complexity of the programable features defined by the network managers is hard to predict, and can also widely vary from packet to packet. These features can easily go beyond standard heavy packet processing tasks such as payload encryption [3], intrusion detection [4], and worm signature generation [5], [6], which already have a complexity that can be two orders of magnitude higher than packet forwarding [7]. Technology improvements can help in part, in particular through the increase in CMOS technology density and the growth of parallel architectures [8]. Still, these new features clearly make it hard for vendors to provide a guarantee about the packet line rates at which their processors will process all packets.

M. Radan and I. Keslassy are with the Department of Electrical Engineering, Technion, Israel (e-mails: {radan@tx, isaac@ee}.technion.ac.il).

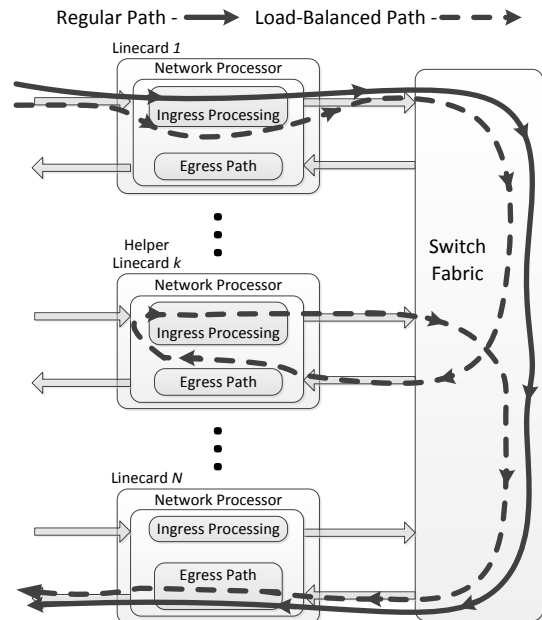


Fig. 1: Generic switch fabric with  $N$  linecards. The load-balanced path allows each packet to be processed at a linecard different from its source linecard, enabling the source linecard to tap into the router's unutilized processing power.

The goal of this paper is to help vendors introduce a new architectural technique that would enable routers to better use their processing resources to deal with these emerging programmable features. We do so by suggesting a *novel multiplexing approach that better utilizes the unused processing power of the router*.

Specifically, Figure 1 illustrates a router with  $N$  linecards, each containing an input and an output. Assume that linecard 1 experiences heavy congestion because its incoming packets require too much processing. Then its queues will fill up, and eventually it will drop packets. Meanwhile, the other router linecards could be in a near-idle state, with their processing power left untapped. This additional processing power would be invaluable to the congested linecard.

*The main idea of this paper is to tap into the unused processing power of the idle router linecards. We do so by load-balancing the packets among the linecards. As illustrated in the figure, suppose linecard 1 is heavily congested while linecard  $k$  is idle. Then a packet arriving at linecard 1 could be sent by linecard 1 to linecard  $k$ , where it would be processed, and later switched to its output destination. Using this technique, linecards with congested ingress paths can send workload to other linecards through the switch fabric. By*

doing so, the available processing power for the traffic arriving to a single linecard would potentially expand up to the entire router's processing power, depending on the other arrival rates and on the switch fabric capacity.

Of course, many challenges exist to implement this idea. For instance, the load-balancing mechanism may require an update mechanism to collect and distribute the information on the congestion of each linecard, together with an optimization algorithm that would take this information and decide when to load-balance packets. There are also additional considerations such as potential packet reordering and a more complex buffer management. We discuss these issues later in the paper.

## B. Contributions

*The main contribution of this paper is a novel technique for load-balancing packet processing between router linecards to achieve better processing performance.*

Our goal is to expand the capacity region of the router by load-balancing and diverting traffic away from congested linecards. However, load-balancing more and more packets may also lead to a congestion of the switch fabric. Therefore, we need to trade off between reducing the *processing load* at the linecards, and reducing the *switching load* at the switch fabric. To help us develop some intuition about the impact of the switch fabric architecture on this trade-off, we present three increasingly-complex queueing models of switch-fabric architectures:

*At first, we neglect the impact of the switch fabric architecture, by assuming that the switch fabric has infinite capacity. We demonstrate that to minimize the total average delay when linecards have equal service rates, the load-balancing algorithm should assign the same processing load to all linecards.*

In addition, we want to minimize the overhead associated with redirecting traffic by minimizing the number of redirection linecard pairs between sending congested linecards and receiving uncongested ones. We prove this to be NP-hard by presenting a reduction from the subset sum problem. We also provide an efficient algorithm that approximates the optimal solution within a factor of two.

*Second, we model a shared-memory switch fabric with finite switching bandwidth. We model the entire switch fabric as a single queue with a finite service rate, representing the finite rate of memory accesses to the shared memory. Then, we prove the Karush-Kuhn-Tucker (KKT) conditions for optimal total average delay in the router. We also present an efficient algorithm to achieve such an optimal solution, and bound its complexity. We further show that our load-balancing scheme significantly increases the capacity region of the router.*

*Third, we provide the most advanced algorithm, given an output-queued switch fabric model. In this model, each switch-fabric output has its own queue, which packets can enter when the linecard is either their helper (middle) linecard or their egress (destination) linecard. We build on previous results to determine the conditions for minimal delay, and quantify the complexity used to achieve the optimal solution.*

In the experimental section, we compare our algorithms to existing load-balancing algorithms using more realistic settings, including OC-192 Internet backbone link traces [9], and an input-queued switch fabric based on VOQs (virtual output queues) with an iSLIP switch-fabric scheduling algorithm [10]. The congestion values of the different linecards are updated at fixed intervals, and packet processing times are based on [7]. We confirm that our algorithms significantly extend the capacity region. In addition, they achieve a lower delay than existing algorithms and than a simple Valiant-based load-balancing algorithm.

Finally, we discuss the different implementation concerns, e.g. buffer management, packet reordering and algorithm overheads. Our suggested schemes are shown to be feasible and able to accommodate various features using only limited modifications, if any, to the existing implementations.

## C. Related Work

There is a large body of literature on improving the performance of network processors. Most recent works have focused on harnessing the benefits of multi-core architectures for network processing [11]–[17]. However, these approaches are local to the processors, and therefore orthogonal to ours. They can be combined for improved performance.

Other efforts have attempted to optimize the switch fabric architecture, e.g. using Valiant-based load-balanced routers [18]–[20]. These efforts have focused on providing a guaranteed switching rate in the switch fabric, while our goal is to optimize the processing capacity of the linecards. Still, in this paper, we also introduce Valiant-like load-balancing algorithms, and show that they often prove to be sub-optimal.

Building on these load-balanced routers, [21] has presented the GreenRouter, an energy-efficient load-balanced router architecture. Instead of load-balancing traffic to all middle linecards, the GreenRouter only load-balances traffic to a subset of the middle linecards, enabling it to shut down idle processing power at other linecards to save power. Therefore, the GreenRouter always load-balances traffic, while we only need to do so in case of congestion. Also, it tries to concentrate traffic into as few processing units as possible, while our goal is to spread it across all the linecards to increase the capacity region and reduce delay. The architecture is also partly modelled by the sub-optimal Valiant-based algorithm that we introduce in the paper, when all linecards are used. Of course, if needed, the energy-efficient goal could be incorporated into our load-balancing scheme as well.

[22] suggested distributing the FIB (Forwarding Information Base) among linecards to support more content prefixes in an ICN (Information-Centric Networking) router. They use static hashing, so only one linecard can provide a given prefix, essentially forcing a Valiant-like load-balancing. On the contrary, all our linecards can process any packet, allowing for more flexibility.

Finally, [23] famously presents the fundamental power of choice in a supermarket model. It shows how choosing the least loaded out of 2 random servers significantly reduces the expected queueing time when compared to a uniform random




Fig. 2: Notation used to describe the traffic within a linecard.

choice. However, in our case, there is no centralized scheduler. Also, after packets arrive at their respective ingress queue, they need to use the switch fabric to be assigned to another queue. Thus, we need to deal with a cost of switch fabric congestion that does not exist in the fundamental model.

## II. MODEL AND PROBLEM FORMULATION

### A. Linecard Model

We start by introducing notations and formally defining the problem. We consider a router with  $N$  linecards connected to a switch fabric. Each linecard is associated with an independent exogenous Poisson arrival process. We denote by  $\gamma_i$  the arrival rate at linecard  $i$ , and by  $\gamma_{i,j}$  the arrival rate at linecard  $i$  of traffic destined to egress  $j$ .

To load-balance processing tasks, we assume that each linecard  $i$  may choose to redirect a fixed portion  $p_{i,j}$  of its incoming traffic to be processed at linecard  $j$ . Thus, each incoming packet is independently redirected to linecard  $j$  with probability  $p_{i,j}$ , with  $\sum_{j=1}^N p_{i,j} = 1$ . For instance, if linecard  $i$  is not redirecting traffic, then  $p_{i,i} = 1$ .

As illustrated in Figure 2, we denote by  $R_i^-$  the redirected traffic away from linecard  $i$ , such that  $R_i^- = \sum_{j \neq i} p_{i,j} \gamma_i$ .

Similarly, we denote by  $R_i^+$  the redirected traffic into linecard  $i$ , such that  $R_i^+ = \sum_{j \neq i} p_{j,i} \gamma_j$ . Note that a packet can be redirected at most once, immediately upon its arrival.

In addition, we denote by  $\lambda_i$  the total effective arrival rate into the processing queue of linecard  $i$ , i.e.

$$\lambda_i = \sum_{j=1}^N p_{j,i} \gamma_j = \gamma_i + R_i^+ - R_i^-. \quad (1)$$

Finally, we assume that the linecard processing time of each packet is exponentially distributed with parameter  $\mu_i$ .

**Example 1.** Assume  $N = 2$  linecards, with equal service rate  $\mu_1 = \mu_2 = 1$ . Further assume that the first linecard is congested with an arrival rate of  $\gamma_1 = 1.2$  and the second linecard is idle, i.e.  $\gamma_2 = 0$ . Then the first linecard could redirect half of its incoming traffic to the second linecard, namely  $p_{1,2} = 0.5$ . Therefore  $R_1^- = R_2^+ = 0.6$ , and the effective arrival rate into each linecard queue is  $\lambda_1 = \lambda_2 = 0.6$ .

### B. Switch-Fabric Model

We now propose three different models for the switch-fabric architecture and the resulting switch-fabric congestion.

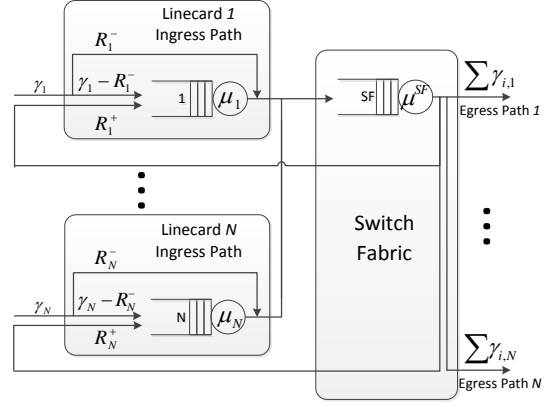


Fig. 3: The single-queue switch-fabric model of a shared-memory router.

