Costly Circuits, Submodular Schedules and Approximate Carathéodory Theorems

Shaileshh Bojja
Venkatakrisnan
University of Illinois
Urbana-Champaign
bjjvnt2@illinois.edu

Mohammad Alizadeh
Massachusetts Institute of
Technology
alizadeh@csail.mit.edu

Pramod Viswanath
University of Illinois
Urbana-Champaign
pramodv@illinois.edu

ABSTRACT

Hybrid switching – in which a high bandwidth circuit switch (optical or wireless) is used in conjunction with a low bandwidth packet switch – is a promising alternative to interconnect servers in today’s large scale data centers. Circuit switches offer a very high link rate, but incur a non-trivial reconfiguration delay which makes their scheduling challenging. In this paper, we demonstrate a lightweight, simple and nearly-optimal scheduling algorithm that trades-off reconfiguration costs with the benefits of reconfiguration that match the traffic demands. Seen alternatively, the algorithm provides a fast and approximate solution towards a constructive version of Carathéodory’s Theorem for the Birkhoff polytope. The algorithm also has strong connections to submodular optimization, achieves a performance at least half that of the optimal schedule and strictly outperforms state of the art in a variety of traffic demand settings. These ideas naturally generalize: we see that indirect routing leads to exponential connectivity; this is another phenomenon of the power of multi-hop routing, distinct from the well-known load balancing effects.

1. INTRODUCTION

Modern data centers are massively scaling up to support demanding applications such as large-scale web services, big data analytics, and cloud computing. The computation in these applications is distributed across tens of thousands of interconnected servers. As the number and speed of servers increases,¹ providing a fast, dynamic, and economic switching interconnect in data centers constitutes a topical networking challenge. Typically, data center networks use multi-rooted tree designs: the servers are arranged in racks and an Ethernet switch at the top of the rack (ToR) connects the rack of servers to one or more aggregation (or spine) layers. These designs use multiple paths between the ToRs

¹Servers with 10Gbps network interfaces are common today and 40/100Gbps servers are being deployed.

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

SIGMETRICS '16, June 14-18, 2016, Antibes Juan-Les-Pins, France

© 2016 ACM. ISBN 978-1-4503-4266-7/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2896377.2901479>

to deliver uniform high bisection bandwidth, and consist of a large number of high speed electronic packet switches that provide fine-grained switching capabilities but at poor speed/cost ratios.

Recent work has proposed the use of high speed circuit switches based on optical [46, 16, 53] or wireless [27, 56, 25] links to interconnect the ToRs. These architectures enable a dynamic topology tuned to actual traffic patterns, and can provide a much higher aggregate capacity than a network of electronic switches at the same price point, consume significantly less power, and reduce cabling complexity. For instance, Farrington et al [15] report 2.8×, 6× and 4.7× lower cost, power, and cabling complexity respectively using optical circuit switching relative to a baseline network of electronic switches.

The drawback of circuit switches, however, is that their switching configuration time is much slower than electronic switches. Depending on the specific technology, reconfiguring the circuit switch can take a few milliseconds (e.g., for 3D MEMS optical circuit switches [46, 16, 53]) to tens of microseconds (e.g., for 2D MEMS wavelength-selective switches [40]). During this reconfiguration period, the circuit switch cannot carry any traffic. By contrast, electronic switches can make per-packet switching decisions at sub-microsecond timescales. This makes the circuit switch suitable for routing stable traffic or bursts of packets (e.g., hundreds to thousands of packets at a time), but not for sporadic traffic or latency sensitive packets. A natural approach is then to have a *hybrid* circuit/packet switch architecture: the circuit switch can handle traffic flows that have heavy intensity but also require sparse connections, while a lower capacity packet switch handles the complementary (low intensity, but densely connected) traffic flows [16].

With this hybrid architecture, the relatively low intensity traffic is taken care of by the packet switch — switch scheduling here can be done dynamically based on the traffic arrival and is a well studied topic [36, 28, 35]. On the other hand, scheduling the circuit switch, based on the heavy traffic demand matrix, is still a fundamental unresolved question. Consider an architecture where a centralized scheduler samples the traffic requirements at each of the ToR ports at regular intervals (W , of the order of 100μs–1ms), and looks to find the schedule of circuit switch configurations over the interval of W that is “matched” to the traffic requirements. The challenge is to balance the overhead of reconfiguring the circuits with the capability to be flexible and meet the traffic demand requirements.

The centralized scheduler must essentially decide a se-

quence of *matchings* between sending and receiving ToRs which the circuit switch then implements. For an optical circuit switch, for instance, the switch realizes the schedule by appropriately configuring its MEMS mirrors. As another example, in a broadcast-select optical ring architecture [9], the ToRs implement the controller’s schedule by tuning in to the appropriate wavelength to receive traffic from their matching sender as dictated by the schedule.

Hence, we need a scheduling algorithm that decides the state (i.e., matching) of the circuit switch at each time and also a routing protocol to decide on an appropriate (direct or indirect) route packets can take to reach their destination ToR port. This is a challenging problem and entails making several choices on: (a) number of matchings, (b) choice of matchings (switch configuration), (c) durations of the matchings and (d) the routing protocol, in each interval W . Mathematically, this leads to a well defined optimization problem, albeit involving both combinatorial and real-valued variables. Even special cases of this problem [31] are NP hard to solve exactly.

Central to understanding this scheduling problem is finding a good sparse representation of the traffic matrix – a fundamental algorithmic question in Carathéodory’s Theorem that has remained largely unanswered so far [38]. Recent papers have proposed heuristic algorithms to address this scheduling problem. In Solstice [34], the authors present a greedy perfect-matching based heuristic for a hybrid electrical-optical switch. Experimental evaluations show Solstice performing well over a simple baseline (where the schedules are provided by a truncated Birkhoff-von Neumann decomposition of the traffic matrix), although no theoretical guarantees are presented. Indirect routing in a distributed setting, but without considerations of configuration switching costs, is studied in another recent work [9].

1.1 Our Contributions

We first focus on routing policies where packets are sent from the source port to the destination port only via a direct link connecting the two ports, leading to *direct* or *single-hop* routing.

Approximate Carathéodory’s Theorem: Our main result here is an approximately optimal, very simple and fast algorithm for computing the switch schedule in each interval. In turn this corresponds to a fast algorithm for computing a sparse approximate representation of a point on the Birkhoff polytope [10]. While Carathéodory’s Theorem guarantees the *existence* of such a representation, an efficient *algorithm* to compute it has remained elusive so far. Our algorithm, which we christen Eclipse, has a performance that is at least half that of optimal for every instance of the traffic demands, and experimentally shows a strict and consistent improvement over the state-of-the-art [34]. A key technical contribution here is the identification of a submodularity structure [5] in the problem, which allows us to make connections between submodular function maximization and the circuit switch scheduling problem with reconfiguration delay.

Indirect Routing: Next, we consider routing policies where packets are allowed to reach their destination after (potentially) transiting through many intermediate ports, leading to *indirect* or *multi-hop* routing. This class of routing policies is motivated by our observation that if the number of matchings is limited, multi-hop routing can *exponentially* improve the *reachability* of nodes; a novel benefit of multi-hop routing

Figure 1: An illustration of our hybrid switch architecture.

is distinct from the classical and well known load balancing effects [42, 49, 23]. We again identify submodularity in the problem, but the constraints for this submodular maximization problem are no longer linear and efficient solutions are challenging to find. However, for the important special case where the sequence of switch configurations have already been calculated (and the indirect routing policy has to be decided) we propose a simple and fast greedy algorithm that is near optimal universally for all traffic requirements. Detailed simulation results demonstrate strong improvements over direct routing, which are especially pronounced when the switch reconfiguration delays are relatively large.

The paper is organized as follows. In Section 2, the model, framework and the problem objective are formally stated along with a succinct summary of the state of the art. Section 3 focuses on direct routing and Section 4 on indirect routing. In Section 5, we present a detailed evaluation of the proposed algorithms on a variety of traffic inputs. Section 6 closes with a brief discussion. Technical aspects of the algorithm and its evaluation, including connections to submodularity and combinatorial optimization problems are deferred to the Appendix.

2. SYSTEM MODEL

In this section, we present our model for a hybrid circuit-packet switched network fabric, and formally define our scheduling problem. Our model closely follows [34].

2.1 Hybrid Switch Model

We consider an n -port network where each port is simultaneously connected to a circuit switch and a packet switch as shown in Figure 1. A set of nodes are attached to the ports and communicate over the network. The nodes could either be individual servers or top-of-rack switches.

We model the circuit switch as an $n \times n$ crossbar comprising of n input ports and n output ports. At any point in time, each input port can send packets to at most one output port and each output port can receive packets from at most one input port over the circuit switch. The circuit switch can be reconfigured to change the input-output connections. We assume that the packets at the input ports are organized in virtual-output-queues [41] (VOQ) which hold packets destined to different output ports.

In practice, the circuit switch is typically an optical switch-

[46, 16, 53].² These switches have a key limitation: changing the circuit configuration imposes a *reconfiguration delay* during which the switch cannot carry any traffic. The reconfiguration delay can range from few milliseconds to tens of microseconds depending on the technology [40, 33]. This makes the circuit switch suitable for routing stable traffic or bursts of packets (e.g., hundreds to thousands of packets at a time), but not for sporadic traffic or latency sensitive packets. Therefore, hybrid networks also use a (electrical) packet switch to carry traffic that cannot be handled by the circuit switch. The packet switch operates on a packet-by-packet basis, but has a much lower capacity than the circuit switch. For example, the circuit and packet switches might respectively run at 100Gbps and 10Gbps per port.

