

# Scheduling Congestion-Free Updates of Multiple Flows with *Chronicle* in Timed SDNs

Jiaqi Zheng<sup>†</sup>, Bo Li<sup>†</sup>, Chen Tian<sup>†</sup>, Klaus-Tycho Foerster<sup>‡</sup>, Stefan Schmid<sup>‡</sup>, Guihai Chen<sup>†</sup>, Jie Wu<sup>§</sup>

<sup>†</sup>*State Key Laboratory for Novel Software Technology, Nanjing University, China*

<sup>‡</sup>*Faculty of Computer Science, University of Vienna, Austria*

<sup>§</sup>*Center for Networked Computing, Temple University, USA*

**Abstract**—The advent of more accurate synchronization in Software-Defined Networks (SDNs) in general and the notion of timed updates in particular, enables operators to fully exploit the potential of the more fine-grained and adaptive traffic engineering, by avoiding disruptions and inconsistencies during the update. However, little is known today about how to schedule the update of multiple flows in such timed SDNs: As flows compete for limited resources, implementing a congestion-free update remains algorithmically challenging, even in timed SDNs.

This paper initiates the study of the fundamental problem of how to reroute the update of *multiple* network flows in a synchronized SDN in a *congestion-free* manner. We show that that the problem is NP-hard already for flows of unit size and network links with unit delay. Our main contribution is a first solution for this problem: *Chronicle*. Our approach is based on a time-extended network construction and resource dependency graph, which is implemented by Openflow 1.5 using the *scheduled bundles feature*. Evaluation results show that *Chronicle* can reduce the makespan by 63% and reduce the number of changed rules by 50% compared to state-of-the-art.

## I. INTRODUCTION

**Motivation:** The more fine-grained and adaptive traffic engineering enabled by Software-Defined Networks (SDN), is one of the key benefits of this new networking paradigm [11]. However, reaping the benefits of a more adaptive network control is still challenging in practice due to the inherent asynchrony in the communication between SDN controller and switches, and also in the switches themselves [17]. Indeed, many potential disruptions, (transient or even permanent) inconsistencies, and instabilities, have been identified and studied empirically and analytically over the last years, making the SDN update problem an active area of research [13], [6].

Network updates are not only relevant in the context of a more adaptive and fine grained traffic engineering (typically, for minimizing network loads), but the ability to quickly and consistently reroute flows is crucial for the correctness, availability, and performance more generally. For example, network update problems arise due to changes in the network's (security) policy, upon link failures, or during maintenance work.

The recent introduction of a notion of time in SDNs and the resulting more accurate synchronization, enabled *timed* updates in OpenFlow [24]: updates which can be scheduled accurately in time and hence allow to mitigate, or even avoid entirely, the above problems. In particular, it has been shown

that so-called *flow swaps* [24], results in significantly less packet loss during updates.

However, while synchronized SDNs *enable* a faster and more consistent network update, they still pose a challenging *algorithmic problem* which is hardly understood today: on the one hand, it is desirable that the update, respectively *rerouting* of flows is completed *fast*. On the other hand, it is important that the update is congestion-free, which implies that the update schedule of flows needs to be *jointly* optimized: due to capacity constraints and given link latencies, flows may temporarily interfere, which results in packet loss and harms performance.

**Our contributions:** This paper initiates the study of the fundamental problem of how to schedule the update of *multiple* network flows in a synchronized SDN in a *congestion-free* manner. We show that the problem is NP-hard already for flows of unit size and network links with unit delay. We also show that a greedy approach to update the network can delay the update significantly. Our approach is based on a time-extended network construction and resource dependency graph. We implement our system in Openflow 1.5 using the *scheduled bundles feature*, and evaluate its feasibility and efficiency on a small-scale testbed and using large-scale simulations.

First and foremost, the scheduling of multiple flow updates raises the question of the to be considered time horizon. We use the time-extended network to capture the dynamic process of flow transmission during the network updates. Based on this, we ask for accurate time schedules—specifying for each switch and flow an update time point—such that the total update time (the *makespan*) is minimized and congestion-freedom is ensured at any moment in time. We formulate this problem as an optimization program in the time-extended network and prove its hardness.

Our second contribution is *Chronicle*, a heuristic scheduling algorithm. The key idea is to first divide all network update instances into small update blocks, then build the dependency relations among blocks. Based on the constructed time-extended network, we adjust the update time and merge the common update blocks accordingly. Finally we construct a resource dependency graph among the update blocks and schedule these blocks in time domain.

We evaluate *Chronicle* using both a prototype implementation and large-scale simulations. We develop a prototype using

the new `scheduled bundles` feature of Openflow 1.5. We use OFSoftSwitch and Dpctl [5] as Openflow switches and the controller. Our evaluation results show that Chronicle can reduce the update time by 63% and reduce the number of changed rules by 50% compared to state-of-the-art. At the same time, Chronicle can avoid transient congestion, save flow table space and provide near optimal solution.

**Novelty:** The need for fast and consistent network update mechanisms has been articulated well in the literature in various contexts, for security [22], performance [16], and dependability [20], [21], [31] reasons. Most existing literature focuses on “interactive” update mechanisms, involving the controller which monitors the progress of the update (e.g., using acknowledgements) before deciding to start the next stage of the update. In particular, existing approaches can be roughly classified into (1) *two-phase* protocols [28], [21], [15], [8] in which the controller first installs the new rules before tagging packets with the new path at the ingress port, ensuring that each packet either takes the old or the new route, but never a combination of both; (2) *node-ordering* protocols [17], [12] where the controller updates (subsets of) switches one-by-one, such that transient inconsistencies are avoided (without the need for tagging). While solutions between the two worlds are emerging [29], these approaches have in common that they need to rely on interactions with the controller to implement synchronization.