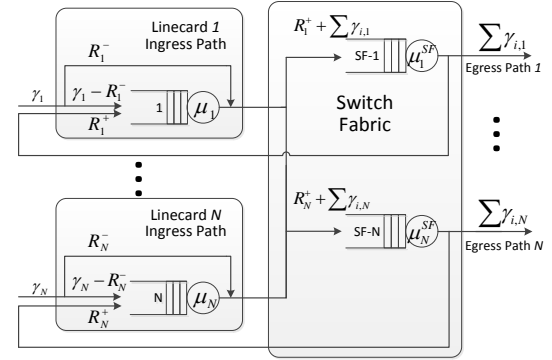


Fig. 4: The  $N$ -queues switch fabric model of an output-queued router. The  $N$  output queues represent the switch-fabric delay that depends only on the packet destination.

In the *first model*, we assume that the switch fabric has an infinite switching capacity, such that packets pass through it without delay. Therefore, the load-balancing algorithm will not take switch congestion into consideration. This model provides us with some intuition on the pure load-balancing approach.

As illustrated in Figure 3, the *second model* represents a shared-memory switch fabric that relies on a single shared queue. We model the service time at this shared switch-fabric queue as exponentially distributed with parameter  $\mu^{SF}$ , representing the switching capacity of the switch fabric. The total arrival rate to this queue is  $\lambda^{SF} = \sum_{j=1}^N \gamma_j + \sum_{i=1}^N R_i^+$ , i.e. the sum of the total incoming traffic rate into the router and the additional rate of redirected packets.

Figure 4 shows the *third model*, which represents an output-queued switch fabric as a set of  $N$  queues, each queue being associated with a different switch-fabric output. We denote by  $\mu_i^{SF}$  the service rate for output queue  $i$ . Its arrival rate is  $\lambda_i^{SF} = \sum_{j=1}^N \gamma_{j,i} + R_i^+$ , where  $\sum_{j=1}^N \gamma_{j,i}$  is the total arrival rate destined towards linecard  $i$ .

In all three switch-fabric models, the stationary distribution of the queue sizes can be modeled as following a *product-form distribution*. This can be seen by directly applying either Kelly's results on general customer routes in open queueing networks (Corollary 3.4 of [24]), or by applying the BCMP theorem, with each flow being a chain, and each redirected

flow being denoted as a new class after its first passage through the switch fabric (Fig. 3 of [25]).

### C. Delay Optimization Problem

Given a router with linecard arrival rates  $\{\gamma_{i,j}\}$ , our goal is to find the fixed redirection probabilities  $p_{i,j}$  that minimize the average packet delay. Whenever defined, we assume that we are in the steady state. We also assume that all the buffers have infinite size and FIFO policy.

Note that by minimizing the average delay, we also maximize the capacity region, the group of feasible arrival vectors with finite expected delay, because outside the feasible region the steady-state average delay is infinite by definition.

## III. INFINITE-CAPACITY SWITCH-FABRIC MODEL

### A. Minimizing the Average Delay

To gain some intuition on the delay optimization problem, we start by studying the first model where the switch-fabric capacity is infinite.

Our first result shows that to minimize the total average delay, the load-balancing algorithm should achieve equal weighted delay derivative for all linecards.

**Theorem 1.** *In the infinite-capacity switch-fabric model with feasible incoming rates (i.e.  $\sum_{i=1}^N \gamma_i < N \cdot \mu$ ), the load-balancing algorithm is optimal for the average delay iff*

$$\forall i, j \in [1, N], \frac{\mu_i}{(\mu_i - \lambda_i)^2} = \frac{\mu_j}{(\mu_j - \lambda_j)^2} \quad (2)$$

*Proof:* In steady state, we can combine Little's law and the above-mentioned product-form distribution for the stationary distribution of all the queue sizes. As a result, the average delay for linecard  $i$  is the same delay as for an M/M/1 queue with a Poisson arrival rate of  $\lambda_i$  and a service rate of  $\mu_i$ . The delay-minimization problem can thus be formulated as:

$$\begin{aligned} \text{minimize } \bar{D}(\{\lambda_i\}) &= \frac{1}{\sum_{i=1}^N \gamma_i} \cdot \sum_{i=1}^N \frac{\lambda_i}{\mu_i - \lambda_i} \\ \text{subject to } \sum_{i=1}^N \lambda_i &= \sum_{i=1}^N \gamma_i \\ \lambda_i - \mu_i &< 0, \forall i \in [1, N] \\ -\lambda_i &\leq 0, \forall i \in [1, N]. \end{aligned}$$

This is a convex optimization problem, and therefore we will find the Karush-Kuhn-Tucker (KKT) conditions. The KKT conditions are the first-order necessary conditions. Since our goal function and region are convex, the KKT conditions are also sufficient for a solution to be optimal.

The first constraint,  $\sum_{i=1}^N \lambda_i = \sum_{i=1}^N \gamma_i$ , means that all the traffic that enters the router is also processed by one of the linecards in the router. The second constraint,  $\lambda_i - \mu_i < 0$ , requires that the arrival rates into the linecards' processing queues be lower than their service rates. The third constraint,  $-\lambda_i \leq 0$ , simply means that the queue arrival rates must be non-negative.

The relevant KKT conditions for each  $j \in [1, N]$  are:

$$\begin{aligned} -\frac{\partial}{\partial \lambda_j} (\bar{D}(\Lambda)) &= \sum_{i=1}^N \eta_i \frac{\partial}{\partial \lambda_j} (-\lambda_i) \\ &+ \sum_{i=N+1}^{2N} \eta_i \frac{\partial}{\partial \lambda_j} (\lambda_{i-N} - \mu_{i-N}) \\ &+ \tau_0 \frac{\partial}{\partial \lambda_j} \left( \sum_{i=1}^N \lambda_i - \sum_{i=1}^N \gamma_i \right) \\ \eta_i &\geq 0, \forall i \in [1, 2N] \\ \eta_i (-\lambda_i) &= 0, \forall i \in [1, N] \\ \eta_i (\lambda_{i-N} - \mu_{i-N}) &= 0, \forall i \in [N+1, 2N] \end{aligned}$$

The solution must be within the feasible region, with  $\mu_i > \lambda_i, \forall i \in \{1, \dots, N\}$ , therefore  $\eta_i = 0, \forall i \in [N+1, 2N]$ . Developing the conditions for a certain index  $j \in [1, N]$  shows that if  $\lambda_j > 0$  then  $\eta_j = 0$ :

$$\frac{\mu_j}{(\mu_j - \lambda_j)^2} = -\tau_0 \sum_{i=1}^N \gamma_i \quad (3)$$

Otherwise  $\lambda_j = 0$ :

$$-\frac{1}{\sum_{i=1}^N \gamma_i} \frac{1}{\mu_j} = -\eta_j + \tau_0 \frac{\eta_j \geq 0}{\mu_j} \geq -\tau_0 \sum_{i=1}^N \gamma_i \quad (4)$$

Eq. (3) shows that all of the linecards with incoming traffic have the same delay derivative, while in Eq. (4) linecards with slow service rate obtain no incoming traffic. ■

As expected, since in this model there is no cost to load-balancing, a corollary is that when all processing rates are equal, each linecard should obtain the same amount of traffic to process.

**Corollary 2.** *In the infinite-capacity switch-fabric model with feasible incoming rates, if all linecards are identical, i.e.  $\forall i \in [1, N] : \mu_i = \mu$ , then they should have the same rate of incoming traffic to process, i.e.  $\forall i \in [1, N] : \lambda_i = \mu$ .*

*Proof:* Eq. (3) holds for all linecards, thus  $\forall j \in [1, N]$ :

$$\lambda_j = \mu - \sqrt{-\frac{\mu}{\tau_0 \sum_{i=1}^N \gamma_i}} = \text{Constant.}$$

In other words, as expected, when all linecards are identical, they have the same rate of incoming traffic to process:  $\forall i \in [1, N], \lambda_i = \lambda_{\text{avg}} = \left( \sum_{i=1}^N \frac{\gamma_i}{N} \right)$ . ■

**Load-balancing algorithm.** A possible algorithm to achieve these conditions is quite simple. Linecards with arrival rates above  $\lambda_{\text{avg}}$  load-balance their excess arrival rates to helper linecards. To do so, they simply pick each helper linecard proportionally to its capacity to help, i.e. to the amount of traffic that it would need to reach  $\lambda_{\text{avg}}$ . Formally, each linecard  $i$  sends  $R_i^- = \max(0, \gamma_i - \lambda_{\text{avg}})$  and receives  $R_i^+ = \max(0, \lambda_{\text{avg}} - \gamma_i)$ , using load-balancing probability

$$p_{i,j \neq i} = \frac{R_i^-}{\gamma_i} \cdot \frac{R_j^+}{\sum_{j=1}^N R_j^+}.$$

**Example 2.** Suppose we have an infinite-capacity switch with arrival rates  $\underline{\gamma} = (1, 2, 2, 3.5, 4.5, 7, 8)$ , and equal service rates  $\mu_i = \mu = 5$ . Then any algorithm that achieves  $\lambda_i = 4$  for all  $i$  is optimal. In particular, the above algorithm yields  $\underline{R}^- = (0, 0, 0, 0, 0.5, 3, 4)$  and  $\underline{R}^+ = (3, 2, 2, 0.5, 0, 0, 0)$ .

### B. Linecard Pairing Problem

We have several degrees of freedom when assigning values to the  $N^2$  variables  $p_{i,j}$ : for each linecard, the total amount of traffic it receives and redirects is defined using the above optimality equations, but these equations do not determine how to split the redirected traffic among linecards. Unfortunately, redirecting traffic from one linecard to another would require additional overhead to synchronize variables and allocate resources. Therefore, we will now study how to *minimize the number of linecard pairings*.

We first assume that if a linecard is receiving redirected traffic, then it must not redirect traffic, and vice versa. Intuitively, this is because we would like to minimize the redirected flow in the system in order to avoid unnecessarily congesting the switch fabric, since realistically its capacity is finite.

As a result, the total redirected traffic towards other linecards must be equal to:

$$\forall i \in [1, N], \begin{cases} R_i^- = \sum_{j=1, i \neq j}^N p_{i,j} \gamma_i = \gamma_i - \lambda_{\text{avg}}, & \gamma_i > \lambda_{\text{avg}} \\ R_i^- = \sum_{j=1, i \neq j}^N p_{i,j} = 0, & \gamma_i \leq \lambda_{\text{avg}} \end{cases}$$

and likewise, the total incoming redirected traffic from other linecards must be equal to:

$$\forall i \in [1, N], \begin{cases} R_i^+ = \sum_{j=1, i \neq j}^N p_{j,i} \gamma_j = \lambda_{\text{avg}} - \gamma_i, & \gamma_i < \lambda_{\text{avg}} \\ R_i^+ = \sum_{j=1, i \neq j}^N p_{j,i} \gamma_j = 0, & \gamma_i \geq \lambda_{\text{avg}} \end{cases}$$

**Definition 1** (Redirection Matrix). A redirection matrix  $R$  is a matrix of  $N$  rows and  $N$  columns.

$$R = \begin{pmatrix} r_{1,1} & r_{1,2} & \cdots & r_{1,N} \\ r_{2,1} & r_{2,2} & \cdots & r_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ r_{N,1} & r_{N,2} & \cdots & r_{N,N} \end{pmatrix}$$

Each element  $r_{i,j}$  represents the rate at which linecard  $i$  is redirecting traffic to linecard  $j$ .

$$r_{i,j} = \begin{cases} p_{i,j} \gamma_i, & i \neq j \\ 0, & i = j \end{cases}$$

$R_i^+$ , the incoming redirected traffic rate into linecard  $i$ , could now also be defined as the sum on column  $i$  of  $R$ ,  $\sum_{j=1}^n r_{j,i}$ . Similarly,  $R_i^-$  is the sum on row  $i$  of  $R$ ,  $\sum_{j=1}^n r_{i,j}$ . Therefore the optimization only dictates the sums of the rows and columns of the redirection matrix, and allows freedom when assigning individual elements with values.