We divide time into slots, with each slot corresponding to a (full-sized) packet transmission time on the circuit switch. We consider a *scheduling window* of $W \in \mathbb{Z}$ time units. A central controller uses measurements of the aggregated traffic demand between different ports to determine a *schedule* for the circuit switch at the start of each scheduling window. The schedule comprises of a sequence of configurations and how long to use each configuration (Section 2.3). The controller communicates the schedule to the circuit switch, which then follows the schedule for the next scheduling window (W) without involving the controller. We assume that the delay for each reconfiguration is $\delta \in \mathbb{Z}$ time units.

2.2 Traffic Demand

Let $T \in \mathbb{Z}^{n \times n}$ denote the accumulated traffic at the start of a scheduling window. We assume T is a feasible traffic demand, i.e., T is such that $\sum_{j=1}^n T(i, j) \leq W$ and $\sum_{i=1}^n T(i, j) \leq W$ for all $i, j \in \{1, 2, \dots, n\}$. The (i, j) th entry of T denotes the amount of traffic that is in the VOQ at node i destined for node j .

We assume that the controller knows T .³ We also assume that non-zero entries in the traffic matrix T are bounded as $2\delta \leq T(i, j) \leq \epsilon W$ for all $i, j \in [n] : T(i, j) > 0$ and some parameter $0 < \epsilon < 1$. This is a mild condition because traffic between pairs of ports that is small relative to δ is better served by the packet switch anyway.

Previous measurement studies have shown that the inter-rack traffic in production data centers is sparse [34, 8, 3, 43]. Over short periods of time (e.g., 10s of milliseconds), most nodes communicate with only a small number of other nodes (e.g., few to low tens). Further, in many cases, a large fraction of the traffic is sent by a small fraction of “elephant” flows [3]. While our algorithms and analysis are general, it is important to note that such sparse traffic patterns are necessary for hybrid networks to perform well (especially with larger reconfiguration delay).

2.3 The Scheduling Problem

Given the traffic demand, T , our goal is to compute a schedule that maximizes the total amount of traffic sent over the circuit switch during the scheduling window W . This is desirable to minimize the load on the slower packet switch. In general, the scheduling problem involves two aspects:

²Designs based on point-to-point wireless links have also been proposed [27, 56]. Our abstract model is general.

³Our work is orthogonal to how the controller obtains the traffic demand estimate. For example, the nodes could simply report their backlogs before each scheduling window, or a more sophisticated prediction algorithm could be used.

1. Determining a schedule of circuit switch configurations: The algorithm must determine a sequence of circuit switch configurations: $(\alpha_1, P_1), (\alpha_2, P_2), \dots, (\alpha_k, P_k)$. Here, $\alpha_i \in \mathbb{Z}$ denotes the duration of the i^{th} switch configuration, and P_i is an $n \times n$ permutation matrix, where $P_i(s, t) = 1$ if input port s is connected to output port t in the i^{th} configuration. For a valid schedule, we must have $\alpha_1 + \alpha_2 + \dots + \alpha_k + k\delta \leq W$ since the total duration of the configurations cannot exceed the scheduling window W .

2. Deciding how to route traffic: The simplest approach is to use only direct routes over the circuit switch. In other words, each node only sends traffic to destinations to which it has a direct circuit during the scheduling window. Alternatively, we can allow nodes to use indirect routes, where some traffic is forwarded via (potentially multiple) intermediate nodes before being delivered to the destination. Here, the intermediate nodes buffer traffic in their VOQs for transmission over a circuit in a subsequent configuration.

In the next section, we begin by formally defining the problem in the simpler setting with direct routing and developing an algorithm for this case. Then, in Section 4, we consider the more general setting with indirect routing.

Remark 1. Prior work [34, 31] has considered the objective of *covering* the entire traffic demand in the least amount of time. For example, the ADJUST algorithm in [31] takes the traffic demand T as input and computes a schedule $(\alpha_1, P_1), \dots, (\alpha_k, P_k)$ such that $\sum_{i=1}^k \alpha_i + k\delta$ is minimized while $\sum_{i=1}^k \alpha_i P_i \geq T$. Our formulation (and solution) is more general, since an algorithm which maximizes throughput over a given time period can also be used to find the shortest duration to cover the traffic demand (e.g., via binary search).

Remark 2. From a systems viewpoint, the traffic demand estimation can be done either by directly polling ToR switches or end-host NICs [33, 53], or indirectly through applications and flow information [16, 2]. For systems at scale, this represents a non-trivial task (and often taking 100s of μs [33]) necessitating the need for a window W to compute the schedule (vis-à-vis dynamic policies; see following Section 2.4).

2.4 Related Work

Before presenting the work in this paper, we briefly summarize related work on this topic. Scheduling in crossbar switches is a classical and well studied topic and traditionally it has been used to model the packet switch where the reconfiguration delay is very small. Hence the scheduling solutions proposed – ranging from centralized Birkhoff-von-Neumann decomposition scheduler [36] on one end to the decentralized load-balanced scheduler [11] on the other – did not account for reconfiguration delay.

With the proposals on hybrid circuit/packet switching systems [16, 53], simplified models that factor for the reconfiguration delay were considered. Early works often assumed the delay to be either zero [26] or infinity [48, 54]. The infinite delay setting corresponds to a problem where the number of matchings is minimized. However they still require $O(n)$ matchings. In a different context (satellite-switched time-division multiple access), works such as [22] also computed schedules that minimized the number of matchings. Moderate reconfiguration delays are considered in DOUBBLE [48] and other algorithms such as [19, 31, 55] that explicitly take reconfiguration delay into account. The al-

gorithm ADJUST [31] minimizes the covering time but still requires around n configurations. All of these algorithms do not benefit from sparse demands and continue to require $O(n)$ configurations [34]. In a complementary approach, [12] considers conditions on the input traffic matrix under which efficient polynomial time algorithms to compute the optimal schedule exists. Yet other approaches have been to introduce speedup [35], or randomization in the algorithms [21], however they do not address the basic optimization problem underlying this scenario head-on. Such is the goal of this paper.

The above algorithms are “batch” policies [51] in which each computational call returns a schedule for an entire window of time. Another research direction is to consider “dynamic” policies where scheduling decisions are made time-slot by time-slot. A variant of the well known MaxWeight algorithm is presented in [52] and is shown to be throughput optimal. Fixed-Frame MaxWeight (FFMW) is a frame based policy proposed in [32] and has good delay performance. However it requires the arrival statistics to be known in advance. A hysteresis based algorithm that adapts many previously proposed algorithms for crossbar switch scheduling to the case with reconfiguration delay is presented in [51]. All these algorithms require perfect queue state information at every instant.

3. DIRECT ROUTING

The centralized scheduler samples the ToR ports and arrives at the traffic demand (matrix) T to be met in the upcoming slot. In this section, we develop an algorithm, named Eclipse, that takes the traffic demand, T , as input and computes a schedule of matchings (circuit configurations) and their durations to maximize throughput over the circuit switch; only direct routing of packets from source to destination ports are allowed here. Eclipse is fast, simple and nearly-optimal in every instance of the traffic matrix T . Towards a formal understanding of the notion of optimality, consider the following optimization problem:

$$\begin{aligned} & \text{maximize} \quad \left\| \min \left(\sum_{i=1}^k \alpha_i P_i, T \right) \right\|_1 \\ & \text{s.t.} \quad \alpha_1 + \alpha_2 + \dots + \alpha_k + k\delta \leq W \\ & \quad k \in \mathbb{N}, P_i \in \mathcal{P}, \alpha_i \geq 0 \forall i \in \{1, 2, \dots, k\}, \end{aligned} \quad (1)$$

where $\mathbb{N} = \{1, 2, \dots\}$ and $\mathcal{P}$ is the set of permutation matrices.

This optimization problem is NP-hard [31], and a recent work [34] in the literature has focused on heuristic solutions. Our proposed algorithm has some similarities to the prior work in [34] in that the matchings and their durations are computed successively in a *greedy* fashion. However, the algorithm is overall quite different in terms of both ideas and details; we uncover and exploit the underlying *submodularity* [44] structure inherent in the problem to design and analyze the algorithm in a principled way.

We also note that this problem can be viewed as finding permutation matrices $P_1, \dots, P_k$ and weightings $\alpha_1, \dots, \alpha_k$ such that their weighted sum is a good *approximation* of the traffic matrix T . Carathéodory’s Theorem applied to the Birkhoff polytope guarantees the *existence* of such a dual representation; however till date we do not know of an efficient *algorithm* to compute this representation. A re-

Algorithm 1: A general greedy algorithm template

Input : Traffic demand T , reconfiguration delay δ and scheduling window size W
Output: Sequence of matchings $P_1, \dots, P_k$ and their corresponding durations $\alpha_1, \dots, \alpha_k$:
 $\text{sch} \leftarrow \{\}$; // schedule
 $k \leftarrow 0$;
 $T_{\text{rem}} \leftarrow T$; // traffic remaining
while $\sum_{i=1}^k (\alpha_i + \delta) \leq W$ **do**
 $k \leftarrow k + 1$;
 Decide on a duration α for the matching;
 $M \leftarrow \text{argmax}_{M \in \mathcal{M}} \|\min(\alpha M, T_{\text{rem}})\|_1$;
 $\text{sch} \leftarrow \text{sch} \cup \{(\alpha, M)\}$;
 $T_{\text{rem}} \leftarrow T_{\text{rem}} - \min(\alpha M, T_{\text{rem}})$;
end
if $\sum_{i=1}^k (\alpha_i + \delta) > W$ **then**
 $\text{sch} \leftarrow \text{sch} \setminus \{(\alpha, M)\}$;
end
 $k \leftarrow k - 1$;

cent work [6] proposes a fast approximation algorithm but they implicitly assume availability of a dual representation to start with. We do not assume such an oracle, but are nevertheless able to provide an efficient algorithm to obtain an approximate Carathéodory expansion.