Our approach is enabled by technologies such as Time4 [24], [26] which allow to synchronize network updates using accurate time [25]. To the best of our knowledge, the only algorithmic study of the network update problem in timed SDNs is by Zheng et al. [30], which however focuses on a *single flow*. While constituting an interesting first step, we anticipate many situations in which multiple flows need to be updated simultaneously, e.g., upon a policy change or link failure, but also due to a new traffic engineering optimization. Scheduling multiple flow updates simultaneously however is significantly harder as different flows can interfere in complex (potentially combinatorial) ways at different links.

## II. BACKGROUND AND MOTIVATION

We consider a network where a controller updates the forwarding rules at the switches whenever a route changes. Fig. 1(a) illustrates a simple example: there are five switches  $v_1, \dots, v_5$  in the network. The link capacity of  $\langle v_2, v_5 \rangle$  is assumed to be two units and the rest is one unit. The propagation delay of link  $\langle v_3, v_4 \rangle$  is assumed to be three time units and the rest is one time unit. That is, if one unit of flow leaves switch  $u$  at time  $t$  on the link  $\langle u, v \rangle$ , one unit of flow arrives at switch  $v$  at time  $t + \sigma_{u,v}$ , where  $\sigma_{u,v}$  is the propagation delay between  $u$  and  $v$ . We use the notion of *dynamic flow* to represent the propagation of packets of a flow in time domain [14]. In our example, the demand of two “dynamic flows” colored as red and green are both one unit, which are both routed from the source  $v_1$  to the destination  $v_5$ . The initial routing are depicted as two solid red and green lines and the final routing are depicted as two dashed red and green lines. With

dynamic flows, the utilization of a link varies over time. As discussed before, prior work on the network update problem usually relies on one of two fundamental update techniques: *two-phase update protocols* and *node ordering protocols*.

**Two-phase update protocols:** In the first phase, new rules—whose matching fields use the new version tag that corresponds to the second stage—are added. During this phase, flows are still forwarded according to existing rules as packets are still stamped with the old version tag of the first stage. Once the update is done for all switches, the protocol enters the second phase, when we stamp every incoming packet with the new version tag. At this point the new rules become functional, and old rules are removed by the controller. Reitblatt et al. [28] initiated this line of research by introducing a two-phase commit protocol that preserves consistency when changing between two different routing configurations. Based on this, SWAN [15] and zUpdate [21] try to find a congestion-free two-phase update plan. SWAN shows that if each link has a certain slack capacity, there always exists a congestion-free update sequence. This condition is too strong to always hold in practice. Furthermore, Brandt et al. [8] analyze the condition that a congestion-free update sequence exists. As the update plan is not unique, Dionysus [17] seeks to determine a fastest update sequence according to different runtime conditions of switches.

A two-phase update procedure in our example of Fig. 2(b) is: the route of new version tag for red and green flows are updated in the first phase; in the second phase, we change the version tag at the source switch  $v_1$  and swap the red and green flows into their final path. Fig. 2(c) shows a possible asynchronous update case, where the time difference between two phases is assumed to be one time unit. We can observe that the congestion happens at the link  $\langle v_4(t_4), v_5(t_5) \rangle$  as one unit capacity of link  $\langle v_4, v_5 \rangle$  cannot accommodate two flows at the same time.

**Node ordering protocols:** At each round, the controller waits until all the switches have completed their updates, and only then invokes the next round. Ludwig et al. [23] aim to minimizing the number of sequential controller interactions when transitioning from the initial to the final update stage. The authors prove that finding a shortest node ordering sequence that avoids forwarding loops is NP-hard. Furthermore, they introduce a notion of relaxed loop-freedom, which provides an interesting consistency-runtime tradeoff. Another work by Ludwig et al. [22] considers secure network updates in the presence of middleboxes [27]. The authors try to find a node ordering sequence that preserves a specific security policy.

A possible node ordering sequence in our example of Fig. 2(b) is: Fig. 1(c1)  $\rightarrow$  (c2). In the first round,  $v_2$  (the route for red flow),  $v_3$  (the route for green flow) and  $v_4$  (the route for green flow) are updated asynchronously. Then  $v_1$  (the route for both red and green flows) and  $v_2$  (the route for green flow) are updated in the second round. Due to the asynchronous nature of the data plane, the new route for  $v_2$  (green flow) may become functional earlier than that for  $v_1$  (red and green flows) in the second round as shown in Fig. 1(c2). At this