**Definition 2** (LPP). The Linecard Pairing Problem attempts to minimize the number of pairings. Formally, define an indicator function:

$$1_{i,j} := \begin{cases} 1, & \text{if } r_{i,j} > 0, \\ 0, & \text{else.} \end{cases}$$

Given the arrival rates into the linecards, we want to find the redirection probabilities that minimize  $P = \sum_{i=1}^N \sum_{j=1}^N 1_{i,j}$ , the number of non-zero values in matrix  $R$ .

For instance, one simple and ineffective assignment is the *Proportional Distribution Pairing Algorithm* (PDPA). In PDPA, each linecard  $j$  with below-average flow receives a portion of the excess flow from linecard  $i$  with above-average flow. The portion size is based on linecard  $j$ 's contribution to the total redirected flow in the system. This achieves average flow on all linecards. The redirection probabilities are calculated using the following formula:

$$\forall i, j \in [1, N], \gamma_i > \lambda_{\text{avg}}, \gamma_j < \lambda_{\text{avg}},$$

$$p_{i,j} = \frac{\gamma_i - \lambda_{\text{avg}}}{\gamma_i} \frac{\lambda_{\text{avg}} - \gamma_j}{\sum_{k=1, \gamma_k < \lambda_{\text{avg}}}^N (\lambda_{\text{avg}} - \gamma_k)}.$$

From a practical point of view, the PDPA assignment is ineffective due to the overhead associated with pairing linecards. The overhead may consist of resource allocation for the redirected traffic, or synchronizing state variables such as patterns for security applications. A better approach in those cases would be to send the redirected flow from each linecard to the lowest possible number of linecards.

**Example 3.** Let's demonstrate PDPA for the system that appears in Example 2. We obtain the following traffic redirection matrix for PDPA:

$$R_{PDPA} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.200 & 0.133 & 0.133 & 0.033 & 0 & 0 & 0 \\ 1.200 & 0.800 & 0.800 & 0.200 & 0 & 0 & 0 \\ 1.600 & 1.066 & 1.066 & 0.266 & 0 & 0 & 0 \end{pmatrix}$$

This assignment could be also represented as:

$$R_{PDPA} = \frac{1}{\sum_i R_i^-} \cdot R^- (R^+)^T = \frac{1}{7.5} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0.5 \\ 3.0 \\ 4.0 \end{pmatrix} \begin{pmatrix} 3.0 \\ 2.0 \\ 2.0 \\ 0.5 \\ 0 \\ 0 \\ 0 \end{pmatrix}^T$$

The pairings are represented by the non-zero values. Therefore, in total,  $P_{PDPA} = 12$ .

However, in this example, we can find that the optimal solution yields  $P_{OPT} = 4$ , using:

$$R_{OPT} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0 & 0 \\ 3.0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2.0 & 2.0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Specifically, linecard 8 redirects to linecards 2 and 3, linecard 7 only redirects to linecard 1, and linecard 6 only redirects to linecard 5. This is optimal because each linecard with a below-average flow needs to be paired at least once, and there are 4 linecards with below-average flow in this example.

In fact, in the general case, we are able to prove the following theorem about LPP, showing that there is no easy solution to the problem.

**Theorem 3.** *LPP is NP-hard.*

*Proof:* If we had an efficient solution to LPP, we would have an efficient solution to the subset sum problem, known to be NP-complete. In the subset sum problem, we are given a set of integers  $S$  and an integer  $k$ . The problem is to find a non-empty subset of  $S$  which sums up to  $k$ .

In order to reduce the subset sum problem to LPP the set of integers  $S$  would represent the size of the overflow in the overflow linecards, the integer  $k$  would be the underflow in one linecard and  $\left(\sum_{i \in S} s_i - k\right)$  would be the underflow in another linecard.

If the algorithm returns an assignment with  $\text{size}(S)$  pairs it means that there is a non-empty subset which sums up to  $k$ , the overflow values of the linecards paired with the  $k$ -value underflow linecard. Otherwise, if the algorithm returns an assignment with  $\text{size}(S) + 1$  pairs, it means that there is not a non-empty subset which sums up to  $k$ . Therefore LPP is an NP-hard problem.

This reduction is illustrated in Figure 5,  $\lambda_{\text{avg}}$  and  $\mu$  are chosen to be large enough to accommodate an underflow of size  $\max\left(\sum_{i \in S} s_i - k, k\right)$  and an overflow of  $\max_{s_i \in S}(s_i)$  ■

We now introduce a simple heuristic algorithm, which we denote (SLPA (Simple Linecard Pairing Algorithm)), to approximate the solution of LPP. SLPA randomly chooses a congested linecard, and iteratively pairs it with uncongested linecards, redirecting as much traffic as possible, until it has reached the desired arrival rate.

Algorithm 1 provides the pseudo-code for SLPA.  $R_{\text{minusP}}$  is the index of the linecard with traffic being redirected away from and  $R_{\text{plusP}}$  is the index of the linecard receiving the redirected traffic. The algorithm at each iteration redirects the maximal amount of traffic it can and advances the different counters accordingly.

**Example 4.** Since SLPA processes linecards from  $N$  to 1 regardless of their incoming flow, its performance varies depending on the exact flows. An SLPA algorithm run for the

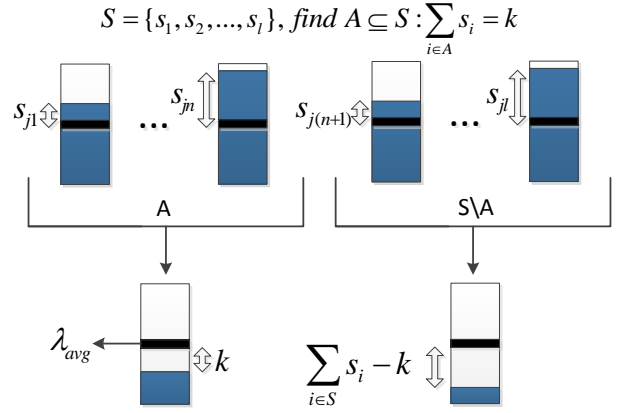


Fig. 5: Subset sum problem reduction to the linecard pairing problem.

---

**Algorithm 1:** SLPA algorithm

---

**Input :** Vectors  $R^-$  and  $R^+$  of length  $N$   
**Output:** Matrix  $R_{\text{SLPA}}$

- 1  $R_{\text{minusP}} = N$ ;
- 2  $R_{\text{plusP}} = 1$ ;
- 3  $R_{\text{SLPA}} = 0_{N,N}$ ;
- 4 **while**  $R_{\text{minusP}} \geq 1$  **do**
- 5     **if**  $R^- [R_{\text{minusP}}] > 0$  **then**
- 6         **if**  $R^- [R_{\text{minusP}}] > R^+ [R_{\text{plusP}}]$  **then**
- 7              $R^- [R_{\text{minusP}}] -= R^+ [R_{\text{plusP}}]$ ;
- 8              $R_{\text{SLPA}} [R_{\text{minusP}}] [R_{\text{plusP}}] = R^+ [R_{\text{plusP}}]$ ;
- 9              $R^+ [R_{\text{plusP}}] = 0$ ;
- 10             $R_{\text{plusP}} += 1$ ;
- 11         **else**
- 12              $R^+ [R_{\text{plusP}}] -= R^- [R_{\text{minusP}}]$ ;
- 13              $R_{\text{SLPA}} [R_{\text{minusP}}] [R_{\text{plusP}}] = R^- [R_{\text{minusP}}]$ ;
- 14              $R^- [R_{\text{minusP}}] = 0$ ;
- 15              $R_{\text{minusP}} -= 1$ ;
- 16         **else**
- 17              $R_{\text{minusP}} -= 1$ ;
- 18 **return**  $R_{\text{SLPA}}$ ;

---

system that appears in Example 2 is as follows: Suppose the algorithm begins with linecard 7 with  $\gamma_7 = 8$ , redirects the most it can to linecard 1 before it reaches the average flow, 3, and remains with an overflow of 1. It then redirects the remainder to linecard 2. Linecard 6 continues the same way and redirects to linecard 2 and 3, and linecard 6 redirects to linecard 5. The resulting redirection matrix would be:

$$R_{\text{SLPA}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0 & 0 \\ 0 & 1.0 & 2.0 & 0 & 0 & 0 & 0 \\ 3.0 & 1.0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

In total  $P_{SLPA} = 5$ . In this example the input vectors were sorted, but this does not help achieve the optimal solution. The worst-case scenario for the SLPA algorithm, in this example a reversed arrival-rate vector, would be  $P_{SLPA} = 6$ , since at each iteration we create one pairing and at least one non-zero value in  $R^+$  and  $R^-$  is zeroed. At the final iteration  $R^+$  and  $R^-$  become zero vectors. There are seven non-zero values in  $R^+$  and  $R^-$  in this example, and therefore there would be at most  $7 - 1 = 6$  iterations for the SLPA algorithm.

**Theorem 4.**  $P_{SLPA} < 2 \cdot P_{OPT}$ .

*Proof:* Let  $M^+$  and  $M^-$  denote the group of linecards with above- and below-average congestion, respectively. The minimum number of pairings is at least  $\max(|M^+|, |M^-|)$  since each linecard in each group needs to be paired at least once.

Using SLPA, the number of linecard pairings is upper-bounded by  $|M^+| + |M^-| - 1$ . This is because at each pairing, either the most congested linecard's rate has been decreased to equal the average or the least congested linecard's rate has been increased to equal the average. At the final pairing, both the linecards' rates are equal to the average and therefore the number of pairings is at most the size of both groups minus one.

Since  $|M^+| + |M^-| - 1 < 2 \cdot \max(M^+, M^-)$ , the upper-bound of SLPA's pairings solution is smaller than twice the lower-bound of the optimal solution. ■

The ratio in this bound is tight. Suppose we have  $N = 2k$  linecards, half of the linecards with above-average flow and half below-average flow. In addition there exists a symmetry of the flows,  $\forall i \in [1, k]$ ,  $\gamma_{k+i} - \lambda_{avg} = \lambda_{avg} - \gamma_{k-i+1}$ , and a single linecard has an overflow equal to the sum of all the other overflows in the system,  $\gamma_N = \lambda_{avg} + \sum_{i \in [k+1, N-1]} (\gamma_i - \lambda_{avg})$ .

$P_{OPT}$  in this case is  $k$ , due to symmetry.  $P_{SLPA}$  could possibly pair all the small underflow linecards with the great overflow linecard, and pair the remaining small overflow linecards with the great underflow linecard, for a total of  $2k - 2$  pairing. Therefore, in this example, the ratio is equal to  $\frac{2k-2}{k}$ , and as there are more linecards in the system the ratio converges to 2.