3.1 Intuition

Before a formal presentation and analysis of the algorithm, we begin with an intuitive and less-formal approach to how one might solve this optimization problem. Consider greedy algorithms with the template shown in Algorithm 1. The template starts with an empty schedule, and proceeds to add a new matching to the schedule in each iteration. This process continues until the total duration of the matchings exceeds the allotted time budget of W , at which point the algorithm terminates and outputs the schedule computed so far. In each iteration, the algorithm first picks the duration of the matching, α . It then selects the *maximum weight matching* in the traffic graph whose edge weights are thresholded by α (i.e., edge weights $> \alpha$ are clipped to α). The traffic graph is a bipartite graph between n input and n output vertices, with an edge of weight $T(i, j)$ between input node i and output node j . It remains to specify how to choose α in each iteration.

Consider an exercise where we vary the matching duration α from 0 to W and compute the maximum weight matching in the thresholded traffic graph for each α . For a typical traffic matrix, this results in a curve similar to the solid-blue line in Fig. 2. Notice that the value of the maximum weight matching is precisely equal to the sum-throughput that can be achieved in that round of the switch schedule. It is straightforward to see that the maximum weight matching curve has the following properties: (a) it is non-decreasing and (b) piecewise linear. These are explained as follows: when α is very small a lot of the edges in the traffic graph have a weight that is saturated at α . Hence it is likely to find a perfect matching with total weight of $n\alpha$. As such the slope of the curve when α is small is n . However, as α becomes large there are increasingly fewer edges whose weights are saturated at α and, correspondingly, the slope reduces. When α is so large that all of the edge weights

Figure 2: Throughput of max. weight matching as a function of threshold duration. The effective utilization curve of the matchings is also shown.

are strictly smaller than α , then the value of the maximum weight matching does not change even with any further increase in α and the curve ultimately flattens out.

Two operating points of interest, considering Fig. 2, are (a) the largest α where the slope of the curve is maximum ($= n$ in the typical case where every ingress/egress port has traffic) and (b) the smallest α where the value of the maximum weight matching is the largest. These points have been denoted by α_1 and α_2 in Fig. 2 respectively. Setting $\alpha = \alpha_1$ is interesting because it results in a matching where the links are all fully utilized. For example, the Solstice algorithm presented in [34] implicitly adopts this operating point. On the other hand, $\alpha = \alpha_2$ gives a matching that achieves the largest possible sum-throughput in that round.

However we note that both choices of α are less than ideal for the following reasons. Recall that after every round of switching we incur a delay of δ time units. As such if the value of α_1 is small (say comparable to δ) in each round, then the number of matchings, and hence the time wasted due to the reconfiguration delay, becomes large. As a concrete example, consider the transpose of the traffic matrix $T_1 = [A_1^t \mathbf{b}_1^t]$ where A_1 is a sparse $(n-1) \times n$ matrix and $\mathbf{b} = [2\delta, 2\delta, \dots, 2\delta, 0, 0, \dots, 0]$ comprises of some k entries of value 2δ and $n-k$ entries of value 0. In other words, we are considering an input where a node or a collection of nodes have a large number of small flows to a particular node or vice-versa. For such an instance it is clear that if we insist on matchings with 100% utilized links, then the maximum duration of the matching is 2δ (i.e., $\alpha_1 = 2\delta$). Thus, continuing the process described in Algorithm 1 results in a sequence of k matchings each of which is only 2δ time units long. Hence in the worst case (if $k > 1/(3\delta)$) about 1/3rd of the entire scheduling window is wasted just due to reconfiguration delay limiting the maximum possible throughput to $2n/3$. On the other hand, if we had ignored the entries in $\mathbf{b}$, then we could have scheduled just A_1 achieving a total throughput of $n - 2k\delta \approx n$ for large n . We point out that the phenomenon described above happens in a large family of instances, of

which T_1 is a specific example. We also emphasize that such instances are pretty likely to occur in practice; for example, [43, Fig.5-b] shows traffic measurements in a Facebook data center where the interactions between Cache and Web servers lead to traffic matrices having this property.

Similarly for the operating point with $\alpha = \alpha_2$, consider the traffic matrix $T_2 = \begin{bmatrix} A_2 & \mathbf{0} \\ \mathbf{0} & B_2 \end{bmatrix}$ where A_2 is a sparse

$(n-2) \times (n-2)$ matrix and $B_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. This is a diametrically opposite situation from T_1 where a small collection of nodes interact only amongst themselves with no interaction outside. Such a situation occurs, for example, in multi-tenant cloud-computing data centers [45] where individual tenants run their jobs on small clusters of servers. In such a case, the value of the maximum weight matching can be maximum for a large α . For T_2 the maximum value occurs at $\alpha = 1 - \delta$ (i.e., $\alpha_2 = 1 - \delta$), resulting in a schedule with just one matching of duration $1 - \delta$ and potentially missing a lot of traffic for A_2 . For example, if A_2 is uniformly k -sparse, we miss out roughly $(k-1)n/k$ units of traffic. On the other hand, by choosing the duration of the matching to be $1/k - \delta$ in each step we can achieve a sum throughput of $n - O(\delta) \approx n$.

In scenarios exemplified by T_2 , setting $\alpha = 1$ is bad because the utilization of the resulting matching is poor, i.e., a vast majority of the matching links carry only a fraction of their capacity. This can be overcome by insisting that we choose only those matchings with utilization of at least 75% (say). However, in the case of T_1 we observe a poor performance in spite of all matchings having a utilization of 100%. The issue in this case is that the duration of the matchings are small compared to the reconfiguration delay. Hence to avoid this scenario we can insist on $\alpha \geq 20\delta$ (say) in Algorithm 1.

Our first main observation is that both of the above heuristics are captured if we consider the *effective utilization* of the matchings. We define effective utilization as the ratio $\text{mwm}(\alpha)/(\alpha + \delta)$ where $\text{mwm}(\alpha)$ denotes the value of the maximum weight matching at α . This ratio indicates the overall efficiency of a matching by including the reconfiguration delay into the duration. In Fig. 2 we plot the effective utilization of the matchings as the red-dotted curve. As can be seen there, the effective utilization at both α_1 and α_2 is suboptimal. We propose an algorithm that selects α to maximize effective utilization; a detailed description is deferred to Section 3.3.

The justification for selecting matchings according to the above is further reinforced by the *submodularity* structure of the problem (we discuss submodularity in Section 3.2). It turns out that for a certain class of submodular maximization problems with linear packing constraints, greedy algorithms take a form that precisely matches the intuitive thought process above [5]: the proposed intuitively correct algorithm is borne out naturally from submodular combinatorial optimization theory. We briefly recall relevant aspects of submodularity and associated optimization algorithms next.

3.2 Submodularity

A set function $f : 2^{[n]} \rightarrow \mathbb{R}$ is said to be *submodular* if it has the following property: for every $A, B \subseteq [n]$ we have $f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$. Alternatively,

submodular functions are also defined through the property of *decreasing marginal values*: for any S, T such that $T \subseteq S \subseteq [n]$ and $j \notin S$, we have

$$f(S \cup \{j\}) - f(S) \leq f(T \cup \{j\}) - f(T).$$

The difference $f(S \cup \{j\}) - f(S)$ is called the incremental marginal value of element j to set S and is denoted by $f_S(j)$. For our purpose we will only focus on submodular functions that are *monotone* and *normalized*, i.e., for any $S \subseteq T \subseteq [n]$ we have $f(S) \leq f(T)$ and further $f(\emptyset) = 0$.

Many applications in computer science involve maximizing submodular functions with *linear packing constraints*. This refers to problems of the form:

$$\max f(S) \quad \text{s.t. } A\mathbf{x}_S \leq b \text{ and } S \subseteq [n],$$

where $A \in [0, 1]^{m \times n}$, $b \in [1, \infty)^m$ and $\mathbf{x}_S$ denotes the characteristic vector of the set S . Each of the A_{ij} 's is a cost incurred for including element j in the solution. The b_i 's represent a total budget constraint. A well-known example of a problem in the above form is the Knapsack problem (the objective function in this case is in fact modular).

With the above background, we formulate the optimization problem under direct routing as one of submodular function maximization. Recall that for any given input traffic matrix T , the schedule that is computed is described by a sequence of matchings and corresponding durations. Consider the set $\mathcal{M}$ of all perfect matchings in the complete bipartite graph $K_{n \times n}$ with n nodes in each partite. Then any round in the schedule is simply $(\alpha, P) \in \mathbb{Z} \times \mathcal{M}$. The key observation we make now is to view the schedules as a subset of $\mathbb{Z} \times \mathcal{M}$. Formally, define a *switch schedule* as any subset $\{(\alpha_1, M_1), \dots, (\alpha_k, M_k)\}$ of $\mathbb{Z} \times \mathcal{M}$. The objective function in our case is the sum-throughput defined as

$$f(\{(\alpha_1, M_1), \dots, (\alpha_k, M_k)\}) = \left\| \min \left(\sum_{i=1}^k \alpha_i M_i, T \right) \right\|_1, \quad (2)$$

where the minimum is taken entrywise and $\|\cdot\|_1$ refers to the entrywise L_1 -norm of the matrix. We observe that the function f is submodular, deferring the proof to the Appendix.

THEOREM 1. *The function $f : 2^{\mathbb{Z} \times \mathcal{M}} \rightarrow \mathbb{R}$ defined by Equation (2) is a monotone, normalized submodular function.*

We have established that optical switch scheduling under the sum-throughput metric is a submodular maximization problem. With this, we are ready to present a greedy algorithm that achieves a sum-throughput of at least a constant factor of the optimal algorithm for *every* instance of the traffic matrix.