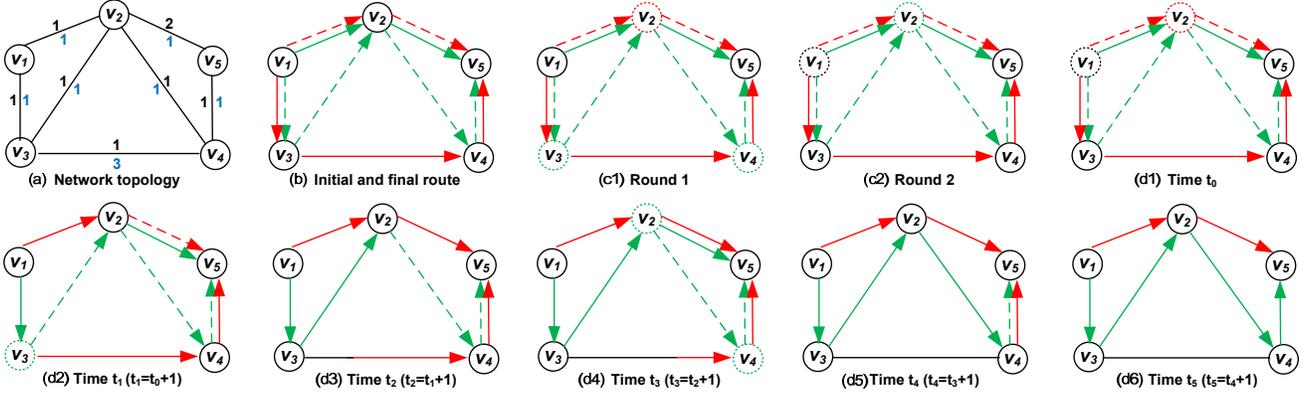


Fig. 1. Illustration of the network update problem considered in this paper. In this example topology,  $v_1$  is the source and  $v_5$  is the destination of both the old (initial) route and the new (final) route. There are two flows in the network, which are depicted as green and red, respectively. The initial routing for the two flows are illustrated as two solid lines with green and red, while the final routing are represented as two dashed lines with green and red. The solid links represent that the load on the link is greater than zero, which indicates that the dynamic flow is passing through this link. In our example, the link capacity of  $\langle v_2, v_5 \rangle$  is assumed to be two units and the rest is one unit. The link propagation delay of  $\langle v_3, v_4 \rangle$  is assumed to be three units and the rest is one unit. The timed update sequence is: Fig. 1(d1)  $\rightarrow$  (d2)  $\rightarrow$  (d3)  $\rightarrow$  (d4)  $\rightarrow$  (d5)  $\rightarrow$  (d6).

point, the red and green flows are routed through the path  $\langle v_1, v_3, v_4, v_5 \rangle$  and  $\langle v_1, v_2, v_4, v_5 \rangle$  respectively. Here the red and green flows together would result in a transient congestion on the link  $\langle v_4, v_5 \rangle$  as the sum of flow demand is two units, which are beyond the one unit link capacity (Fig. 2(b)).

**Timed update protocols:** Mizrahi et al. [24], [26] propose a time synchronization protocol between controller and data plane, which uses accurate timing to trigger network updates. Though the idea of timed update has surfaced in the literature recently, the only algorithmic study of the timed update is by Zheng et al. [30], which however focuses on a *single flow*. To this end, Zheng et al. prove that minimizing the makespan is NP-hard, and design a heuristic algorithm to perform network update for a single flow. The novelty of our work lies in the first study and comprehensive exploration of the design of timed scheduling algorithms for *multiple flows*. We also extend the hardness results of Zheng et al. [30] to the case of multiple flows, in particular we show that it is NP-hard even for flows of unit size and network links with unit delay.

A timed update schedule (Fig. 1(d1)  $\rightarrow$  (d2)  $\rightarrow$  (d3)  $\rightarrow$  (d4)  $\rightarrow$  (d5)  $\rightarrow$  (d6)) can effectively solve our problem. Firstly, the route of  $v_1$  for both red and green flows are updated at  $t_0$ . And then  $v_3$  for green flow is updated at  $t_1$ . Finally the route of  $v_2$  and  $v_4$  for green flow is updated simultaneously at  $t_3$ . The congestion-free condition are ensured at any moment in time as shown in the time-extended network of Fig. 2(d). This timed schedule can be acceptable in practice because the flow table rules can be updated accurately on the order of one microsecond [25]. In addition, the controller can send all the update commands at a time and the update behavior for each switch is triggered by a pre-defined time instant, which can significantly decrease the time overhead resulting from wait-invoke mechanism of *node ordering protocols*. Also we only modify the action in the flow table during the update process, where we do not require additional flow table space headroom and overcome the drawback of *two-phase update protocols*.

TABLE I  
KEY NOTATIONS IN THIS PAPER.

$F$	The set of dynamic flow $f$
$V$	The set of switches $v$
$E$	The set of links $\langle u, v \rangle$
$G$	The acyclic directed network graph $G = (V, E)$
$t_i$	The time point. $t_{i+1} > t_i$
$T$	The set of time point. $T = \{t_0, t_1, \dots, t_n\}$
$F^T$	The set of flows in the time-extended network
$V^T$	The set of switches $v(t)$ , where $v \in V$ and $t \in T$
$E^T$	The set of links $\langle u(t_i), v(t_j) \rangle$
$G^T$	The time-extended network $G^T = (V^T, E^T)$
$C_{u,v}$	The capacity of link $\langle u, v \rangle$
$p_{init}^f$	The initial path for the dynamic flow $f$
$p_{fin}^f$	The final path for the dynamic flow $f$
$d_f$	The demand of the dynamic flow $f$
$n$	The number of the switches. $n =  V $
$\sigma_{u,v}$	The transmission delay for the link $\langle u, v \rangle$ .

### III. AN OPTIMIZATION FRAMEWORK

#### A. Dynamic Flow Model and Problem Formulation

Before formulating the problem, we first present our network model. A network is a directed graph  $G = (V, E)$ , where  $V$  is the set of switches and  $E$  the set of links with capacities  $C_{u,v}$  and transmission time  $\sigma_{u,v}$  for each link  $\langle u, v \rangle \in E$ . For each flow  $f$ , the network contains two paths:  $p_{init}^f$  and  $p_{fin}^f$ . The former is the old routing path which is depicted as a solid line in our example and the latter is the new routing path depicted as a dashed line. We use different colors to distinguish different flows. Both of  $p_{init}^f$  and  $p_{fin}^f$  have a common source  $v^+$  and destination  $v^-$ . For convenience, we summarize important notations in Table I. Let us introduce three related notations first.

**Definition III.1. Dynamic flow [14]:** A dynamic flow on  $G$  is a function  $f : E \times T \rightarrow Z_+$  ( $Z_+$  represents the set of positive integers) that satisfies the following conditions for  $\forall v \in V - \{v^-, v^+\}$  and  $\forall t \in T$ .