#### IV. SINGLE-QUEUE SWITCH-FABRIC MODEL

##### A. Minimizing the Average Delay

We now want to start taking into account the congestion at the switch fabric when deciding whether to load-balance traffic to reduce processing congestion. As illustrated in Figure 3, we model a shared-memory switch fabric using a single shared queue. Therefore, using the product-form distribution, the average total delay  $\bar{D}(R^+, R^-)$  of a packet through the router can be expressed as:

$$\bar{D} = \frac{1}{\sum_{i=1}^N \gamma_i} \left[ \left( \sum_{i=1}^N \frac{\overbrace{\lambda_i}^{\text{Linecard Delay}}}{\mu_i - \lambda_i} \right) + \left( \frac{\overbrace{\sum_{i=1}^N (\gamma_i + R_i^+)}^{\text{Switch Fabric Delay}}}{\mu^{\text{SF}} - \sum_{i=1}^N (\gamma_i + R_i^+) } \right) \right]$$

i.e., as the sum of the average delay through the linecards and through the switch fabric, averaged over all the traffic. The linecard delay is simply the delay through a queue with an arrival rate  $\lambda_i$  (from Eq. (1)) and service rate  $\mu_i$ . Likewise, the switch-fabric delay is simply the delay through a queue with an arrival rate  $\sum_i (\gamma_i + R_i^+)$  and service rate  $\mu^{\text{SF}}$ .

As a result, the optimization problem is given by:

$$\begin{aligned} & \text{minimize } \bar{D}(R^+, R^-) \\ & \text{subject to } \sum_{i=1}^N R_i^+ - \sum_{i=1}^N R_i^- = 0 \\ & \sum_{i=1}^N \gamma_i + \sum_{i=1}^N R_i^+ - \mu^{\text{SF}} \leq 0 \\ & \gamma_i + R_i^+ - R_i^- - \mu_i \leq 0, i \in [1, N] \\ & -\gamma_i - R_i^+ + R_i^- \leq 0, -R_i^+ \leq 0, -R_i^- \leq 0, i \in [1, N] \end{aligned}$$

where the first condition expresses flow conservation, the next two conditions signify that the queues are stable, and the last three conditions keep the flows non-negative. Note that formally, we further substitute  $\sum_{i=1}^N R_i^+$  by  $\frac{1}{2} \left( \sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^- \right)$  for symmetry in the expressions of the derivative value, and also replace  $\lambda_i$  in the expression of  $\bar{D}(R^+, R^-)$  by the expression in Eq. (1).

We find that in our load-balancing scheme, an arrival vector  $\underline{\gamma}$  is feasible when both the processing and switching loads are feasible, i.e. (a) the combined service rate of the linecards is greater than the total arrival rate, and (b) the switch-fabric service rate is greater than the total arrival rate combined with the second pass of redirected traffic. Formally,

$$\begin{cases} \sum_{k=1}^N \gamma_k < \sum_{k=1}^N \mu_k \\ \sum_{k=1}^N \min(\gamma_k, \mu_k) + 2 \cdot \sum_{k=1}^N \max(\gamma_k - \mu_k, 0) < \mu^{\text{SF}} \end{cases}$$

We obtain the following result for an optimal load-balancing:

**Theorem 5.** A solution to the average delay minimization problem for the single-queue switch fabric model is optimal if and only if the arrival rate vector is feasible and there exists a constant  $\tau_0$  such that for each linecard  $j$ , if  $R_j^+ > 0$ :

$$\begin{aligned} & \frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = \\ & -\tau_0 \sum_{i=1}^N \gamma_i - \frac{1}{2} \frac{\mu^{\text{SF}}}{\left( \mu^{\text{SF}} - \sum_{i=1}^N \gamma_i - \frac{1}{2} \left( \sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^- \right) \right)^2} \end{aligned} \quad (5)$$

$$\begin{aligned} & \text{If } R_j^- > 0 : \frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = \\ & -\tau_0 \sum_{i=1}^N \gamma_i + \frac{1}{2} \frac{\mu^{\text{SF}}}{\left( \mu^{\text{SF}} - \sum_{i=1}^N \gamma_i - \frac{1}{2} \left( \sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^- \right) \right)^2} \end{aligned} \quad (6)$$

*Proof:* Once again, the Karush-Kuhn-Tucker conditions are the first order necessary conditions, which in our case of convex function and region are also sufficient conditions, for a solution to be optimal.

The KKT multipliers are:  $\eta_i, i \in \{1, \dots, 4N + 1\}, \tau_0$ . There are  $4N + 1$  constants,  $\eta_i$ , one for each inequality constraint of the problem.  $\tau_0$  is used for the single equality constraint. The stationarity equations are for all  $j$ :

$$\begin{aligned}
& - \frac{\partial}{\partial R_j^{+/-}} (\bar{D}(R^+, R^-)) \\
& = \sum_{i=1}^N \eta_i \frac{\partial}{\partial R_j^{+/-}} (-\gamma_i - R_i^+ + R_i^-) \\
& + \sum_{i=N+1}^{2N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (-R_{i-N}^+) \\
& + \sum_{i=2N+1}^{3N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (-R_{i-2N}^-) \\
& + \sum_{i=3N+1}^{4N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (\gamma_{i-4N} \\
& + R_{i-4N}^+ - R_{i-4N}^- - \mu_{i-4N}) \\
& + \eta_i \frac{\partial}{\partial R_j^{+/-}} \left( -\mu^{\text{SF}} + \sum_{k=1}^N \gamma_k + \frac{1}{2} \left( \sum_{k=1}^N R_k^+ + \sum_{k=1}^N R_k^- \right) \right) \\
& + \tau_0 \frac{\partial}{\partial R_j^{+/-}} \left( \sum_{i=1}^N R_i^+ - \sum_{i=1}^N R_i^- \right)
\end{aligned}$$

The dual feasibility and complementary slackness equations are:

$$\begin{aligned}
& \eta_i \geq 0, \forall i \in \{1, \dots, 4N + 1\} \\
& \eta_i (-\gamma_i - R_i^+ + R_i^-) = 0, \forall i \in [1, N] \\
& \eta_i (-R_{i-N}^+) = 0, \forall i \in [N + 1, 2N] \\
& \eta_i (-R_{i-2N}^-) = 0, \forall i \in [2N + 1, 3N] \\
& \eta_i (\gamma_{i-4N} + R_{i-4N}^+ - R_{i-4N}^- - \mu_{i-4N}) = 0, \\
& \quad \quad \quad \forall i \in [3N + 1, 4N] \\
& \eta_{4N+1} \left( -\mu^{\text{SF}} + \sum_{k=1}^N \gamma_k + \frac{1}{2} \left( \sum_{k=1}^N R_k^+ + \sum_{k=1}^N R_k^- \right) \right) = 0
\end{aligned}$$

If the traffic is feasible, the service rates of the linecards and switch fabric are greater than the arrival rate into the linecard queues, thus  $\eta_{3N+j} = 0, \forall j \in [1, N]$ , and  $\eta_{4N+1} = 0$ . For a certain linecard  $j$ , if  $\gamma_i + R_i^+ - R_i^- > 0$ , meaning it is processing packets, then  $\eta_j = 0$ . If  $R_j^+ > 0$  then  $\eta_{N+j} = 0$  and we obtain:

$$\begin{aligned}
& \frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = \\
& - \tau_0 \sum_{i=1}^N \gamma_i - \frac{1}{2} \frac{\mu^{\text{SF}}}{\left( \mu^{\text{SF}} - \sum_{i=1}^N \gamma_i - \frac{1}{2} \left( \sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^- \right) \right)^2}
\end{aligned}$$

Similarly for  $R_j^- > 0$ :

$$\begin{aligned}
& \frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = \\
& - \tau_0 \sum_{i=1}^N \gamma_i + \frac{1}{2} \frac{\mu^{\text{SF}}}{\left( \mu^{\text{SF}} - \sum_{i=1}^N \gamma_i - \frac{1}{2} \left( \sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^- \right) \right)^2}
\end{aligned}$$

Intuitively, we define the (weighted) delay derivative of a linecard  $i$  as  $\frac{\mu_i}{(\mu_i - \lambda_i)^2}$ . Then Theorem 5 states that *the delay derivative of all linecards that receive traffic must equal the same value in Eq. (5)*. Similarly, the delay derivative value of redirecting linecards with excess traffic must equal the same value in Eq. (6). The difference between Eq. (5) and Eq. (6) is the delay derivative of the switch fabric. This is expected because the switch fabric acts as a penalty for redirecting traffic. Also, as expected,  $R_j^+ > 0$  and  $R_j^- > 0$  cannot both be true for a specific linecard  $j$ .

Moreover, an interesting result is that *it is possible for a linecard not to participate in the load-balancing*. Specifically, if for a linecard  $j$ ,  $\frac{\mu_j}{(\mu_j - \gamma_j)^2}$  is between the two values in Eq. (5) and Eq. (6), then  $R_j^+ = 0$  and  $R_j^- = 0$ . Intuitively, the linecard is not congested enough to send traffic and congest the switch fabric, but also not idle enough to gain from receiving traffic that would further congest the switch fabric.

The algorithm needed to achieve the KKT optimality conditions is relatively simple. It orders all linecards by their delay derivative, and *progressively sends more flow from the linecard(s) with the highest delay derivative to the linecard(s) with the lowest*, until it achieves the conditions specified in Theorem 5.

In Algorithm 2 we first initialize the different variables:  $n_i, s_i, r_i$  are the weighted delay derivative values of the linecards that are not participating in the load-balancing, sending linecards, and receiving linecards respectively.  $\Gamma_{\max}$  (resp.  $\Gamma_{\min}$ ) is the set of linecards with the maximal (resp. minimal) weighted delay derivative.  $\Delta$  is the weighted delay derivative value of the switch fabric.  $B_i^{\max}$  and  $B_i^{\min}$  are used to keep the values of the linecards in  $\Gamma_{\max}$  and  $\Gamma_{\min}$  equal at each iteration, respectively.

Then at each iteration of the *while* loop, either a linecard joins one of the groups  $\Gamma_{\max/\min}$  or the algorithm halts. Calculating the distance from  $\Gamma_{\max}$  or  $\Gamma_{\min}$  to the closest linecards is trivial:

$$\begin{aligned}
& \frac{\mu_i}{(\mu_i - \gamma_i - R_i^+ + R_i^-)^2} = \frac{\mu_j}{(\mu_j - \gamma_j)^2} \\
& \frac{\sqrt{\mu_i}}{(\mu_i - \gamma_i - R_i^+ + R_i^-)} = \frac{\sqrt{\mu_j}}{(\mu_j - \gamma_j)} \\
& \frac{\sqrt{\mu_i}}{\sqrt{\mu_j}} (\mu_j - \gamma_j) - (\mu_i - \gamma_i) = -R_i^+ + R_i^-
\end{aligned}$$

Distributing the redirected traffic to the linecards in  $\Gamma_{\max}$  and  $\Gamma_{\min}$  while maintaining equal value is also easily achieved. We begin with  $i, j \in \Gamma_{\max/\min}$ ,  $\frac{\mu_i}{(\mu_i - \gamma_i)^2} = \frac{\mu_j}{(\mu_j - \gamma_j)^2}$ , we need to distribute the redirected traffic such that:

$$\frac{\mu_i}{(\mu_i - \gamma_i + R_i)^2} = \frac{\mu_j}{(\mu_j - \gamma_j + R_j)^2}$$