3.3 Algorithm

Algorithm 2 – Eclipse – captures our proposed solution under direct routing. Eclipse takes the traffic matrix T , the time window W and reconfiguration delay δ as inputs, and computes a sequence of matchings and durations as the output. The algorithm proceeds in rounds (the “while loop”), where in each round a new matching is added to the existing sequence of matchings. The sequence terminates whenever the sum of the matching durations exceeds the allocated time window W or whenever the traffic matrix T is fully covered.

Algorithm 2: Eclipse: greedy direct routing algorithm

Input : Traffic demand T , reconfiguration delay δ and scheduling window size W
Output: Sequence of matchings $P_1, \dots, P_k$ and their corresponding durations $\alpha_1, \dots, \alpha_k$:
 $\text{sch} \leftarrow \{\}$; // schedule
 $k \leftarrow 0$;
 $T_{\text{rem}} \leftarrow T$; // traffic remaining
while $\sum_{i=1}^k (\alpha_i + \delta) \leq W$ **do**
 $k \leftarrow k + 1$;
 $(\alpha, M) \leftarrow \operatorname{argmax}_{M \in \mathcal{M}, \alpha \in \mathbb{R}_+} \frac{\|\min(\alpha M, T_{\text{rem}})\|_1}{(\alpha + \delta)}$;
 $\text{sch} \leftarrow \text{sch} \cup \{(\alpha, M)\}$;
 $T_{\text{rem}} \leftarrow T_{\text{rem}} - \min(\alpha M, T_{\text{rem}})$;
end
if $\sum_{i=1}^k (\alpha_i + \delta) > W$ **then**
 $\text{sch} \leftarrow \text{sch} \setminus \{(\alpha, M)\}$;
end
 $k \leftarrow k - 1$;

Consider any round t in the algorithm; let $(\alpha_1, M_1), \dots, (\alpha_{t-1}, M_{t-1})$ denote the schedule computed so far in $t - 1$ rounds (stored in variable sch) and let $T_{\text{rem}}(t)$ denote the amount of traffic yet to be routed. The matching that is selected in the t -th round is the one for which *utilization* – the percentage of the total matching capacity that is actually used – is maximum. Mathematically, we choose an (α, M) pair such that $\frac{\|\min(\alpha M, T_{\text{rem}})\|_1}{\alpha + \delta}$ is maximized. In [50], an extended version of this paper, we have given a proof that the maximum (for α) occurs on the support of T_{rem} . Hence this can be easily found by looking at the support of the (sparse) matrix T_{rem} . We also propose a simple binary-search procedure, discussed in Algorithm 3, that finds only a local maximum but performs extremely well in our evaluations (Section 5). This process of selecting a matching is repeated in each round until the sum-duration of the matchings exceeds the scheduling window W , when the last chosen matching is discarded and the remaining set of matchings are returned. Eclipse is simple and also fast, a fact the following calculation demonstrates.

Complexity: We begin with the complexity of Algorithm 3. Since i_{ub} is no more than the number of distinct entries of T , we have $i_{\text{ub}} \leq n^2$. In each iteration, the algorithm only considers entries of H that have indices between i_{lb} and i_{ub} . However, binary-search halves the effective size of H (i.e., those numbers in H with array indices $i_{\text{lb}}, i_{\text{lb}+1}, \dots, i_{\text{ub}}$), and the number of iterations of the **while** loop is bounded by $\log n^2 = 2 \log n$. Within the **while** loop, computing the maximum weight matching can be done in $O(dn^{3/2} \log(W\epsilon))$ time (a basic fact of submodular optimization [14, 44]) where dn is the number of edges in bipartite graph formed by T (i.e., d is the average sparsity). Further $(1 - \epsilon)$ approximate maximum weight matching can be computed in linear time, e.g. $O(dn\epsilon^{-1} \log \epsilon^{-1})$ [13], and efficient implementations in practice have been studied extensively in the literature [37, 39, 17]. Hence the overall time complexity is $O(dn^{3/2} \log n \log(W\epsilon))$. Now, in Algorithm 2 the number of iterations in the **while** loop is bounded by W/δ . As such the total complexity of the algorithm is $\tilde{O}(dn^{3/2} \frac{W}{\delta})$. An exact search over the support of T_{rem} in the maximization step results in a overall complexity of $\tilde{O}(d^2 n^{5/2} \frac{W}{\delta})$.

Algorithm 3: Finding the greedy maximum

Input : Traffic demand T , reconfiguration delay δ
Output: $(\alpha, M) \in \mathbb{Z} \times \mathcal{M}$ such that
 $(\alpha, M) = \operatorname{argmax}_{M \in \mathcal{M}, \alpha \in \mathbb{Z}} \frac{\|\min(\alpha M, T)\|_1}{(\alpha + \delta)}$
 $H \leftarrow$ distinct entries of T sorted in ascending order;
 $i_{lb} \leftarrow 1$ and $i_{ub} \leftarrow \text{length}(H)$;
while $i_{lb} < i_{ub}$ **do**
 $i \leftarrow (i_{lb} + i_{ub})/2$;
 $T_1 \leftarrow \min\{T, H(i)\}$; // thresholding T to $H(i)$
 $T_2 \leftarrow \min\{T, H(i+1)\}$;
 $v_1 \leftarrow$ (max. weight matching in T_1)/($H(i) + \delta$) ;
 $v_2 \leftarrow$ (max. weight matching in T_2)/($H(i+1) + \delta$) ;
 if $v_1 < v_2$ **then**
 $i_{lb} \leftarrow i$;
 else if $v_1 > v_2$ **then**
 $i_{ub} \leftarrow i$;
 else
 return ($H(i)$, max. weight matching in T_1);
 end
end

Approximation Guarantee: Since the proposed direct routing algorithm is connected to submodular maximization with linear constraints, we can adapt standard combinatorial optimization techniques to show an approximation factor of $1 - 1/e$. Let **OPT** denote the sum-throughput of the optimal algorithm for given inputs T, δ and W . Let **ALG2** denote the sum-throughput achieved by Eclipse. We then have the following.

THEOREM 2. *If the entries of T are bounded by $\epsilon W + \delta$ then Eclipse approximates the optimal algorithm to within a factor of $1 - 1/e^{(1-\epsilon)}$, i.e., $\text{ALG2} \geq (1 - 1/e^{(1-\epsilon)})\text{OPT}$.*

The proof of the above Theorem is deferred to the Appendix. As a concluding remark, we note that the constant ϵ in the approximation factor comes from the requirement that $\alpha + \delta \leq \epsilon W$ hold. We observe that this mild technical condition, required to show that Eclipse is a constant factor approximation of the optimal algorithm, has an added implication. Informally, it ensures that no single matching occupies the bulk of the scheduling window.

4. INDIRECT ROUTING

In the previous section, we focused on direct routing where packets are forwarded to their destination ports only if a link *directly* connecting the source port to the destination port appeared in the schedule – this is essentially a “single-hop” protocol. In this section, we explore allowing packets to be forwarded to (potentially) multiple intermediate ports before arriving at its final destination. In terms of implementing this more involved protocol, we note that there is no extra overhead needed: the destination of any received packet is read first upon reception and since the queues are maintained on a per-destination basis at each ToR port, any received packet can be diverted to the appropriate queue. The key point of allowing indirect routing is the vastly increased range of ports that can be reached from a small number of matchings.

Consider Fig. 3 which illustrates a 6-port network and a sequence of 3 consecutive matchings in the schedule. With

Figure 3: Reachability of nodes under multi-hop routing.

direct routing, port 3 can only forward packets to ports 2, 5 and 4 in rounds 1, 2 and 3 respectively, i.e., the set of egress ports *reachable* by port 3 is $\{2, 4, 5\}$. In the indirect routing framework of this section, port 3 can also forward packets to port 1. This can be achieved by first forwarding the packets to port 2 in the first round where the packets are queued. Then in the second round we let port 2 forward those packets to the destination port 1. Thus the *reachability* of the nodes is enhanced by allowing for indirect routing. Indirect routing can also be viewed as “multi-hop” routing.

Traditionally multi-hop routing has been used as a means of *load balancing*. This is known to be true in the context of networks such as the Internet where the benefits of “Valiant load-balancing” are legion [42, 49, 23]. The benefits of load balancing are also well known in the switching context – a classic example is the two-stage load-balancing algorithm in crossbar switches without reconfiguration delay [47]. The benefit of multi-hop routing in our context is *markedly different*: the reachability benefits of indirect routing are especially well suited to the setting where input ports are directly connected to only a few output ports due to the *small number* of matchings in the scheduling window. In fact, an elementary calculation shows that over a period of k matchings in the schedule, indirect routing can allow a node to forward packets to $O(2^k)$ other nodes, compared to only $O(k)$ nodes possible with direct routing. This is because of the recursion $f(k) = 2f(k-1) + 1$ where $f(k)$ denotes the number of nodes reachable by any node in k rounds. If a node (say, node 1) can reach $f(k-1)$ nodes in $k-1$ rounds, then in the k -th round (i) there is a new node directly connected to node 1 and (ii) each of the $f(k-1)$ nodes can be connected to a new node. Thus the number of nodes connected to node 1 in the k -th round becomes $f(k-1) + (f(k-1) + 1)$. Fig. 3 also illustrates this phenomenon where reachability from node 3 is shown. As a corollary we observe that $O(\log_2 n)$ rounds of matchings are sufficient to reach *all* other nodes in a n -port network.