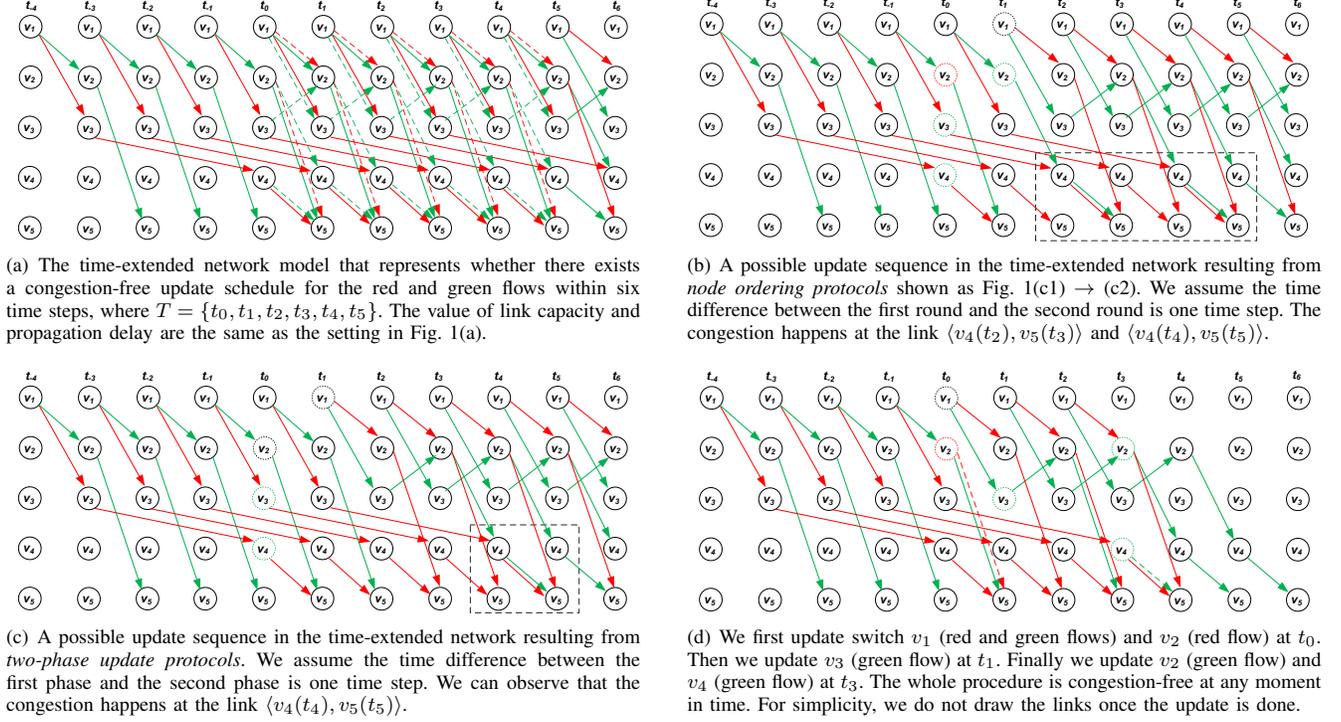


Fig. 2. Illustration of *two-phase update protocols*, *node ordering protocols* and our timed schedule shown in the time-extended network.

$$\begin{aligned}
& \sum_{(u,v) \in E^+(v), t - \sigma_{u,v} \geq 0} x_{u,v}(t - \sigma_{u,v}) - \sum_{(u,v) \in E^-(v)} x_{u,v}(t) \\
& = \begin{cases} -d_f & v = v^-, \forall t \in T \\ 0 & \forall v \in V - \{v^-, v^+\}, \forall t \in T \\ d_f & v = v^+, \forall t \in T \end{cases} \quad (1)
\end{aligned}$$

The conservation condition (1) indicates that if one unit of dynamic flow leaves switch  $u$  at  $t - \sigma_{u,v}$  on link  $\langle u, v \rangle$ , one unit of flow arrives  $v$  at  $t$ . Here  $E^+(v)$  and  $E^-(v)$  represent the set of links incoming and outgoing to the switch  $v$ , respectively. The notation  $d_f$  is the flow demand, which is a positive integer. The set  $T$  is measured in discrete steps, where  $T = \{t_0, t_1, \dots, t_n\}$ .  $x_{u,v}(t)$  characterizes the load on the link  $\langle u, v \rangle$  at  $t$ , which cannot go beyond the link capacity at each moment in time.

$$0 \leq x_{u,v}(t) \leq C_{u,v}, \forall \langle u, v \rangle \in E, \forall t \in T \quad (2)$$

Condition (2) ensures that the link capacity  $C_{u,v}$  cannot be violated for  $\forall t \in T$ .

**Definition III.2. Congestion-free condition:** *The congestion-free condition holds if and only if condition (2) always holds for  $\forall t \in T$  throughout the update process.*

Our model and approach can be visualized nicely with a *time-extended network concept*: a network in which there is a

copy of each switch for every time step  $t_i \in T$  and the links are redrawn between these copies to express their transmission delay. Succinctly:

**Definition III.3. The time-extended network:** *The time-extended network  $G^T$  is a directed graph  $G$  with switches  $v(t)$  for all  $v \in V$  and  $t \in T$ . For each link  $\langle u, v \rangle \in E$  with transmission delay  $\sigma_{u,v}$  and capacity  $C_{u,v}$ , the network  $G^T$  has link  $\langle u(t), v(t + \sigma_{u,v}) \rangle$  with capacity  $C_{u,v}$ .*