Linecard	$\lambda=\gamma$	delay derivative	$\lambda$	delay derivative	$\lambda$	delay derivative	$\lambda$	delay derivative
1	1.0	0.31	2.0	0.55	3.5	2.22	3.58	<u>2.64</u>
2	2.0	0.55	2.0	0.55	3.5	2.22	3.58	<u>2.64</u>
3	2.0	0.55	2.0	0.55	3.5	2.22	3.58	<u>2.64</u>
4	3.5	2.22	3.5	2.22	3.5	2.22	3.58	<u>2.64</u>
5	4.5	20.0	4.5	20.0	4.5	20.0	4.5	20.0
6	7.0	—	7.0	—	4.75	80.0	4.58	<u>28.4</u>
7	8.0	—	7.0	—	4.75	80.0	4.58	<u>28.4</u>
	$\Delta=0.714$		$\Delta=0.972$		$\Delta=15.55$		$\Delta=25.94$	

Fig. 6: Successive steps of the algorithm in the single-queue switch-fabric model, based on Example 5. At first, each linecard receives some traffic, and assumes there is no load-balancing. It also computes its resulting processing delay derivative. Then, in step (a), linecard 7 (with an infinite derivative and the highest amount of traffic) load-balances one unit of flow to linecard 1 (which has the lowest derivative). Next, in step (b), saturated linecards 6 and 7 send 4.5 units of flow to linecards 1, 2 and 3. Finally, in step (c), all derivatives are finite. Linecards 6 and 7 send 0.34 additional units of flow to linecards 1, 2 3 and 4. The amount in the final step (c) is found using a binary search, other steps use a simple calculation on the derivatives. The algorithm stops, since the difference between the heavy and light derivatives is  $\Delta$ , the switch fabric delay derivative value, i.e. *reducing the processing congestion will already deteriorate too much the switch congestion*. Also, for this reason, linecard 5 does not participate in the load-balancing. The underlined values are the final weighted derivative values of the linecards in  $\Gamma_{\min}$ (top) and  $\Gamma_{\max}$ (bottom).

$$\sqrt{\frac{\mu_i}{\mu_j}} = \frac{\mu_i - \gamma_i + R_i}{\mu_j - \gamma_j + R_j}$$

$$\left[ \frac{\mu_i}{(\mu_i - \gamma_i)^2} = \frac{\mu_j}{(\mu_j - \gamma_j)^2}, \sqrt{\frac{\mu_i}{\mu_j}} (\mu_j - \gamma_j) = \mu_i - \gamma_i \right]$$

$$\sqrt{\frac{\mu_i}{\mu_j}} (\mu_j - \gamma_j + R_j) = \mu_i - \gamma_i + R_i \sqrt{\frac{\mu_i}{\mu_j}} (\mu_j - \gamma_j) = \mu_i - \gamma_i$$

$$\frac{R_j}{\sqrt{\mu_j}} = \frac{R_i}{\sqrt{\mu_i}}, B_i^{\max/\min} = \frac{\sqrt{\mu_i}}{\sum_{k \in \Gamma_{\max/\min}} \sqrt{\mu_k}}$$

Let's provide an intuitive example to show how it works, as further illustrated in Figure 6.

**Example 5.** Consider the same arrival rates as in Example 2, together with a finite-capacity switch fabric of switching rate equal to the sum of the linecard processing rates, i.e.  $\mu^{\text{SF}} = 5 \cdot 7$ .

As shown in Figure 6, the algorithm progressively decides to load-balance more and more flows. We can see how at the end of the algorithm, the resulting arrival rates into the linecard queues are not all equal to  $\lambda_{\text{avg}}$ , as they were in Example 2. It is because the switch fabric is limiting the amount of redirected traffic. This illustrates the tradeoff between switch-fabric congestion and linecard processing congestion. In fact, linecard 5 ends up not participating in the load-balancing.

The following result shows that the algorithm complexity is relatively low.

**Theorem 6.** The complexity of reaching the optimal solution within an error bound of  $\varepsilon$  is  $O(N + \log_2(\frac{1}{\varepsilon} \cdot \sum_{i \in [1, N]} \mu_i))$ .

*Proof:*

The algorithm adds linecards to  $\Gamma_{\max}$  and  $\Gamma_{\min}$  for at most  $N - 2$  iterations, because the groups are initialized to contain at least one member. At the last iteration, the gap between the values of  $s_i$  and  $r_i$  is smaller than  $\Delta$ , therefore the algorithm does not perform it. Instead it searches for a redirected traffic size such that:

$$\underbrace{s_i}_{i \in \Gamma_{\max}} = \underbrace{r_j}_{j \in \Gamma_{\min}} + \Delta$$

---

### Algorithm 2: Single-Queue algorithm

---

**Input :** Vectors  $\{\gamma_i\}$  and  $\{\mu_i\}$  of length  $N$  and  $\mu^{\text{SF}}$

**Output:** Vectors  $R^-$  and  $R^+$  of length  $N$

---

- 1  $R^- = 0_N$ ;
  - 2  $R^+ = 0_N$ ;
  - 3  $R = 0$ ;
  - 4  $\forall i : n_i = \frac{\mu_i}{(\mu_i - \gamma_i)^2}$  (neutral);
  - 5  $\Gamma_{\max} = \{i : n_i = \max_i(n_i)\}$ ;
  - 6  $\Gamma_{\min} = \{i : n_i = \min_i(n_i)\}$ ;
  - 7  $B_i^{\max/\min} = \frac{\sqrt{\mu_i}}{\sum_{k \in \Gamma_{\max/\min}} \sqrt{\mu_k}}$ ;
  - 8  $\Delta = \frac{\mu^{\text{SF}}}{(\mu^{\text{SF}} - \sum_{i=1}^N \gamma_i - \frac{1}{2}(\sum_{i=1}^N R_i^+ + \sum_{i=1}^N R_i^-) - R)^2}$ ;
  - 9  $\forall i \in \Gamma_{\max} : s_i = \frac{\mu_i}{(\mu_i - \gamma_i + R_i^- + B_i^{\max} R)^2}$  (sending);
  - 10  $\forall i \in \Gamma_{\min} : r_i = \frac{\mu_i}{(\mu_i - \gamma_i - R_i^+ - B_i^{\min} R)^2}$  (receiving);
  - 11  $\forall i \in \Gamma_{\min} \cup \Gamma_{\max} : \text{remove } n_i$ ;
  - 12 **while**  $\min(\{r_i\}) + \Delta < \max(\{s_i\})$  **do**
  - 13     Update  $\Gamma_{\max}, \Gamma_{\min}, B_i^{\max/\min}$ ;
  - 14      $\forall i \in \Gamma_{\max} : R_i^- += B_i^{\max} R$ ;
  - 15      $\forall i \in \Gamma_{\min} : R_i^+ += B_i^{\min} R$ ;
  - 16      $\text{distMax} \leftarrow \text{Find } R : \max(\{s_i\}) == \max(\{n_i\})$ ;
  - 17      $\text{distMin} \leftarrow \text{Find } R : \min(\{r_i\}) == \min(\{n_i\})$ ;
  - 18      $R = \min(\text{distMax}, \text{distMin})$ ;
  - 19     Update  $\{s_i\}, \{r_i\}, \Delta$ ;
  - 20 **Find**  $R : \min(\{r_i\}) + \Delta = \max(\{s_i\})$ ;
  - 21  $\forall i \in \Gamma_{\max} : R_i^- += B_i^{\max} R$ ;
  - 22  $\forall i \in \Gamma_{\min} : R_i^+ += B_i^{\min} R$ ;
  - 23 **return**  $R^-, R^+$ ;
- 

$$\frac{\frac{\mu_i}{(\mu_i - \gamma_i + R_i^- + B_i^{\max} R)^2}}{\frac{\mu_i}{(\mu_j - \gamma_j - R_j^+ - B_j^{\min} R)^2} + \frac{\mu^{\text{SF}}}{\left(\mu^{\text{SF}} - \sum_{k=1}^N \gamma_k - \frac{1}{2} \left(\sum_{k=1}^N R_k^+ + \sum_{k=1}^N R_k^- + 2R\right)\right)^2}}$$

This is solvable using the bisection method on the possible

values of  $R$ , from 0 to  $\min\left(\frac{\mu_j - \gamma_j - R_j^+}{B_i^{\min}}, \frac{-\gamma_i + R_i^-}{B_i^{\max}}\right)$ , this is the bound on  $R$  to maintain a feasible flow in both the sending and receiving linecards. This value is bounded by the maximal redirected flow in the router  $\sum_{i \in [1, N]} \mu_i - \min_{i \in [1, N]} \mu_i$ , the total processing capacity minus the minimal local processing capacity that will not be redirected. Hence the upper-bound on the complexity of the bisection method is  $\log_2\left(\frac{1}{\epsilon} \cdot \sum_{i \in [1, N]} \mu_i\right)$ . The complexity is therefore at most  $N - 2$  iterations for the  $N - 2$  linecards joining  $\Gamma_{\max/\min}$ , and  $\log_2\left(\frac{1}{\epsilon} \cdot \sum_{i \in [1, N]} \mu_i\right)$  for the bisection method, which is the worst case for root-finding methods. ■

### B. Capacity Region.

We define the capacity region as the set of all non-negative arrival rate vectors  $\{\gamma_i\}$  that are feasible (as in [26], [27]). We want to compare the capacity region of our algorithm against that of a basic *Current* scheme without load-balancing, and of a *Valiant-based* load-balancing scheme that we now introduce (based on [28]).

The capacity region of each linecard without load-balancing is bounded by its own processing capability, regardless of the traffic in other linecards. Using the load-balancing scheme, we can expand the capacity region based upon the *total available processing power* of the router.

**Definition 3** (Valiant-based load-balancing). *Let the Valiant-based algorithm be defined as an oblivious load-balancing scheme where each linecard  $i$  automatically load-balances a uniform fraction  $\frac{1}{N}$  of its incoming traffic to be processed at each linecard  $j$  (and processes locally  $\frac{1}{N}$  as well), regardless of the current load in the different linecards.*

**Theorem 7.** *The capacity region using the different algorithms is as follows:*

- i) Current:  $\forall i : \gamma_i < \mu_i$ , and  $\sum_i \gamma_i < \mu^{SF}$ .
- ii) Valiant:  $\forall i : \frac{\sum_i \gamma_i}{N} < \mu_i$ , and  $\frac{2N-1}{N} \cdot \sum_i \gamma_i < \mu^{SF}$ .
- iii) Single Queue (ours):  $\sum_i \gamma_i < \sum_i \mu_i$ , and


$$\underbrace{\sum_i \min(\gamma_i, \mu_i)}_{\text{processed locally}} + 2 \cdot \underbrace{\sum_i \max(\gamma_i - \mu_i, 0)}_{\text{redirected traffic}} < \mu^{SF}.$$

*Proof:* (i) The result for the current implementation is straightforward, since it cannot redirect traffic.

(ii) This derives from the fact that Valiant load-balancing uniformly redirects  $\frac{N-1}{N}$  of the arrival rate of each linecard to all other linecards.

(iii) Finally, in our scheme, the first equation derives from the processing multiplexing, and the second equation states that the switch-fabric arrival rate is smaller than its service rate. ■

Figure 7 illustrates the theorem results. It plots the capacity region when either  $N = 2$  or  $N = 10$ , using a uniform



(a) Capacity region with  $N = 2$  (b) Capacity region with  $N = 10$ , only 2 linecards have incoming traffic.