As in the direct routing case, computing the optimal schedule remains a challenging problem. While it is clear that we can achieve a performance at least as good as with direct routing, the gain is different for each instance of the traffic matrix – precisely quantifying the gain in an instance-specific way appears to be challenging. Our main result is that the submodularity property of the objective function continues to hold, provided the variables are considered in an appropriate format. Further, if we restrict some of the

variables (the matchings and their durations), then there is also a natural, simple and fast greedy algorithm to compute the switching schedules and routing policies that is approximately optimal for *each* instance of the traffic matrix. This algorithm serves as a heuristic solution to the more general problem of jointly finding the number of switchings, their durations, the switching schedules and routing policies. We present these results, following the same format as in the direct routing section, leaving numerical evaluations to a later section. We follow the model as discussed in Section 2.

4.1 Submodularity of objective function

We first adopt an alternative way of describing the switch schedule by specifying the multi-hop *path* taken by each packet. Such a formulation serves us well in the *causal* structure of the routed traffic patterns that naturally occur here.

For simplicity let us fix the number of rounds k in the schedule. Consider a fully connected k -round time-layered directed graph G consisting of $k+1$ partites, $V_0, V_1, \dots, V_k$ (of n nodes each), with nodes in each partite i having directed edges to all the nodes in partite $i+1$. Let $\mathcal{P}$ denote the set of all paths in G that begin at a node in V_0 and end at a node in V_k . Any such $p \in \mathcal{P}$ describes a multi-hop route for a packet in the system. If we are able to choose a path for every packet in the traffic matrix T , subject to capacity constraints, then we have a valid sequence of switch configurations and routing policy for the schedule. Now, for a set of paths $(\beta_1, p_1), \dots, (\beta_m, p_m)$, where β_i denotes the number of packets sharing the same path p_i , consider the sum-throughput given by a function $f : 2^{\mathbb{Z} \times \mathcal{P}} \rightarrow \mathbb{Z}$ defined as $f(\{(\beta_1, p_1), \dots, (\beta_m, p_m)\}) \triangleq$

$$\sum_{i,j \in [n]} \min \left(\sum_{l=1}^m \beta_l \mathbf{1}_{\substack{p_l(0)=i, \\ p_l(k+1)=j}}, T_{ij} \right), \quad (3)$$

where $p(0)$ and $p(k+1)$ denote the starting and ending nodes of path p and $\mathbf{1}_{\{\cdot\}}$ is the indicator function. Then the main observation is that f is submodular.

THEOREM 3. *The function $f : 2^{\mathbb{Z} \times \mathcal{P}} \rightarrow \mathbb{Z}$ defined by Equation (3) is submodular.*

The proof is analogous to Theorem 1 and we omit it in the interests of space.⁴ So far we have not imposed any restrictions on the set of paths that we choose for the schedule. This can be incorporated in the form of constraints to the problem, thus rephrasing the objective as a constrained submodular maximization problem.

Constraints: Since we are only choosing weighted paths, we can write constraints that ensure (i) the set of paths form a matching in each round and (ii) the total duration of the matchings is at most $W - k\delta$ [50]. However, the key challenge here is that the constraints are *nonlinear* – it is not clear whether an efficient (approximation) algorithm exists. The nonlinearities appear only in the sense of membership tests and a corresponding thresholding function – so it is possible that an efficient nearly-optimal greedy algorithm exists, but we leave this study for future work. We do note, however, that for the special case in which the configurations are fixed and we only have to decide on the indirect routing policies, the constraints take on a linear form – in this setting, we are able to construct fast and efficient greedy algorithms. This

⁴The proof can be found in [50].

case represents a composition of direct routing (where switch schedules are computed) and indirect routing (where the multi-hop routing policies are described), and is discussed next.

Multi-Hop Routing Policies: Consider a fixed sequence $(\alpha_1, M_1), \dots, (\alpha_k, M_k)$ of switch configurations and an input traffic demand matrix T . Let G denote the time-layered edge-capacitated graph obtained from the sequence of matchings, i.e., G consists of $k+1$ partites $V_0, \dots, V_k$ with n nodes each, and M_i is the matching between partites V_{i-1} and V_i with edge capacity α_i on the matching edges. In addition to the matching edges, there are also edges, with unlimited edge capacities, connecting the j -th nodes of V_{i-1} and V_i for all $j \in [n], i \in [k]$. Let $R(e)$ denote the capacity of edge $e \in G$. In this setting, the capacity constraints on the end-to-end paths are the sole constraints in the optimization problem – we consider subsets $\{(\beta_1, p_1), \dots, (\beta_m, p_m)\}$ that obey

$$\sum_{i=1}^m \beta_i \mathbf{1}_{\{e \in p_i\}} \leq R(e) \quad \forall e \in G. \quad (4)$$

Notice that the constraints above have a linear form, and there are a total of kn such constraints (one for each edge). Such a setting allows for a natural, simple, fast and nearly-optimal algorithm which we discuss below. Prior to that discussion, we remark that a naive approach to resolve the setting here is to formulate a *linear program* that maximizes the required objective. Indeed linear programming based approaches was the predominant technique used to solve this classical *multicommodity flow* problem [1, 24, 7]. However, despite many years of research in this direction the proposed algorithms were often too slow even for moderate sized instances [30]. Since then there has been a renewed effort in providing efficient *approximate* solutions to the multicommodity flow problem [20, 4]. The algorithm we propose is also a step in this direction, favoring efficiency over exactness of the solution. Note that in the path-formulation of the linear program there can be an exponential (in k) number of variables. Thus, in this form, there is no obvious efficient (exact) solution to this LP. On the other hand, the formulation we work with is able to handle this issue by appropriately *weighing the edges* and *picking the path* with the smallest weight; this latter step can be done efficiently – for example, using Dijkstra’s algorithm.

4.2 Algorithm

We propose Eclipse++ (Algorithm 4) directly motivated from [5], which presents a fast and efficient multiplicative weights algorithm for submodular maximization under linear constraints. The structure of Algorithm 4 is similar in spirit to Eclipse (Algorithm 2) in the sense that (a) the algorithm proceeds in rounds, where one new path is added to the schedule in each round and (b) we select a path that offers the greatest utility per unit of cost incurred. However, unlike Algorithm 2 where there was only one linear constraint, we have multiple linear constraints now. This is addressed by assigning weights to the constraints and considering a *linear combination* of the costs as the true cost in each round. In the following, we describe the salient features of Eclipse++.

Recall the capacity constraints in Equation (4) for each edge $e \in G$; let w_e denote the weight assigned to the constraint involving edge e . We set $w_e = 1/R(e)$ for all e ini-

Algorithm 4: Eclipse++ : greedy indirect routing algorithm

Input : Traffic demand T , switch configurations with residue capacities $R_1, \dots, R_k$, update factor λ
Output: Sequence of paths $p_1, \dots, p_m$ and corresponding weights $\beta_1, \dots, \beta_m$

```

sch ← {}; // schedule
Trem ← T; // traffic remaining
we ← 1/R(e) for all e ∈ E;
m ← 1;
while ∑e∈E R(e)we ≤ λ and ||Trem||1 > 0 do
    (βm, pm) ← argmaxp∈P, β∈Z  $\frac{\min(\beta, T_{\text{rem}}(p(0), p(k+1)))}{\sum_{e \in E} \beta \mathbf{1}_{\{e \in p\}} w_e}$ ;
    sch ← sch ∪ {(βm, pm)};
    Trem(pm(0), pm(k+1)) ← Trem(pm(0), pm(k+1)) - β;
    ;
    we ← weλβm1{e∈p}/R(e)} ∀ e ∈ G;
    m ← m + 1;
end
if ∑i=1m-1 βi1{e∈pi} ≤ R(e) ∀ e ∈ E then
    return sch
else
    return sch \ (βm-1, pm-1)
end

```

tially, i.e., edges with a large capacity are assigned a small weight and vice-versa. We can now have another graph G_w (with same topology as G) whose edges are weighted by w_e . Now, for any path p the “effective cost” of the path per packet is simply the total cost of p in G_w . Thus for the path (β, p) carrying β packets, the effective cost is given by $\sum_{e \in E} \beta w_e \mathbf{1}_{\{e \in p\}}$. On the other hand, the benefit we get due to adding path (β, p) is given by $\min(\beta, T(p(0), p(k+1)))$ where $p(0)$ and $p(k+1)$ stand for the starting and terminating nodes along path p . Thus, the ratio $\frac{\min(\beta, T(p(0), p(k+1)))}{\sum_{e \in E} \beta w_e \mathbf{1}_{\{e \in p\}}}$ denotes the benefit of path p per unit cost incurred. In Algorithm 4 we select p such that the utility per unit cost is maximized.

Now, once we have selected a weighted path (β_1, p_1) in the first round, we update the weights w_e on the edges. This is done as $w_e \leftarrow w_e \lambda^{\beta_1 / R(e)}$ for each edge $e \in p$, where λ is an input parameter. For the remaining edges the weights remain unchanged. Thus repeating the above iteratively until the **while** loop condition $\sum_{e \in E} R(e)w_e \leq \lambda$ becomes invalid, we get a schedule that is the output of the algorithm. It can also be shown that if the schedule returned **sch** violates any of the constraints (Equation (4)) then it must have happened at the very last iteration and hence we return a schedule with the last added path removed from it. A detailed analysis and correctness of the algorithm proposed can be found in [50].