The time-extended network captures the dynamic process of flow transmission in the network. Fig. 2(a) gives a time-extended network example of Fig. 1(a), where  $t_{-1}, \dots, t_{-3}$  and  $t_{-4}$  represent the past time steps,  $t_0$  represents the current time step,  $t_1, t_2, \dots$  represent the future time steps. The green flow on the link  $\langle v_1(t_0), v_2(t_1) \rangle$  starts at current time step  $t_0$ , while the green flow on the link  $\langle v_2(t_0), v_5(t_1) \rangle$  and the red flow on the link  $\langle v_3(t_0), v_4(t_3) \rangle$  both start at past time step  $t_{-1}$ . The red solid line between  $v_3$  and  $v_4$  strides over three time steps as its link transmission delay we assumed in Fig. 1(a) are three time units. The rest only stride over one time step as its link transmission delay is one time unit. We can only update the switches in the current and future time steps and cannot update them in the past steps. The reason why we illustrate the past time steps there is that we require to check the congestion-free condition defined in (III.2). In Fig. 2(a), the red flow starting at past time step  $t_{-4}$  will occupy the link bandwidth between  $v_4(t_0)$  and  $v_5(t_1)$ , which makes a difference to the updates at current time step  $t_0$ .

Based on the above model and definition, we formulate the *Minimum Update Time Problem for Multiple Flows (MUTP-MF)* as an integer linear program (3) in the time-extended network, where the initial (solid lines) and final (dashed lines) routing paths for each flow  $f$  are given. We seek to find an optimal timed update sequence so as to minimize the total update steps, such that the congestion-free condition holds at any moment in time.

$$\begin{aligned} & \text{minimize} && |T| && (3) \\ & \text{subject to} && \sum_{f \in F} d_f \cdot y_{u(t_i), v(t_j)}^f \leq C_{u(t_i), v(t_j)}, \\ & && \forall \langle u(t_i), v(t_j) \rangle \in E^T, && (3a) \\ & && (3b), (3c), (3d), (3e), (3f). \end{aligned}$$

The formulation of the minimum update time problem for multiple flows is shown in (3). The objective aims to minimize the number of elements in set  $T$ , i.e., the time steps during the update. At the beginning, the element of set  $T$  is  $t_0$ . We iteratively add one time step  $t_i$  into the set  $T$  each time until we find a feasible solution or the number of elements in  $T$  reaches a pre-defined threshold. The upper bound analysis about the number of elements in  $T$  will be discussed soon in Theorem III.3. The LHS of constraint (3a) in the formulation characterizes the load of total flows at link  $\langle u(t_i), v(t_j) \rangle$ , which must be less than or equal to its capacity in order to meet the congestion-free condition defined in (III.2).

$$x_{u(t_i)}^f \in \{0, 1\}, \quad \forall f \in F, \forall u(t_i) \in V^{T \setminus t_n}, \quad (3b)$$

$$x_{u(t_n)}^f = 1, \quad \forall f \in F, \forall u(t_n) \in V^T. \quad (3c)$$

The zero-one integer variable  $x_{u(t_i)}^f$  equals one when the routing configuration of switch  $u$  for flow  $f$  is updated at  $t_i$  in the time-extended network, and equals zero otherwise. This optimization variable determines that which switch should be updated at which time point. The optimization variables  $x_{u(t_i)}^f$  ( $t_i \in \{T \setminus t_n\}$ ) need to be determined, while the variable  $x_{u(t_n)}^f$  ( $t_n$  is the last time step in set  $T$ ) is known to be one as formulated in Constraint (3c) since all updates should be complete before the last time steps  $t_n$ .

$$y_{u(t_i), v(t_j)}^f = 1 - x_{u(t_i)}^f, \quad \forall f \in F, \langle u(t_i), v(t_j) \rangle \in p_{init}^f, \quad (3d)$$

The zero-one integer variable  $y_{u(t_i), v(t_j)}^f$  in Constraint (3d) indicates that whether the flow  $f$  is routed through the link  $\langle u(t_i), v(t_j) \rangle$  belong to the initial path  $p_{init}^f$ . Obviously, it equals zero when the switch  $u$  is updated at  $t_i$ , and equals one otherwise.

$$y_{u(t_i), v(t_j)}^f = x_{u(t_i)}^f, \quad \forall f \in F, \langle u(t_i), v(t_j) \rangle \in p_{fin}^f, \quad (3e)$$

On the contrary, the zero-one integer variable  $y_{u(t_i), v(t_j)}^f$  in Constraint (3e) indicates that whether the flow  $f$  is routed through the link  $\langle u(t_i), v(t_j) \rangle$  belong to the final path  $p_{fin}^f$ .

It equals one when the switch  $u$  is updated at  $t_i$ , and equals zero otherwise.

$$x_{u(t_i)}^f \geq x_{u(t_j)}^f, \quad \forall f \in F, t_i \geq t_j, \quad (3f)$$

Constraint (3f) captures a fact that the routing configuration at a specific switch for the flow  $f$  remains unchanged once the update is complete. That is to say, we can only update the route from the initial to final path, nor the other. For example in Fig. 2(d), once the route of switch  $v_1$  for red flow is updated at  $t_0$  ( $x_{u(t_0)}^f = 1$ ), it will stay the same at the next time steps ( $x_{u(t_1)}^f = 1, x_{u(t_2)}^f = 1, \dots$ ).

Note that the switches processing delay for updating a forwarding rule in TCAM [4] can influence the accuracy of our model. However, recent work [7] shows that the update time can be predictable and a constant. This suggests that we can subtract a corresponding time offset from the outputs of our model, indicating that the time point triggering the update should be earlier than the resulting update time calculated from our model.

## B. Theoretical Analysis

We establish the hardness of our problem MUTP-MF below.