Fig. 7: Single-Queue switch fabric model with  $\mu^{SF} = N$ . The colored fills represent the simulation results for the current scheme without load-balancing (dark grey) and our Single-Queue delay-optimal scheme (light grey). The plotted borders represent the theoretical values for the two schemes and for the Valiant-based scheme. We can see how our scheme achieves a larger capacity region in both cases, since the Current scheme does not do load-balancing, and the Valiant scheme performs oblivious and sometimes harmful load-balancing.

service rate  $\mu_i = 1$ ,  $\mu^{SF} = N$ , and assuming that only two linecards have arrivals. We can see how when one of the linecards is idle, the feasible processing rate of the other linecard can significantly increase using our load-balancing scheme. The two other schemes appear sub-optimal, since the Current scheme does not do load-balancing, and the Valiant scheme relies on an oblivious and sometimes harmful load-balancing. The performance gain with respect to Current is especially large when there is dissymmetry between linecard rates, and against Valiant when there is symmetry and load-balancing becomes harmful.

## V. N-QUEUES SWITCH-FABRIC MODEL

We now consider a more accurate switch-fabric model, where the congestion in the switch fabric depends on the destination. As illustrated in Fig. 4, we model an output-queued switch fabric by assigning a queue at each output. This helps us capture output hot-spot scenarios.

### A. Minimizing the Average Delay

As in the previous sections, we want to minimize the average delay of packets in the router, which equals the sum of the delays through all queues, weighted by the flows going through these queues, and normalized by the total incoming exogenous flow. We obtain a similar result on the average-delay minimization problem.

The average delay of a packet in the system consists of the flow through each queue multiplied by the delay of the queue and finally normalized with the total flow in the system.

$$\text{minimize } \bar{D}(R^+, R^-) = \frac{1}{\sum_{i=1}^N \gamma_i} \left( \sum_{i=1}^N \frac{\gamma_i + R_i^+ - R_i^-}{\mu_i - \gamma_i - R_i^+ + R_i^-} + \sum_{i=1}^N \frac{\sum_{j=1}^N \gamma_{j,i} + R_i^+}{\mu_i^{SF} - \sum_{j=1}^N \gamma_{j,i} - R_i^+} \right)$$

$$\begin{aligned}
\text{subject to } & \sum_{i=1}^N R_i^+ - \sum_{i=1}^N R_i^- = 0 \\
& -\gamma_i - R_i^+ + R_i^- \leq 0, i \in [1, N] \\
& \gamma_i + R_i^+ - R_i^- - \mu_i \leq 0, i \in [1, N] \\
& -R_i^+ \leq 0, -R_i^- \leq 0, i \in [1, N] \\
& -\mu_i^{\text{SF}} + \sum_{j=1}^N \gamma_{j,i} + R_i^+ \leq 0, i \in [1, N]
\end{aligned}$$

**Theorem 8.** A solution to the average-delay minimization problem for the  $N$ -queues switch fabric model is optimal iff (i) the flows are feasible, and (ii) there exists a constant  $\tau_0$  such that for each linecard  $j$ , if  $R_j^+ > 0$  :

$$\frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} + \frac{\mu_j^{\text{SF}}}{\left(\mu_j^{\text{SF}} - \sum_{i=1}^N \gamma_{i,j} - R_j^+\right)^2} = -\tau_0 \sum_{i=1}^N \gamma_i;$$

$$\text{and if } R_j^- > 0 : \frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = -\tau_0 \sum_{i=1}^N \gamma_i.$$

*Proof:* We will write the KKT conditions, which are the necessary and sufficient conditions for a solution to be optimal in our case of convex function and region.

The KKT multipliers are:  $\tau_0, \eta_i, i \in \{1, \dots, 5N\}$ . There are  $5N$  constants,  $\eta_i$ , one for each inequality constraint of the problem.  $\tau_0$  is used for the single equality constraint.

The stationarity equations are, for all  $j$ :

$$\begin{aligned}
& -\frac{\partial}{\partial R_j^{+/-}} \bar{D}(R^+, R^-) = \\
& \sum_{i=1}^N \eta_i \frac{\partial}{\partial R_j^{+/-}} (-\gamma_i - R_i^+ + R_i^-) \\
& + \sum_{i=N+1}^{2N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (-R_{i-N}^+) \\
& + \sum_{i=2N+1}^{3N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (-R_{i-2N}^-) \\
& + \sum_{i=3N+1}^{4N} \eta_i \frac{\partial}{\partial R_j^{+/-}} \left( -\mu_{i-3N}^{\text{SF}} + \sum_{k=1}^N \gamma_{k,i-3N} + R_{i-3N}^+ \right) \\
& + \sum_{i=4N+1}^{5N} \eta_i \frac{\partial}{\partial R_j^{+/-}} (\gamma_{i-4N} + R_{i-4N}^+ - R_{i-4N}^- - \mu_{i-4N}) \\
& + \tau_0 \frac{\partial}{\partial R_j^{+/-}} \left( \sum_{i=1}^N R_i^+ - \sum_{i=1}^N R_i^- \right)
\end{aligned}$$

The dual feasibility and complementary slackness equations are:

$$\begin{aligned}
& \eta_i \geq 0, \forall i \in \{1, \dots, 5N\} \\
& \eta_i (-\gamma_i - R_i^+ + R_i^-) = 0, \forall i \in [1, N] \\
& \eta_i (-R_{i-N}^+) = 0, \forall i \in [N+1, 2N] \\
& \eta_i (-R_{i-2N}^-) = 0, \forall i \in [2N+1, 3N] \\
& \eta_i (-\mu_{i-3N}^{\text{SF}} + \sum_{k=1}^N \gamma_{k,i-3N} + R_{i-3N}^+) = 0, \\
& \quad \forall i \in [3N+1, 4N] \\
& \eta_i (\gamma_{i-4N} + R_{i-4N}^+ - R_{i-4N}^- - \mu_{i-4N}) = 0, \\
& \quad \forall i \in \{4N+1, \dots, 5N\}
\end{aligned}$$

Similarly to before, if the traffic is feasible, the service rate of the linecards and switch fabric is greater than the arrival rate into their queues, thus  $\eta_{3N+j} = 0$  and  $\eta_{4N+j} = 0, \forall j \in [1, N]$ . For a certain linecard  $j$ , if  $\gamma_i + R_i^+ - R_i^- > 0$ , meaning it is processing packets, then  $\eta_j = 0$ . If  $R_j^+ > 0$  then  $\eta_{N+j} = 0$ :

$$\frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} + \frac{\mu_j^{\text{SF}}}{\left(\mu_j^{\text{SF}} - \sum_{i=1}^N \gamma_{i,j} - R_j^+\right)^2} = -\tau_0 \sum_{i=1}^N \gamma_i$$

If  $R_j^- > 0$ , then  $\eta_{2N+j} = 0$ :  $\frac{\mu_j}{(\mu_j - \gamma_j - R_j^+ + R_j^-)^2} = -\tau_0 \sum_{i=1}^N \gamma_i$ . ■

From the optimization conditions, it is easy to deduce an algorithm that works quite similarly to the algorithm in the previous section. The algorithm solves the optimization problem by performing a binary search on the possible values of  $R_k^+$ , where  $k$  is the linecard with the smallest combined (linecard and switch fabric) delay derivative. At each step of the binary search, all the linecards with a lesser (respectively, greater) combined derivative are considered to belong to the minimal (maximal) derivative group, and have traffic redirected to them (away from them) until their combined derivative value equals that of linecard  $k$ .

In Algorithm 3 we first initialize the different algorithm variables. The main difference from Algorithm 2 is seen in the  $\Gamma_{\min}$  definition, with the added contribution of the switch fabric to the calculation. The algorithm performs a binary search on the redirected traffic size into the minimal congestion valued linecard, at each step recalculating the  $\Gamma_{\min}$  and  $\Gamma_{\max}$  groups.

The index  $k$  denotes the linecard with minimal congestion, the linecard that allows the maximal redirected traffic into it. All the linecards that fulfill  $n_i > r_k + \Delta_k$  belong to  $\Gamma_{\max}$  since it is beneficial for them to redirect traffic into linecard  $k$ . All the linecards that fulfill  $r_k + \Delta_k > n_i + \Delta_i$  belong to  $\Gamma_{\min}$  since their derivative value is lower than that of linecard  $k$  and therefore they will also receive redirected traffic.

After we have calculated  $\Gamma_{\max}$  and  $\Gamma_{\min}$  we compare the total redirected traffic and the total received traffic at each group. If the redirected traffic is greater, linecard  $k$  needs to receive more traffic, if the redirected traffic is smaller then linecard  $k$  needs to receive less.

**Theorem 9.** The complexity of reaching the optimal solution within an error bound of  $\varepsilon$  is  $O(N \cdot \log_2(\frac{\mu}{\varepsilon})^2)$ .

*Proof:* The above formula assumes equal service rate linecards, the general complexity term is

$$O\left(\max_{i \in [1, N]} (\log_2(\frac{\mu_i}{\varepsilon})) \cdot \sum_{i \in [1, N]} \log_2(\frac{\mu_i}{\varepsilon})\right).$$


---

**Algorithm 3: N-Queues algorithm (short)**


---

**Input :** Vectors  $\{\mu_i\}$  and  $\{\mu_i^{\text{SF}}\}$  of length  $N$   
Matrix  $N \times N \{\gamma_{i,j}\}$

**Output:** Vectors  $R^-$  and  $R^+$  of length  $N$

```

1  $R^- = 0_N$ ;
2  $R^+ = 0_N$ ;
3  $\forall i : n_i = \frac{\mu_i}{(\mu_i - \gamma_i)^2}$  (neutral);
4  $\forall i \in \Gamma_{\max} : s_i = \frac{\mu_i}{(\mu_i - \gamma_i + R_i^-)^2}$  (sending);
5  $\forall i \in \Gamma_{\min} : r_i = \frac{\mu_i}{(\mu_i - \gamma_i - R_i^+)^2}$  (receiving);
6  $\forall i : \Delta_i = \frac{\mu_i^{\text{SF}}}{(\mu_i^{\text{SF}} - \sum_{k=1}^N \gamma_{k,i} - R_i^+)^2}$ ;
7  $\Gamma_{\max} = \{i : n_i = \max_i(n_i)\}$ ;
8  $\Gamma_{\min} = \{i : n_i + \Delta_i = \min_i(n_i + \Delta_i)\}$ ;
9  $\forall i \in \Gamma_{\min} : R_i^* = \min(\mu_i - \gamma_i, \mu_i^{\text{SF}} - \sum_{k=1}^N \gamma_{k,i})$ ;
10  $k = \arg \max_i(R_i^*)$ ;
11 iter = 1;
12 sign = 1;
13  $R_k^+ = 0$ ;
14 while  $\frac{R_k^*}{2^{\text{iter}}} > \varepsilon$  do
15    $R_k^+ += \text{sign} \cdot \frac{R_k^*}{2^{\text{iter}}}$ ;
16    $\forall i \neq k : R_i^+ = 0$ ;
17    $R^- = 0_N$ ;
18   Update  $\{n_i\}, \Delta_i$ ;
19    $\forall i : n_i > r_k + \Delta_k$ , Find  $R_i^- : s_i = r_k + \Delta_k$ 
    $\forall i : r_k + \Delta_k > n_i + \Delta_i$ , Find  $R_i^+ : r_k + \Delta_k = r_i + \Delta_i$ 
20    $R_{\text{total}}^- = \sum R^-$ ;
21    $R_{\text{total}}^+ = \sum R^+$ ;
22   if  $R_{\text{total}}^- > R_{\text{total}}^+$  then
23     | sign = 1;
24   else
25     | sign = -1;
26 return  $R^-, R^+$ ;

```

---

In the previous sections the switch fabric was shared between all linecards and therefore simplified the calculations. It is no longer possible to extract a simple  $B_i^{\min}$  expression, the expression used to maintain equal value across all linecards in  $\Gamma_{\min}$  as we redirect traffic to them.  $B_i^{\max}$  on the other hand still fulfills the equation  $B_i^{\max} = \frac{\sqrt{\mu_i}}{\sum_{k \in \Gamma_{\max}} \sqrt{\mu_k}}$  since the maximal expression remains the same. A significant change is required to the algorithm presented at the previous section.