It only remains to show how the maximizer of

$$\frac{\min(\beta, T_{\text{rem}}(p(0), p(k+1)))}{\sum_{e \in E} \beta \mathbf{1}_{\{e \in p\}} w_e} \quad (5)$$

is computed efficiently in each round (first line inside the **while** loop). Consider the set of shortest paths in G_w (smallest w_e -weighted path) from vertices in V_0 to vertices in V_k . Let p^* denote the shortest among them. Then by setting $\beta^* \leftarrow T_{\text{rem}}(p^*(0), p^*(k+1))$ we claim that Equation (5) is

maximized. This is because,

$$\begin{aligned} \frac{\min(\beta, T_{\text{rem}}(p(0), p(k+1)))}{\sum_{e \in E} \beta \mathbf{1}_{\{e \in p\}} w_e} &\leq \frac{\beta}{\sum_{e \in E} \beta \mathbf{1}_{\{e \in p\}} w_e} \\ &\leq \frac{1}{\min \sum_{e \in E} \mathbf{1}_{\{e \in p\}} w_e}. \end{aligned}$$

If $T_{\text{rem}}(p^*(0), p^*(k+1)) = 0$ we proceed to the second smallest shortest path and so on. This allows a very efficient implementation of the internal maximization step.

Approximation Guarantee: We show, as in the direct-routing scenario, that Eclipse++ has a constant factor approximation guarantee. Specifically, for a fixed instance of the traffic matrix, let **OPT** and **ALG4** denote the value of the objectives achieved by the optimal algorithm and Eclipse++ respectively. Let $\eta := \max_{i,j \in [n], e \in E} T(i, j)/R(e)$. Then one can show that $\text{ALG4} = \Omega(1/(nk)^\eta) \text{OPT}$ for $\lambda = e^{1/\eta} nk$; the proof is analogous to the direct-routing case and follows [5, Theorem 1.1]. Further, if $\eta = O(\epsilon^2 / \log(nk))$ for some fixed $\epsilon > 0$ then we get a approximation ratio of $(1 - \epsilon)(1 - 1/e)$ by letting $\lambda = e^{\epsilon/(4\eta)}$ (using [5, Theorem 1.2]). An interesting regime where this occurs is when the traffic matrices are dense with small skew. For example, we get a constant factor approximation if the sparsity of the traffic matrix grows at least logarithmically fast. This is in stark contrast to direct routing, where sparse matrices generally perform better.

Complexity: The proposed algorithm is simple and fast. In this subsection, we explicitly enumerate the time complexity of the full algorithm and show that the complexity is at most cubic in n and nearly linear in k . Let W denote a bound on the total incoming or outgoing traffic for a node. In each iteration of the **while** loop at least one packet is sent. Therefore there are at most W iterations of the **while** loop. Now, in each iteration finding the shortest paths between nodes in V_0 to nodes in V_k takes $kn^2(\log k + \log n)$ operations using Dijkstra’s algorithm [18]. Sorting the computed distances takes $kn^2(\log k + \log n)^2$ time and at most n^2 more operations to find a pair i, j such that $T_{\text{rem}}(i, j) > 0$. Finally the weights update step takes kn time. Therefore overall it takes $O(kn^2(\log k + \log n)^2)$ time per iteration. Hence the time complexity of the complete algorithm is $O(Wkn^3(\log k + \log n)^2)$.

5. EVALUATION

In this section, we complement our analytical results with numerical simulations to explore the effectiveness of our algorithms and compare them to state-of-the-art techniques in the literature. We empirically evaluate both the direct routing algorithm (Eclipse; Algorithm 2) and the indirect routing algorithm (Eclipse++; Algorithm 4).

Metric: We consider the total fraction of traffic delivered via the circuit switch (sum-throughput) over the duration of a fixed scheduling window as the performance metric throughout this section. Evaluating our algorithms under continuous traffic arrival models remains an important future direction.

Schemes compared: Our experiments compare Eclipse against two existing algorithms for direct routing:

Solstice [34]: This is the state-of-the-art hybrid circuit/packet scheduling algorithm for data centers. The key idea in Solstice is to choose matchings with 100% utilization. This is achieved by thresholding the demand matrix and selecting

Figure 4: Performance comparison of Eclipse under single-block inputs.

a perfect matching in each round. The algorithm presented in [34] tries to minimize the total duration of the window such that the entire traffic demand is covered. In this paper, we have considered a more general setting where the scheduling window W is constrained. To compare against Solstice in this setting, we truncate its output once the total schedule duration exceeds W .

Truncated Birkhoff-von-Neumann (BvN) decomposition [10]: The second algorithm we compare against is the truncated BvN decomposition algorithm. BvN decomposition refers to expressing a doubly stochastic matrix as a convex combination of permutation matrices and this decomposition procedure has been extensively used in the context of packet switch scheduling [34, 29, 10]. However BvN decomposition is oblivious to reconfiguration delay and can produce a potentially large ($O(n^2)$) number of matchings. Indeed, in our simulations BvN performs poorly.

Indirect routing is relatively new and to the best of our knowledge our work is the first to consider use of indirect routing for centralized scheduling.⁵ In our second set of simulations, we show that the benefits of indirect routing are in addition to the ones obtained from switch configurations scheduling. To this end, we compare Eclipse with Eclipse++ to quantify the additional throughput obtained by performing indirect routing (Algorithm 4) on a schedule that has been (pre)computed using Eclipse.

Traffic demands: We consider two classes of inputs: (a) single-block inputs and (b) multi-block inputs (explained in Section 5.1). Intuitively, single-block inputs are matrices which consist of one $n \times n$ ‘block’ that is sparse and skewed, and are similar to the traffic demands evaluated in the Solstice paper [34]. Multi-block inputs, on the other hand, denote traffic matrices that are composed of many sub-matrices each with disparate properties such as sparsity and skew.

Network size: The number of ports is fixed in the range of 50–200. We find that the relative performances stayed numerically stable over this range as well as for increased number of ports.

⁵Indirect routing in a distributed setting but without consideration of switch reconfiguration delay was studied in a recent work [9].

5.1 Direct Routing

While maintaining the sum-throughput as the performance metric, we vary the various parameters of the system model to gauge the performance in different situations.

Single-block Inputs

For a single-block input, our simulation setup consists of a network with 100 ports. The link rate of the circuit switch is normalized to 1, and the scheduling window length is also 1 ($W = 1$). We consider traffic inputs where the maximum traffic to or from any port is bounded by W . Further, we let the reconfiguration delay $\delta = W/100$. The traffic matrix is generated similar to [34] as follows. We assume 4 large flows and 12 small flows to each input or output port. The large flows are assumed to carry 70% of the link bandwidth, while the small flows deliver the remaining 30% of the traffic. To do this, we let each flow be represented by a random weighted permutation matrix, i.e., we have

$$T = \sum_{i=1}^{n_L} \frac{c_L}{n_L} P_i + \sum_{i'=1}^{n_S} \frac{c_S}{n_S} P_{i'} + N, \quad (6)$$

where $n_L(n_S)$ denotes the number of large (small) flows and $c_L(c_S)$ denotes the total percentage of traffic carried by the large (small) flows. In this case, we have $n_L = 4, n_S = 12$ and $c_L = 0.7, c_S = 0.3$. Further, we have added a small amount of noise N — additive Gaussian noise with standard deviation equal to 0.3% of the link capacity — to the non-zero entries to introduce some perturbation. Each experiment below has been repeated 25 times.

Reconfiguration delay: In Fig. 4(a) we plot sum-throughput while varying the reconfiguration delay from $W/3200$ to $4W/100$. Eclipse achieves a throughput of at least 90% for $\delta \leq W/100$. We observe Eclipse to be consistently better than Solstice although the difference is not pronounced until $\delta > W/100$. The BvN decomposition algorithm has a large throughput when the reconfiguration delay is small. As δ increases, its performance gradually worsens.

Skew: We control the skew by varying the ratio of the amount of traffic carried by small and large flows in the input traffic demand matrix (c_L/c_S in Equation (6)). Fig. 4(b) captures the scenario where the percentage traffic carried by the small flows is varied from 5 to 75. We observe that Eclipse is very robust to skew variations and is able to con-

Figure 5: Performance comparison of Eclipse under multi-block inputs.

sistently maintain a throughput of about 85%. Solstice has a slightly better performance at low skew (when small-flows carry $\sim 75\%$ of traffic); but overall, is dominated by Eclipse. **Sparsity:** Finally, we tested the algorithms' dependence on sparsity and plotted the results in Fig. 4(c). The total number of flows is varied from 4 to 32, while fixing the ratio of the number of large to small flows at 1:3. As the input matrix becomes less sparse, the performance of algorithms degrade as expected. However, for Eclipse, the reduction in the throughput is never more than 10% over the range of sparsity parameters considered. Solstice, on the other hand, is affected much more severely by decreased sparsity.

Multi-Block Inputs

Next, we consider a more complex traffic model for a 200 node network with block diagonal inputs of the form

$$T = \begin{bmatrix} B_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & B_m \end{bmatrix},$$

where each of the component blocks $B_1, B_2, \dots, B_m$ can have its own sparsity (number of flows) and skew (fraction of traffic carried by large versus small flows) parameters. The different blocks model the traffic demands of different tenants in a shared data center network such as a public cloud data center. To begin with, we consider inputs with two blocks $T = \begin{bmatrix} B_1 & \mathbf{0} \\ \mathbf{0} & B_2 \end{bmatrix}$ where B_1 is a $n_1 \times n_1$ matrix with 4 large flows (carrying 70% of the traffic) and 12 small flows (carrying 30% of the traffic) and B_2 is a $(200 - n_1) \times (200 - n_1)$ matrix with uniform entries (up to sampling noise).

Size of block: Fig. 5(a) plots the throughput as the block size of B_2 is increased from 0 to 70. We observe a very pronounced difference in the performance of Eclipse and Solstice: Eclipse has roughly $1.5 - 2\times$ the performance of Solstice. These findings are in tune with the intuition discussed in Section 3.1 — the deteriorated performance of Solstice is due its insistence on perfect matchings in each round.

Reconfiguration delay: Fig. 5(b) plots throughput while varying the reconfiguration delay, for fixed size of B_2 to be 50×50 . As expected, the throughput of Solstice and Eclipse both degrade as the reconfiguration delay δ increases. However, Eclipse throughput degrades at a much slower rate than Solstice. The gap between the two is particularly pro-

nounced for $\delta/W \geq 0.02$, a numerical value that is well within range of practical system settings.