**Theorem III.1.** *The feasibility of MUTP-MF is NP-hard.*

*Proof:* Consider a special case of MUTP-MF as shown in Fig. 3. The capacity of all links is  $C$  units, and the delay of all links is one unit. There are  $k$  green flows each with demand  $d_i$  and  $\sum_{i \in \{1, 2, \dots, k\}} d_i = C$ . They are routed through the initial path  $\langle s_1, v, w, t \rangle$  and will be moved into their final path  $\langle s_1, u, t \rangle$ . The red flow with demand  $\frac{C}{2}$  is routed through the initial path  $\langle s_2, u, t \rangle$  and will be moved into the final path  $\langle s_2, w, t \rangle$ . The objective is to assign a update time point for  $k+1$  flows such that the congestion-free condition holds at any moment in time. The only way is that we firstly update the green flows with total demand  $\frac{C}{2}$  at  $t_0$  and keep half of the link capacity at  $\langle w, t \rangle$  vacant. Finally we update the red flow and the rest green flows simultaneously at  $t_1$ , where the time difference between  $t_1$  and  $t_0$  is one time unit.

We construct a polynomial reduction from the set partition problem [9] to it. Consider a partition instance consisting of  $k$  items, each with a value  $a_i$  and  $\sum_{i \in \{1, 2, \dots, k\}} a_i = C$ . Each item  $i$  corresponds to one of  $k$  green flows in the example of Fig. 3, where  $a_i = d_i$ ,  $i \in \{1, 2, \dots, k\}$ . Therefore, any feasible partition of the items corresponds to the updates of  $k$  green flows in two time steps, and vice versa. The routing update in the first time step forms one set of the partition, and that in the second time step forms the other. ■

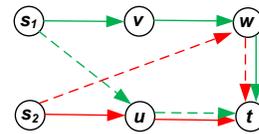


Fig. 3. Topology used for the reduction from Partition to MUTP-MF.

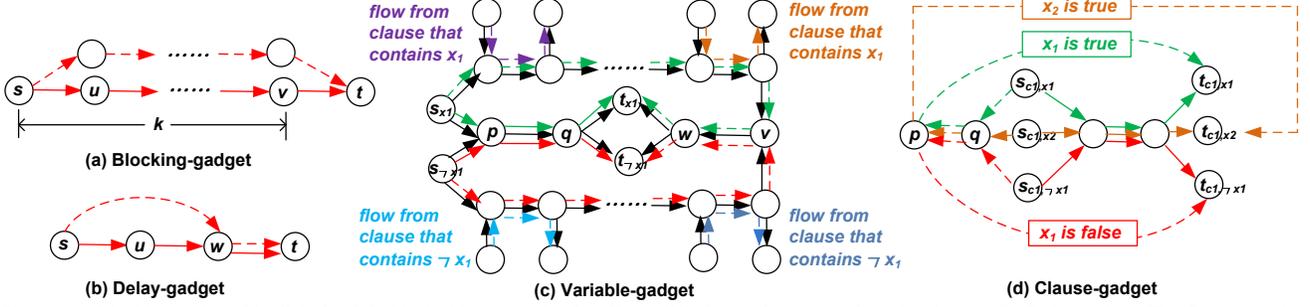


Fig. 4. (a) Blocking-gadget. The link  $(v, t)$  is blocked by one unit of traffic until  $k$  time units when  $s$  is updated at  $t_0$ . (b) Delay-gadget. The flow can only be updated when there is additional free capacity at link  $(w, t)$ . (c) Variable-gadget. The capacity of link  $(p, q)$  is two units and the rest is one unit. Due to the capacity of link  $(v, w)$  being one unit, we cannot update  $s_{x_1}$  and  $s_{\neg x_1}$  at the same time. (d) Clause-gadget for  $(x_1 \vee x_2 \vee \neg x_1)$ . As the capacity of link  $(p, q)$  is one unit, only one of the three flows can be updated. However, this flow needs free capacity in its variable-gadget path.

The hardness of Theorem III.1 relies on different flow sizes, but even for unit size flows, MUTP-MF is NP-hard as well.

**Theorem III.2.** *MUTP-MF is NP-hard, already for flows and delays of unit size.*

*Proof:* Our reduction from 3-SAT [18] will feature four gadgets, called blocking-gadget (Fig. 4(a)), delay-gadget (Fig. 4(b)), variable-gadget (Fig. 4(c)), clause-gadget (Fig. 4(d)). The demand of all flows and the delay of all links will be of unit size in this proof.

A blocking-gadget consists of a flow where the only update is performed at the source, with the old and new path being link-disjoint. By adjusting the length of the old path, we can block one unit of capacity on a link for any desired time. As such, for the other (gadget-) constructions, we can assume any capacity restrictions for new paths, as the new paths will become eventually feasible once the blocking flows leave.

A variable-gadget consists of two flows (truth assignments) that share a path to their destination on the old path of capacity two, but the new paths are split into a “true” and a “false” path, merging into a joint path to their destination, all of capacity one. As such, only either the true or the false flow can update for now, as they collide on the joint path of capacity one.

A clause-gadget consists of three flows (literals) that share an old path to their destination of capacity three, but on the new path, they first share a link of capacity one, and then split their paths, each traversing the corresponding true or false path of their variable-gadget for one link, then reaching their destination. By making the true and false paths sufficiently long enough in the variable-gadget, the paths of literal-flows from different clause-gadgets will not overlap.

Observe that only one literal-flow of each clause can update, but if no flow from a variable-gadget updates, every clause-gadget can update one literal-flow without eventual congestion. To prevent this, we introduce the delay-gadget with one flow, where the old path has a length of three and the new path has a length of two, but they share only the last link. Updating the flow introduces twice the utilization of the last link for one time unit, after one time unit. We add a delay-gadget each to the end of the old paths of the variable- and clause-gadgets, only sharing the last link, increasing that link’s capacity by

one. As the old paths are fully utilized, the delay-gadget cannot update until a flow from the respective gadgets updates.