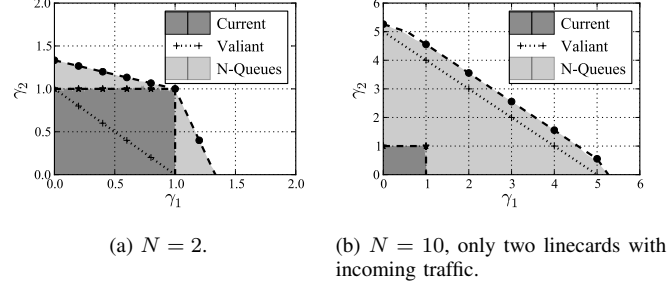


Fig. 8: Capacity region using the N-Queues switch fabric model,  $\mu_i^{\text{SF}} = 1$ . The plotted borders are the theoretical values and the colored fill is the simulation result.

This is due to the fact that at each step:

$$\forall i, j \in \Gamma_{\min},$$

$$\frac{\mu_i}{(\mu_i - \gamma_i)^2} + \frac{\mu_i^{\text{SF}}}{\left(\mu_i^{\text{SF}} - \sum_{k=1}^N \gamma_{k,i}\right)^2} =$$

$$\frac{\mu_j}{(\mu_j - \gamma_j)^2} + \frac{\mu_j^{\text{SF}}}{\left(\mu_j^{\text{SF}} - \sum_{k=1}^N \gamma_{k,j}\right)^2}$$

The algorithm needs to find  $R_i^+$  and  $R_j^+$  such that  $\sum_{k \in \Gamma_{\min}} R_k^+ = \sum_{k \in \Gamma_{\max}} R_k^-$  and:

$$\frac{\mu_i}{(\mu_i - \gamma_i - R_i^+)^2} + \frac{\mu_i^{\text{SF}}}{\left(\mu_i^{\text{SF}} - \sum_{k=1}^N \gamma_{k,i} - R_i^+\right)^2} =$$

$$\frac{\mu_j}{(\mu_j - \gamma_j - R_j^+)^2} + \frac{\mu_j^{\text{SF}}}{\left(\mu_j^{\text{SF}} - \sum_{k=1}^N \gamma_{k,j} - R_j^+\right)^2}$$

There are at most  $\max_{i \in [1, N]} (\log_2(\frac{\mu_i}{\varepsilon}))$  while loop iterations. Each iteration requires a binary search for each linecard in  $\Gamma_{\max}$  and  $\Gamma_{\min}$ , at most  $\sum_{i \in [1, N]} \log_2(\frac{\mu_i}{\varepsilon})$  iterations. Therefore the total complexity is

$$O\left(\max_{i \in [1, N]} (\log_2(\frac{\mu_i}{\varepsilon})) \cdot \sum_{i \in [1, N]} \log_2(\frac{\mu_i}{\varepsilon})\right)$$

### B. Capacity Region.

Similarly to the single-queue model, we again plot the capacity region. We assume the same settings, and also assume that the packet destinations are uniformly distributed and that the output queue services rates are equal to  $\mu_i^{\text{SF}} = 1$ . As illustrated in Figure 8, the results are largely similar, although the two load-balancing algorithms now slightly suffer from the increased switch fabric congestion, since the service rate of the switch fabric is now split between  $N$  queues.

## VI. SIMULATION RESULTS

In this section we compare the performance of our suggested algorithms to the existing algorithms. We first simulate the *N-Queues* switch-fabric queuing architecture using synthetic traces. Afterwards, we simulate an iSLIP input-queued switch-fabric scheduling algorithm with traces from high-speed Internet backbone links.

To clarify, we compare all algorithms *on the same architecture*, but each algorithm determines its load-balancing policy based on its own switch-fabric model. Thus, the algorithm model is not necessarily identical to the simulated architecture.

### A. *N-Queues* Switch Fabric Algorithm

Figure 9 simulates all algorithms on the output-queued switch-fabric architectural model that appears in Figure 4. The figures plot the performance of five algorithms: (i) the baseline Current algorithm without load-balancing, (ii) the Valiant-based algorithm, (iii) our load-balancing algorithm that assumes an infinite switch-fabric service rate, (iv) our single-queue switch-fabric algorithm, and (v) our *N-queues* switch-fabric algorithm, which is expected to outperform since its model is closest to the architecture. The service rates are  $\forall i \in [1, N], \mu_i = \mu_i^{\text{SF}} = 1$ . The destination linecard is chosen uniformly.

Figure 9(a) confirms the capacity region in Figure 8(b), where two linecards have positive incoming arrival rates and the other linecards are idle. As expected, the Current algorithm cannot accommodate an arrival rate that is larger than the service rate of a single linecard. Furthermore, as in Figure 8(b), the capacity region for the Valiant-based load-balancing is capped at  $\gamma_{1,2} = 2.5$ , and for our algorithm at  $\gamma_{1,2} = 2.8$ .

Figure 9(b) shows the delays when five out of ten linecards have an equal arrival rate and the rest are idle. The Valiant-based load-balancing performs significantly worse due to the fact that now  $\frac{4}{N} = 40\%$  of all the redirections it performs are redundant and towards active linecards, whereas with two active linecards only 10% of the redirections were redundant. An interesting result is at the low arrival rates, where the Valiant-based algorithm, and to a lesser degree our Infinite algorithm, perform poorly and have a larger average delay than without load-balancing. The reason is that the switch fabric service rate is split between  $N = 10$  queues, therefore each packet experiences a larger delay when passing through the switch fabric and these algorithms have many redundant redirections at low arrival rates. On the other hand, our optimal algorithm decides at these low arrival rates not to redirect.

In Figure 9(c) the arrival rates are assumed to be distributed linearly at equal intervals. For instance, when the  $x$ -axis is at 1.8, the arrival rates to the linecards are  $(0.0, 0.2, 0.4, \dots, 1.8)$ . The result appears similar.

Incidentally, note that while our *Infinite* algorithm performs slightly worse than our more accurate algorithms, it may be interesting when complexity is an issue. In addition, it is possible to artificially improve the infinite switch fabric algorithm by simply preventing it from redirecting when the maximal load below a certain constant. This will not add to

the complexity of the algorithm, and it would even prevent its redundant full calculation at low loads.

### B. Real Traces

We now simulate a real input-queued switch fabric with VOQs and an implementation of the iSLIP [10] scheduling algorithm. Upon arrival at the switch fabric, each packet is divided into evenly-sized cells and sent to a queue based on its source and destination linecards. At each scheduling step, the iSLIP scheduling algorithm matches sources to destinations, and the switch fabric transmits cells based on these pairings. The destination linecard must receive all of the packet's cells before the packet can continue.

For this simulation we rely on real traces from [9], which were recorded at two different monitors on high-speed Internet backbone links. One of the traces has an average OC192 line utilization of 24%, and the other 3.2%.

Note that in these simulations, all the arrival rates into the linecards are positive, while in reality many router linecards are simply disconnected. Thus, *in real-life, we would expect load-balancing schemes to be even more beneficial*.


The switch fabric rate is set to be equivalent to OC192 per matching. The processing time is calculated based on the packet size using the number of instructions and memory-accesses estimate from [7] for header and payload processing applications. Memory access times are assumed to be constant at 4ns.

The algorithms in previous sections require several adjustments to accommodate the more realistic features of this simulation. The average arrival rates at each linecard are no longer known in advance, and no longer constant. Instead, the algorithm performs periodic updates. At each interval the algorithm measures the arrival rates to all linecards, and based on previous intervals it predicts the arrival rates in the next interval and calculates the required load-balancing policy based on this prediction. In this section the updates were made at thousand-packet intervals.

Additionally, the linecards' processing time depends on the number of instructions and memory accesses, which are derived from the packet's payload size, while the packet's time at the switch fabric depends directly on the total packet size. In contrast, the load-balancing algorithms assume the switch fabric consists of a single queue or  $N$  queues, therefore a simple conversion of the actual switch fabric rate to the algorithm's simplified switch fabric model rate is required.

Figure 10(a) shows a comparison between the Current implementation, Valiant-based load-balancing, and our algorithms. The packet arrival times were divided by two to simulate a more congested network with 48% and 6.4% OC192 line utilization. There are  $N$  linecards in the system, five of which are connected to the relatively highly utilized line rates and the remaining linecards are connected to the lower arrival rate links. The  $x$ -axis shows *the average processing load of the most congested line card*. The change in the  $x$ -axis only affects the service rate of the line cards while the switch fabric rate and the arrival rates stay constant.

The Valiant-based delay is high, even at low loads. This is due to the significant additional congestion at the switch fabric.



(a) Two linecards out of ten have incoming traffic. (b) Five linecards out of ten have incoming traffic. (c) Arrival rates into linecards distributed at equal intervals between 0 and x-axis value.

Fig. 9: N-Queues switch fabric model, with  $N = 10$  and a Poisson arrival model.

This delay is somewhat constant up to very high loads, where the delay begins to be mostly influenced by the congested linecards. Our algorithm does not redirect traffic at low loads and maintains a delay similar to the delay without load-balancing. However, for a small range around a processing load of 0.6, it has a slightly higher delay, because it redirects traffic earlier than it should due to inaccurate arrival rate estimation. Significantly, our algorithm almost doubles the capacity region of the most congested linecard in the system when compared to the Current algorithm. Of course, load-balancing algorithms with less accurate models perform with slightly higher delays.


Figure 10(b) shows the same comparison, but now *the service rates of the linecards are constant while the arrival rates are variable*. The N-queues algorithm performs significantly better than the other algorithms in this scenario because it takes into consideration the different loads at the switch fabric queues. As the arrival rate increases, the switch fabric delay grows along with the processing delay in the linecards.

Figure 11 shows the total utilization of the processing resources in the router, defined as the fraction of the time the linecard processors are busy. Figure 11(a) shows how for a constant arrival rate and a variable processing rate, the utilization of the Current implementation falls off at higher loads, while the Valiant-based algorithm performs similarly to the more complex algorithms. The reason for this is due to the switch fabric being able to withstand the additional load of the Valiant algorithm. On the other hand, in Figure 11(b), with a variable arrival rate and a constant processing rate, both Current and Valiant perform worse. In particular, at high arrival rates, the switch fabric limits the Valiant algorithm, and therefore its utilization of the processors does not scale as well as our algorithms.

## VII. DISCUSSION


### A. Queue-Length-Based Algorithms

Our algorithms balance the traffic load *based on the arrival rates to the linecards*. A different approach would be to load-balance *based on the queue lengths*. For instance, if the switch-fabric capacity were infinite, as in our first switch architecture, and the queue lengths were known at all times, then a simple redirect to the shortest queue approach would be *optimal* [29]. However, such global and constantly updated knowledge would of course be costly to maintain across all the



(a) Constant arrival rate, variable packet processing load. (b) Variable arrival rate, constant packet processing load.

Fig. 10: Simulation with OC-192 Internet backbone link traces [9], iSLIP switch fabric scheduling algorithm [10] and packet processing times based on [7], comparing our three algorithms with the Current and Valiant algorithms.