Varying numbers of flow: In the final experiment we consider block diagonal inputs with 8 blocks of size 25×25 each. Each block carries $10 + \lfloor \sigma \cdot (U - 0.5) \rfloor$ equi-valued flows where $U \sim \text{unif}(0, 1)$ and σ is a parameter that controls the variation in the number of flows. When increasing σ from 0 to 20 we see from Fig. 5(c) that Eclipse is more or less able to sustain its throughput at close to 80%; whereas Solstice is significantly affected by the variation.

5.2 Indirect Routing

In this section, we consider a 50 node network with traffic matrices having varying number of large and small flows as before. We compare the performance of the direct routing algorithm and the indirect routing algorithm that is run on the schedule computed by Eclipse. To understand the benefits of indirect routing, we focus on the regime where the reconfiguration delay δ/W is relatively large and the scheduling window W is relatively long compared to the traffic demand. This regime corresponds to realistic scenarios where the circuit switch is not fully utilized (Real data center networks often have low to moderate utilization; e.g. 10–50% [8]), but the reconfiguration delay is large. In this setting of relatively large δ/W , switch schedules are forced to have only a small number of matchings, and indirect routing is critical to support (non-sparse) demand matrices. The following experiments numerically demonstrate the added gains of indirect routing.

Sparsity: Fig. 6(a) considers a demand with 5 large flows and number of small flows varying from 7 to 49. The large and the small flows each carry 50% of the traffic. We let $\delta = 16W/100$ and consider a load of 20% (i.e., $W = 5$, and traffic load at each port is 1). We observe that the performance of the Eclipse++ is roughly 10% better than Eclipse.

Load: As the network load increases (Fig. 6(b)), we see that indirect routing becomes less effective. This is because at high load, the circuits do not have much spare capacity to support indirect traffic. However, at low to moderate levels of load, indirect routing provides a notable throughput gain over direct routing. For example, we see close to 20% improvement with Eclipse++ over Eclipse at 15% load.

Reconfiguration delay: Finally, we observe the effect of

Figure 6: Performance of Eclipse++ and Eclipse. Here Eclipse++ uses the schedule computed by Eclipse.

δ/W on throughput by varying δ from $3W/100$ to $21W/100$. At smaller values of reconfiguration delay δ both Eclipse and Eclipse++ are able to achieve near 100% throughput. With increasing δ both algorithms degrade with Eclipse++ providing an additional gain of roughly 20% over Eclipse.

6. FINAL REMARKS

We have studied scheduling in hybrid switch architectures with reconfiguration delays in the circuit switch, by taking a fundamental and first-principles approach. The connections to submodular optimization theory allows us to design simple and fast scheduling algorithms and show that they are near optimal — these results hold in the direct routing scenario and indirect routing provided switch configurations are calculated separately. The problem of jointly deciding the switch configurations and indirect routing policies remains open. While submodular function optimization with nonlinear constraints is in general intractable, the specific constraints discussed in Section 4.1 perhaps have enough structure that they can be handled in a principled way; this is a direction for future work.

In between the scheduling windows of W time units, traffic builds up at the ToR ports. This dynamic traffic buildup is known locally to each of the ToR ports and perhaps this local knowledge can be used to pick appropriate indirect routing policies in a distributed, dynamic fashion. Such a study of indirect routing policies is initiated in a recent work [9], but this work omitted the switching reconfiguration delays. A joint study of distributed dynamic traffic scheduling in conjunction with a static schedule of switch configurations that account for reconfiguration delays is also an interesting direction of future research.

7. ACKNOWLEDGMENTS

The authors would like to thank Prof. Chandra Chekuri, Prof. George Porter and Prof. R. Srikant for the many helpful discussions. This work was partially supported by NSF grant CCF-1409106 and Army grant W911NF-14-1-0220.

8. REFERENCES

- [1] R. Ahuja, T. Magnanti, and J. Orlin. *Network flows: theory, algorithms, and applications*. Prentice Hall, 1993.
- [2] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. Hedera: Dynamic flow scheduling for data center networks. In *NSDI*, volume 10, pages 19–19, 2010.
- [3] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan. Data center tcp (dctcp). *SIGCOMM*, 2011.
- [4] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- [5] Y. Azar and I. Gamzu. Efficient submodular function maximization under linear packing constraints. In *Automata, Languages, and Programming*, pages 38–50. Springer, 2012.
- [6] S. Barman. Approximating nash equilibria and dense bipartite subgraphs via an approximate version of caratheodory’s theorem. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 361–369. ACM, 2015.
- [7] C. Barnhart and Y. Sheffi. A network-based primal-dual heuristic for the solution of multicommodity network flow problems. *Transportation Science*, 27(2):102–117, 1993.
- [8] T. Benson, A. Akella, and D. A. Maltz. Network traffic characteristics of data centers in the wild. In *SIGCOMM*, 2010.
- [9] Z. Cao, M. Kodialam, and T. Lakshman. Joint static and dynamic traffic scheduling in data center networks. In *INFOCOM*, 2014.
- [10] C.-S. Chang, W.-J. Chen, and H.-Y. Huang. Birkhoff-von neumann input buffered crossbar switches. In *INFOCOM*, 2000.
- [11] C.-S. Chang, D.-S. Lee, and Y.-S. Jou. Load balanced birkhoff-von neumann switches. In *High Performance Switching and Routing, 2001 IEEE Workshop on*, pages 276–280. IEEE, 2001.
- [12] A. Dasylyva and R. Srikant. Optimal wdm schedules for optical star networks. *IEEE/ACM Transactions on Networking (TON)*, 7(3):446–456, 1999.
- [13] R. Duan and S. Pettie. Linear-time approximation for maximum weight matching. *J. ACM*, 61(1):1:1–1:23, 2018.

- Jan. 2014.
- [14] R. Duan and H.-H. Su. A scaling algorithm for maximum weight matching in bipartite graphs. In *SODA*, 2012.
 - [15] N. Farrington. *Optics in data center network architecture*. PhD thesis, Citeseer, 2012.
 - [16] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat. Helios: a hybrid electrical/optical switch architecture for modular data centers. *SIGCOMM*, 2011.
 - [17] P. F. Felzenszwalb and R. Zabih. Dynamic programming and graph algorithms in computer vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(4):721–740, 2011.
 - [18] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM (JACM)*, 34(3):596–615, 1987.
 - [19] S. Fu, B. Wu, X. Jiang, A. Pattavina, L. Zhang, and S. Xu. Cost and delay tradeoff in three-stage switch architecture for data center networks. In *HPSR*, 2013.
 - [20] N. Garg and J. Koenemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *SIAM Journal on Computing*, 37(2):630–652, 2007.
 - [21] P. Giaccone, B. Prabhakar, and D. Shah. Randomized scheduling algorithms for high-aggregate bandwidth switches. *Selected Areas in Communications, IEEE Journal on*, 21(4):546–559, 2003.
 - [22] I. S. Gopal and C. K. Wong. Minimizing the number of switchings in an ss/tdma system. *Communications, IEEE Transactions on*, 33(6):497–501, 1985.
 - [23] A. Greenberg, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta. Towards a next generation data center architecture: scalability and commoditization. In *Proceedings of the ACM workshop on Programmable routers for extensible services of tomorrow*, pages 57–62. ACM, 2008.
 - [24] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*. Algorithms and combinatorics. Springer-Verlag, 1993.
 - [25] N. Hamedazimi, Z. Qazi, H. Gupta, V. Sekar, S. R. Das, J. P. Longtin, H. Shah, and A. Tanwer. Firefly: a reconfigurable wireless data center fabric using free-space optics. In *SIGCOMM*, 2014.
 - [26] T. Inukai. An efficient ss/tdma time slot assignment algorithm. *Communications, IEEE Transactions on*, 27(10):1449–1455, 1979.
 - [27] S. Kandula, J. Padhye, and P. Bahl. Flyways to de-congest data center networks. 2009.
 - [28] I. Keslassy, C.-S. Chang, N. McKeown, and D.-S. Lee. Optimal load-balancing. In *INFOCOM*, 2005.
 - [29] I. Keslassy, M. Kodialam, T. Lakshman, and D. Stiliadis. On guaranteed smooth scheduling for input-queued switches. In *INFOCOM*, 2003.
 - [30] T. Leighton, F. Makedon, S. Plotkin, C. Stein, É. Tardos, and S. Tragoudas. Fast approximation algorithms for multicommodity flow problems. *Journal of Computer and System Sciences*, 50(2):228–243, 1995.
 - [31] X. Li and M. Hamdi. On scheduling optical packet switches with reconfiguration delay. *Selected Areas in Communications, IEEE Journal on*, 21(7):1156–1164, 2003.
 - [32] Y. Li, S. Panwar, and H. J. Chao. Frame-based matching algorithms for optical switches. In *High Performance Switching and Routing, 2003, HPSR. Workshop on*, pages 97–102. IEEE, 2003.
 - [33] H. Liu, F. Lu, A. Forencich, R. Kapoor, M. Tewari, G. M. Voelker, G. Papen, A. C. Snoeren, and G. Porter. Circuit switching under the radar with reactor. In *NSDI*, 2014.
 - [34] H. Liu, M. K. Mukerjee, C. Li, N. Feltman, G. Papen, S. Savage, S. Seshan, G. M. Voelker, D. G. Andersen, M. Kaminsky, G. Porter, and A. C. Snoeren. Scheduling techniques for hybrid circuit/packet networks. In *ACM CoNEXT*, 2015.
 - [35] N. McKeown. The islip scheduling algorithm for input-queued switches. *Networking, IEEE/ACM Transactions on*, 7(2):188–201, 1999.
 - [36] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. *Communications, IEEE Transactions on*, 47(8):1260–1267, 1999.
 - [37] A. Mekkittikul and N. McKeown. A practical scheduling algorithm to achieve 100% throughput in input-queued switches. In *INFOCOM*, 1998.
 - [38] V. Mirrokni, R. P. Leme, A. Vladu, and S. C.-w. Wong. Tight bounds for approximate carathéodory and beyond. *arXiv preprint arXiv:1512.08602*, 2015.
 - [39] S. Pettie and P. Sanders. A simpler linear time $2/3 - \epsilon$ approximation for maximum weight matching. *Information Processing Letters*, 91(6):271–276, 2004.
 - [40] G. Porter, R. Strong, N. Farrington, A. Forencich, P. Chen-Sun, T. Rosing, Y. Fainman, G. Papen, and A. Vahdat. Integrating microsecond circuit switching into the data center. *SIGCOMM*, 2013.
 - [41] B. Prabhakar and N. McKeown. On the speedup required for combined input-and output-queued switching. *Automatica*, 35(12):1909–1920, 1999.
 - [42] M. O. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of the ACM (JACM)*, 36(2):335–348, 1989.
 - [43] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren. Inside the social network’s (datacenter) network. In *SIGCOMM*, 2015.
 - [44] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
 - [45] A. Shieh, S. Kandula, A. G. Greenberg, and C. Kim. Seawall: Performance isolation for cloud datacenter networks. In *HotCloud*, 2010.
 - [46] A. Singla, A. Singh, and Y. Chen. Osa: An optical switching architecture for data center networks with unprecedented flexibility. In *NSDI*, 2012.
 - [47] R. Srikant and L. Ying. *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, New York, NY, USA, 2014.
 - [48] B. Towles and W. J. Dally. Guaranteed scheduling for switches with configuration overhead. *Networking, IEEE/ACM Transactions on*, 11(5):835–847, 2003.