Hence, all delay-gadgets can’t update for now, unless one literal-flow from each clause and one flow from each variable updates. However, finding such an update is equivalent to solving the corresponding 3-SAT instance. Let  $t$  be the earliest time when the last delay-gadget can update in case the 3-SAT instance is satisfiable and the blocking-gadgets have infinite path lengths. By adjusting the blocking-gadgets appropriately to “free” their blocked capacity for time  $t$ , it is NP-hard to decide if all nodes can update by time  $t$ . ■

**Theorem III.3.** *The maximum time steps in set  $T$  are bounded by  $\sum_{f \in F} \sum_{v \in p_{init}^f \cap p_{fin}^f} \arg \max_{p: v \in p, v \in B} \phi(p)$ , where  $p_{init}^f$  and  $p_{fin}^f$  represent initial and final path for flow  $f$ , function  $\phi(\cdot)$  refers to the sum of link transmission delay.*

*Proof:* The switches in the time-extended network have to wait some time steps in order to ensure that the congestion-free condition holds at any moment in time. We denote by  $t_{v,f}$  the waiting time steps for switch  $v$  and flow  $f$  in the time-extended network and thus we have

$$|T| \leq \sum_{f \in F} \sum_{v \in B} t_{v,f} \leq \sum_{f \in F} \left( \sum_{v \in A} t_{v,f} + \sum_{v \in (B \setminus A)} t_{v,f} \right) \quad (4)$$

where  $A = \{v | v \in p_{init}^f \cap p_{fin}^f\}$  that represents the set of switches both in the initial path  $p_{init}^f$  and final path  $p_{fin}^f$ , and  $B = \{v | v \in p_{init}^f \cup p_{fin}^f\}$  that represents the set of switches either in the initial path  $p_{init}^f$  or final path  $p_{fin}^f$ . Obviously, each update in set  $\{B \setminus A\}$  does not need to wait and we obtain  $\sum_{v \in B \setminus A} t_{v,f} = 0, \forall f \in F$ . Combining inequation (4), we have the following.

$$|T| \leq \sum_{f \in F} \sum_{v \in A} t_v \leq \sum_{f \in F} \sum_{v \in A} \arg \max_{p: v \in p, v \in B} \phi(p)$$

where the function  $\phi(p)$  refers to the sum of link transmission delay on the path  $p$ . The notation  $p$  represents a mixed path traveling through the switches either in the initial path  $p_{init}^f$  or final path  $p_{fin}^f$ . ■

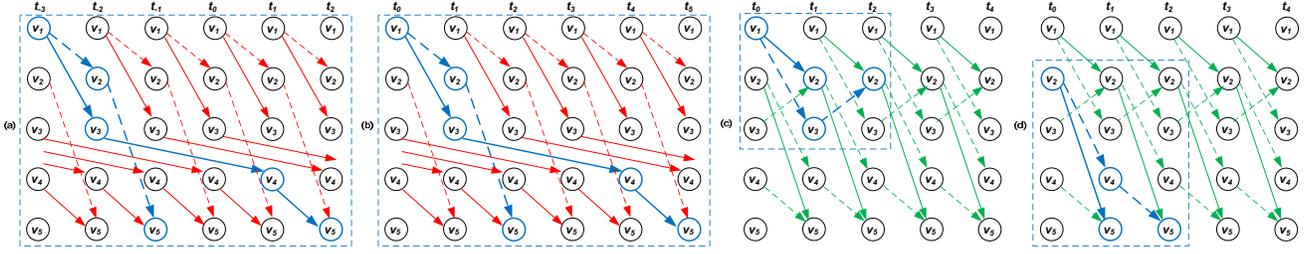


Fig. 5. Four update blocks (colored blue) (a)  $o_1^{f_r}(t_{-3})$ , (b)  $o_1^{f_r}(t_0)$ , (c)  $o_1^{f_g}(t_0)$  and (d)  $o_2^{f_g}(t_0)$  in the time-extended network for red flow  $f_r$  and green flow  $f_g$  of Fig. 1(b).

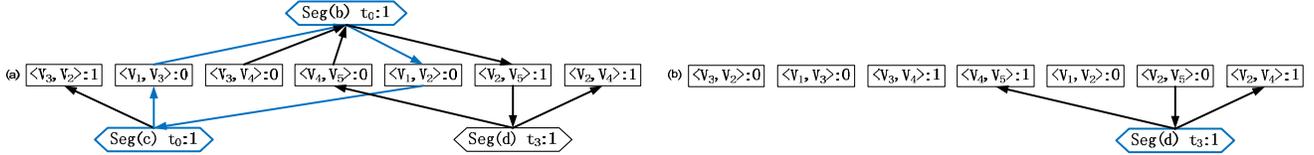


Fig. 6. Illustration of the resource dependency graph.