(a) Constant arrival rate, variable packet processing load. (b) Variable arrival rate, constant packet processing load.

Fig. 11: Simulation with same settings as Figure 10, comparing the utilization of the algorithms.

linecards, and therefore updates regarding the current queue lengths would be sent periodically, similarly to the updates of the arrival rates.

Based on the *power of two choices* [30], we introduce two simple algorithms. First, in the *2-Choices* algorithm, each packet compares the queue sizes at two random linecards and is redirected to the shortest one. Second, in the *2-Modified* algorithm, each incoming packet only compares the queue sizes of its ingress linecard and of another random linecard, and chooses the linecard with the smallest queue.

Figures 12(a) and 12(b) compare the different algorithms. Figure 12(a) uses a constant arrival rate and variable service rates at the linecards, while Figure 12(b) assumes a variable arrival rate and constant service rates at the linecards. The 2-Modified algorithm performs better than 2-Choices at low processing loads, since it needs less redundant redirects. In any case, our N-Queues algorithm seems to outperform both.



(a) Constant arrival rate, variable packet processing load. (b) Variable arrival rate, constant packet processing load.

Fig. 12: Simulation with same settings as Figure 10, comparing the performance of our best algorithm with queue-length-based algorithms.

	Memory	Complexity
No L-B	-	-
Valiant-Based	-	$O(1)$
2 Choices	$O(N)$	$O(1)$
Infinite SF	$O(N)$	$O(N)$
Single Queue SF	$O(N)$	$O(N + \log_2(\frac{N\mu}{\epsilon}))$
N-Queues SF	$O(N)$	$O(N \cdot \log_2(\frac{\mu}{\epsilon})^2)$

TABLE I: Comparison of the different algorithm requirements.

## B. Implementation

We initially had several significant *implementation concerns* regarding our algorithms, in particular regarding (a) *buffer management* and (b) *reordering*. However, these concerns were alleviated following discussions with industry vendors.

First, our algorithms may cause problems if redirecting traffic requires to first *reserve the buffer* at the helper linecard. But this is not significantly different from reserving the buffer at the output, and therefore seems reasonable. The algorithms could also *bypass the buffer* at the input linecard when redirecting, unlike what is currently done. But this is simply about updating the implementation, and does not appear to cause any fundamental issues.

A second potential concern is that our algorithms may cause *reordering*, thus damaging the performance of TCP flows [31]. But there are many papers and vendor implementations that address this concern [21], [32], [33]. For instance, it is possible to change the granularity of the load-balancing from a per-packet decision to a per-flow decision. In fact, we reran our simulations based on OC192 traces, with over  $10^7$  different flows appearing at each second in each linecard, and found that this change did not influence the performance (within an approximation of  $10^{-5}$ ). This is expected because the granularity of the load-balancing with this number of flows is fine enough for any practical use. An alternative technique would be to maintain reordering buffers at the egress paths, although this would of course require additional memory [33]. In addition, [21], [32] have found that in the millisecond level, only hundreds or thousands of flows are active, being processed or buffered, out of millions of alive flows. Allowing us to redirect packets mid-flow without causing reordering when the inter-packet gap is larger than the maximal processing time of a packet.

More minor concerns are about the *algorithm overhead*. Our algorithms need to periodically collect information about the various linecards and in addition incur computational

	Memory	Complexity
No L-B	-	-
Valiant-Based	-	$O(1)$
2 Choices	$O(N)$	$O(1)$
Infinite SF	$O(N)$	$O(N)$
Single Queue SF	$O(N)$	$O(N + \log_2(\frac{N\mu}{\epsilon}))$
N-Queues SF	$O(N)$	$O(N \cdot \log_2(\frac{\mu}{\epsilon})^2)$

TABLE II: Comparison of the different algorithm requirements.

complexity and memory overheads. However, in the worst case every linecard transmits an update to every other linecard. Assuming  $N = 64$  linecards, 1,000 updates per second and 32 bits representing the flow size, yields a total overhead of 131Mbps, which is negligible in routers with total rates of over 100 Gbps. In addition, our algorithms incur other overheads, including complexity and memory overheads. The complexity overhead of our most complex algorithm is roughly of the same order of magnitude as running a typical payload processing application on several packets based on [7]. This computation is done every update interval, i.e. every 10,000 packets over the whole switch in most of our simulations, and therefore the amortized per-packet additional computation is less than 0.1%. Also, the required memory per linecard is at most  $2N \cdot 32 = 3,200$  bits to represent the redirection probabilities and congestion values per linecard. Therefore, we find that these overheads appear quite reasonable.

## VIII. CONCLUSION

In this paper, we have introduced a novel technique for tapping into the processing power of underutilized linecards. Given several switch-fabric models, we introduced load-balancing algorithms that were shown to be delay-optimal and throughput-optimal. We also illustrated their strong performance in simulations using realistic switch architectures.

Of course, there remain many interesting directions for future work. In particular, we are interested in exploring the possibility for router vendors of adding helper linecards with no incoming traffic, with the unique goal of assisting existing linecards to process incoming traffic. In addition, this idea may be merged with the recent *NFV (Network Function Virtualization) proposals*: such helper linecards may in fact contain multi-core processors that would also be devoted to additional programmable tasks, thus avoiding the need to send these tasks to remote servers.

## IX. ACKNOWLEDGMENT

The authors would like to thank Shay Vargaftik, Aran Bergman, Rami Atar and Amir Rosen for their helpful comments. This work was partly supported by the Mel Berlin Fellowship, the Gordon Fund for Systems Engineering, the Neptune Magnet Consortium, the Israel Ministry of Science and Technology, the Intel ICRI-CI Center, the Hasso Plattner Institute Research School, the Technion Funds for Security Research, and the Erteschik and Greenberg Research Funds.

## REFERENCES

- [1] EZchip. (2013) EZchip Technologies — NPS Family. [Online]. Available: [http://www.ezchip.com/p\\_nps\\_family.htm](http://www.ezchip.com/p_nps_family.htm)

- [2] Cisco. (2013) Cisco 4451-x integrated services router data sheet. [Online]. Available: [http://www.cisco.com/en/US/prod/collateral/routers/ps10906/ps12522/ps12626/data\\_sheet\\_c78-728190.html](http://www.cisco.com/en/US/prod/collateral/routers/ps10906/ps12522/ps12626/data_sheet_c78-728190.html)
- [3] R. Atkinson and S. Kent, "RFC 2406 – IP encapsulating security payload (ESP)," 1998.
- [4] M. Roesch, "Snort: Lightweight intrusion detection for networks." in *LISA*, vol. 99, 1999, pp. 229–238.
- [5] H.-A. Kim and B. Karp, "Autograph: Toward automated, distributed worm signature detection." in *USENIX security symposium*, 2004.
- [6] S. Singh, C. Estan, G. Varghese, and S. Savage, "Automated worm fingerprinting." in *OSDI*, vol. 4, 2004, pp. 4–4.
- [7] R. Ramaswamy, N. Weng, and T. Wolf, "Analysis of network processing workloads," *Journal of Systems Architecture*, vol. 55, no. 10-12, 2009.
- [8] F. Abel *et al.*, "Design issues in next-generation merchant switch fabrics," *IEEE/ACM Trans. Netw.*, vol. 15, no. 6, pp. 1603–1615, 2007.
- [9] The CAIDA UCSD anonymized Internet traces 2011 - 20110217-130100-130500. [Online]. Available: [http://www.caida.org/data/passive/passive\\_2011\\_dataset.xml](http://www.caida.org/data/passive/passive_2011_dataset.xml)
- [10] N. McKeown, "The iSLIP scheduling algorithm for input-queued switches," *IEEE/ACM Trans. Netw.*, vol. 7, no. 2, pp. 188–201, 1999.
- [11] R. Ennals, R. Sharp, and A. Mycroft, "Task partitioning for multi-core network processors," in *Compiler Construction*, 2005.
- [12] Y. Qi *et al.*, "Towards high-performance flow-level packet processing on multi-core network processors," in *ACM/IEEE ANCS*, 2007.
- [13] R. Sommer, V. Paxson, and N. Weaver, "An architecture for exploiting multi-core processors to parallelize network intrusion prevention," *Concurrency and Computation*, vol. 21, no. 10, pp. 1255–1279, 2009.
- [14] I. Keslassy, K. Kogan, G. Scalosub, and M. Segal, "Providing performance guarantees in multipass network processors," *IEEE/ACM Trans. Netw.*, vol. 20, no. 6, pp. 1895–1909, Dec. 2012.
- [15] Y. Afek *et al.*, "MCA2: multi-core architecture for mitigating complexity attacks," *ACM/IEEE ANCS*, pp. 235–246, 2012.
- [16] O. Rottenstreich, I. Keslassy, Y. Revah, and A. Kadosh, "Minimizing delay in shared pipelines," *IEEE Hot Interconnects*, pp. 9–16, 2013.
- [17] A. Shpiner, I. Keslassy, and R. Cohen, "Scaling multi-core network processors without the reordering bottleneck," *IEEE HPSR*, pp. 146–153, Jul. 2014.
- [18] C.-S. Chang, D.-S. Lee, and Y.-S. Jou, "Load balanced Birkhoff-von Neumann switches," *Computer Communications*, 2002.
- [19] I. Keslassy, "The load-balanced router," Ph.D. dissertation, Stanford University, 2004.
- [20] B. Lin and I. Keslassy, "The concurrent matching switch architecture," *IEEE/ACM Trans. Netw.*, vol. 18, no. 4, pp. 1330–1343, 2010.
- [21] Y. Kai, Y. Wang, and B. Liu, "GreenRouter: Reducing power by innovating router's architecture," *IEEE CAL*, 2013.
- [22] M. Varvello, D. Perino, and J. Esteban, "Caesar: a content router for high speed forwarding," in *ACM workshop on Information-centric networking*, 2012, pp. 73–78.
- [23] M. Mitzenmacher, "The power of two choices in randomized load balancing," *IEEE TPDS*, vol. 12, no. 10, pp. 1094–1104, 2001.
- [24] F. P. Kelly, *Reversibility and stochastic networks*. Cambridge University Press, 2011.
- [25] S. Balsamo, "Product form queueing networks," *Performance Evaluation*, pp. 377–401, 2000.
- [26] M. J. Neely, "Delay-based network utility maximization," *IEEE/ACM Trans. Netw.*, vol. 21, no. 1, pp. 41–54, 2013.
- [27] —, "Stability and capacity regions for discrete time queueing networks," *arXiv preprint arXiv:1003.3396*, 2010.
- [28] L. G. Valiant and G. J. Brebner, "Universal schemes for parallel communication," in *ACM Symp. on Theory of Computing*, 1981.
- [29] W. Winston, "Optimality of the shortest line discipline," *Journal of Applied Probability*, pp. 181–189, 1977.
- [30] M. Mitzenmacher, "How useful is old information?" *IEEE TPDS*, vol. 11, no. 1, pp. 6–20, 2000.
- [31] J. Bennett, C. Partridge, and N. Shectman, "Packet reordering is not pathological network behavior," *IEEE/ACM Trans. Netw.*, 1999.
- [32] C. Hu, Y. Tang, X. Chen, and B. Liu, "Per-flow queueing by dynamic queue sharing," in *IEEE Infocom*, 2007, pp. 1613–1621.
- [33] O. Rottenstreich *et al.*, "The switch reordering contagion: Preventing a few late packets from ruining the whole party," *IEEE Trans. Comput.*, vol. 63, no. 5, 2014.