- [49] L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- [50] S. B. Venkatakrishnan, M. Alizadeh, and P. Viswanath. Costly circuits, submodular schedules: Hybrid switch scheduling for data centers. *CoRR*, abs/1512.01271, 2015.
- [51] C.-H. Wang and T. Javidi. Adaptive policies for scheduling with reconfiguration delay: An end-to-end solution for all-optical data centers. *arXiv preprint arXiv:1511.03417*, 2015.
- [52] C.-H. Wang, T. Javidi, and G. Porter. End-to-end scheduling for all-optical data centers. In *Computer Communications (INFOCOM), 2015 IEEE Conference on*, pages 406–414. IEEE, 2015.
- [53] G. Wang, D. G. Andersen, M. Kaminsky, K. Papagiannaki, T. Ng, M. Kozuch, and M. Ryan. c-through: Part-time optics in data centers. *SIGCOMM*, 2011.
- [54] B. Wu and K. L. Yeung. Nxg05-6: Minimum delay scheduling in scalable hybrid electronic/optical packet switches. In *GLOBECOM*, 2006.
- [55] B. Wu, K. L. Yeung, and X. Wang. Nxg06-4: Improving scheduling efficiency for high-speed routers with optical switch fabrics. In *GLOBECOM*, 2006.
- [56] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. *SIGCOMM*, 2012.

APPENDIX

A. DIRECT ROUTING - PROOFS

A.1 Proof of Theorem 1

PROOF. Clearly $f(\{\}) = 0$. Also for any $S \subseteq S' \in 2^{\mathbb{Z} \times \mathcal{M}}$ we have $\min \left\{ \sum_{(\alpha, M) \in S} \alpha M, T \right\} \leq \min \left\{ \sum_{(\alpha, M) \in S'} \alpha M, T \right\}$ implying $f(S) \leq f(S')$. Hence f is normalized and monotone. Next, for non-negative reals a, b, c the identity $\min(a + b, c) = \min(a, c) + \min(b, c - \min(a, c))$ holds. Therefore for $S \in 2^{\mathbb{Z} \times \mathcal{M}}$ and $(\alpha_0, M_0) \notin S$, we have $f(S \cup \{(\alpha_0, M_0)\}) =$

$$\begin{aligned} & \left\| \min \left(\sum_{(\alpha, M) \in S} \alpha M + \alpha_0 M_0, T \right) \right\|_1 = \left\| \min \left(\sum_{(\alpha, M) \in S} \alpha M, T \right) \right. \\ & \quad \left. + \min(\alpha_0 M_0, T - \min(\sum_{(\alpha, M) \in S} \alpha M, T)) \right\|_1 \Rightarrow \\ & f_S((\alpha_0, M_0)) = \left\| \min(\alpha_0 M_0, T - \min(\sum_{(\alpha, M) \in S} \alpha M, T)) \right\|_1, \end{aligned} \quad (7)$$

where $f_S((\alpha_0, M_0))$ denotes the incremental marginal value of adding (α_0, M_0) to the set S (see Section 3.2). Now, for $S \subseteq S' \in 2^{\mathbb{Z} \times \mathcal{M}}$ and $(\alpha_0, M_0) \notin S'$, since

$$T - \min \left(\sum_{i \in S'} \alpha_i M_i, T \right) \leq T - \min \left(\sum_{i \in S} \alpha_i M_i, T \right),$$

this together with Equation (7) implies

$$f_{S'}(\{(\alpha_0, M_0)\}) \leq f_S(\{(\alpha_0, M_0)\}).$$

Hence f is submodular. $\square$

A.2 Proof of Theorem 2

PROOF. Recall the submodular sum-throughput function f defined in Equation (2). Let $\{(\alpha_1, M_1), \dots, (\alpha_k, M_k)\}$ be the schedule returned by Algorithm 2. Let $S_i = \{(\alpha_1, M_1), \dots, (\alpha_i, M_i)\}$ denote the schedule computed at the end of i iterations of the **while** loop and let S^* denote the optimal schedule. Now, since in the $i + 1$ -th iteration (α_{i+1}, M_{i+1}) maximizes $\frac{\min(\alpha M, T_{\text{rem}}(i+1))\|_1}{(\alpha + \delta)} = \frac{f_{S_i}(\{(\alpha, M)\})}{(\alpha + \delta)}$ we have for any $(\alpha, M) \notin S_i$,

$$\begin{aligned} \frac{f_{S_i}(\{(\alpha, M)\})}{(\alpha + \delta)} & \leq \frac{f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\})}{(\alpha_{i+1} + \delta_{i+1})} \\ \Rightarrow f_{S_i}(\{(\alpha, M)\}) & \leq \frac{(\alpha + \delta)}{(\alpha_{i+1} + \delta_{i+1})} f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\}). \end{aligned} \quad (8)$$

Now consider $\text{OPT} - f(S_i)$ for some $i < k$. Since f is monotone we have

$$\begin{aligned} \text{OPT} - f(S_i) & = f(S^*) - f(S_i) \leq f(S_i \cup S^*) - f(S_i) \\ & \leq \sum_{(\alpha, M) \in J^*} f_{S_i}(\{(\alpha, M)\}) \\ & \leq \sum_{(\alpha, M) \in J^*} \frac{(\alpha + \delta)}{(\alpha_{i+1} + \delta_{i+1})} f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\}) \\ & \leq \frac{W}{(\alpha_{i+1} + \delta_{i+1})} f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\}), \end{aligned} \quad (9)$$

where $J^* := S^* \setminus S_i$ denotes the set of matchings that are present in the optimal solution but not in S_i , Equation (9) follows from Equation (8), and Equation (10) follows because $\sum_{(\alpha, M) \in J^*} (\alpha + \delta) \leq \sum_{(\alpha, M) \in S^*} (\alpha + \delta) \leq W$. Next, observe that

$$\begin{aligned} f(S_{i+1}) & = f(S_i) + f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\}) \\ \Rightarrow \text{OPT} - f(S_{i+1}) & = \text{OPT} - f(S_i) - f_{S_i}(\{(\alpha_{i+1}, M_{i+1})\}) \\ & \leq (\text{OPT} - f(S_i)) \left(1 - \frac{(\alpha_{i+1} + \delta)}{W} \right) \quad (11) \\ & \leq (\text{OPT} - f(S_0)) \prod_{i'=1}^{i+1} \left(1 - \frac{(\alpha_{i'} + \delta)}{W} \right) \\ & \leq \text{OPT} \times e^{-\sum_{i'=1}^{i+1} (\alpha_{i'} + \delta)/W}, \end{aligned} \quad (12)$$

where Equation (11) follows from Equation (10) and Equation (12) follows because of the identity $1 - x \leq e^{-x}$. Now, since after the k -th iteration the **while** loop terminates, this implies $\sum_{i'=1}^k (\alpha_{i'} + \delta) > W$. However, if the entries of the input traffic matrix T are bounded by $\epsilon W + \delta$, then no matching has a duration longer than ϵW . In particular $\alpha_k + \delta \leq \epsilon W \Rightarrow \sum_{i'=1}^{k-1} (\alpha_{i'} + \delta) \geq W(1 - \epsilon)$. Thus, setting $i = k - 2$ in Equation (12) we have

$$\begin{aligned} \text{OPT} - f(S_{k-1}) & \leq \text{OPT} \times e^{-\sum_{i'=1}^{k-1} (\alpha_{i'} + \delta)/W} \leq \text{OPT} \times e^{-(1-\epsilon)} \\ \Rightarrow \text{OPT} - \text{ALG2} & \leq \text{OPT} \times e^{-(1-\epsilon)}. \end{aligned}$$

Hence we conclude $\text{ALG2} \geq \text{OPT}(1 - e^{-(1-\epsilon)})$. $\square$