#### IV. THE CHRONICLE ALGORITHM

In this section we design a scheduling algorithm to find a feasible update sequence in polynomial time. We firstly explain the high level working of our algorithm. To increase the flexibility, we divide the network update instance into small update blocks that can be scheduled individually in the time-extended network. Taking the green flow in Fig. 1(b) as an example, two blocks are  $(\langle v_1, v_2 \rangle, \langle v_1, v_3, v_2 \rangle)$  and  $(\langle v_2, v_5 \rangle, \langle v_2, v_4, v_5 \rangle)$ , respectively. Regarding for red flow, it has only one block  $(\langle v_1, v_3, v_4, v_5 \rangle, \langle v_1, v_2, v_5 \rangle)$ . The block concept introduces more convenience as we only need to update the first switch in a block. Then we put these blocks into time-extended network (Fig. 5) and establish the dependency relations among them, in which one's initial path and the other's final path have common links. In Fig. 5, the update of block (a) should be earlier than that of block (d) as the one unit capacity of link  $\langle v_4, v_5 \rangle$  cannot accommodate red and green flows simultaneously. Since the block (a) starts at the past time point  $t_{-3}$ , we add the offset of three time units for each update block in order to ensure that all the blocks start at current or further time steps. After the adjustment process is done, we merge the common blocks and construct the resource dependency graph as shown in Fig. 6. Based on this graph, we can detect the deadlocks (dependency cycles) and output a feasible update schedule.

Let us introduce three related notations first.

**Definition IV.1. Update block:** A update block  $o_j^f(t_i)$  for flow  $f$  contains two edge disjoint paths  $p_1$  and  $p_2$  starting and ending at the common node  $u$  and  $v$ , where  $p_1 \cap p_2 = \{u, v\}$ ,  $p_1 \in P_{init}^f$ ,  $p_2 \in P_{fin}^f$ .

Combining the motivating example in Fig. 1(b), four update blocks (Fig. 5) in the time-extended network are (a)  $o_1^{f_r}(t_{-4})$ , (b)  $o_1^{f_r}(t_0)$ , (c)  $o_1^{f_g}(t_0)$ , and (d)  $o_2^{f_g}(t_0)$ . Each update block starts and ends at a common switch in the initial and final path. The  $\rightarrow$  operator captures the update order between two

blocks. For example, the notation  $o_1 \rightarrow o_2$  represents that the update time of block  $o_1$  should not be later than that of block  $o_2$ . Otherwise, the congestion-free condition will be violated.

**Definition IV.2. Resource dependency graph:** A resource dependency graph captures the dependent relation between the block  $o_j^f(t_i)$  and the link  $\langle u, v \rangle$ .

Fig. 6 shows a resource dependency graph example corresponding to the update blocks in Fig. 5. There are two types of rectangles in the graph — the link rectangle and the update block rectangle. The number in the link rectangle indicates the residual capacity  $C'_{u,v}$  on the link  $\langle u, v \rangle$  at the current time step, while that in the update block rectangle represents the flow demand. For a specific update block  $o_j^f(t_i)$ , the incoming edges come from the links in the initial path, while the outgoing edges point to the links in the final path. The construction procedure is shown in Algorithm 1.

---

#### Algorithm 1: Constructing the resource dependency graph

---

- Input:** The set of all update blocks  $O$ ; the initial path  $p_{init}^f$  and the final path  $p_{fin}^f$  for each flow  $f \in F$ .  
**Output:** The resource dependency graph  $G_o$ .
- 1:  $G_o = \emptyset$ ,  $C'_{u,v} = 0$
  - 2: **for** each  $o_j^f(t_i) \in O$  **do**
  - 3:   **for** each  $\langle u, v \rangle \in p_{init}^f$  **do**
  - 4:      $G_o = G_o \cup \{\langle u, v \rangle \rightarrow o_j^f(t_i)\}$
  - 5:      $C'_{u,v} = C_{u,v} - d_f$
  - 6:   **for** each  $\langle u, v \rangle \in p_{fin}^f$  **do**
  - 7:      $G_o = G_o \cup \{o_j^f(t_i) \rightarrow \langle u, v \rangle\}$
- 

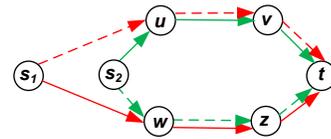


Fig. 7. A deadlock example.

**Definition IV.3. Deadlock:** A deadlock indicates that we cannot find a feasible update schedule in the network.

A deadlock forms if two conditions  $\tau_{s_1,w} > \tau_{s_2,w}$  and  $\tau_{s_1,u} > \tau_{s_2,u}$  hold at the same time shown in Fig. 7. On the one hand, the condition  $\tau_{s_1,w} > \tau_{s_2,w}$  indicates that the update time of  $s_1$  should be earlier than that of  $s_2$  to avoid the congestion at path  $\langle w, z, t \rangle$ . On the other hand, the condition  $\tau_{s_1,u} > \tau_{s_2,u}$  indicates that the update time of  $s_1$  should be later than that of  $s_2$  in order to avoid the congestion at path  $\langle u, w, t \rangle$ . This is a contradiction as we cannot find a feasible update time point for each switch.

---

**Algorithm 2:** Calculating a timed update sequence

---

**Input:** The directed acyclic network  $G$ ; the initial path  $p_{init}^f$  and the final path  $p_{fin}^f$  for each flow  $f \in F$ .

**Output:** A boolean variance that indicates whether there exists a feasible update sequence or not.

- 1: Construct the set of all update blocks  $\{o_j^f(t_0)\}$  starting at  $t_0$  in the time-extended network
  - 2: Construct the set of update blocks  $\{o_j^f(t-x)\}_x$  starting at past time steps, where  $o_j^f(t-x) \rightarrow o_j^f(t_0)$
  - 3:  $O = \{o_j^f(t_0)\} \cup \{o_j^f(t-x)\}$
  - 4: Adjust each element in set  $O$  such that all the update blocks start at current or future time steps
  - 5: Merge the common elements in set  $O$
  - 6: **if** there exists an integer  $\alpha$  such that  $o_j^f(t-\alpha) = o_j^f(t)$  ( $o_j^f(t-\alpha), o_j^f(t) \in O$ ) **then**
  - 7:     **return false**
  - 8: Apply Algorithm 1 to construct the resource dependency graph  $oG$
  - 9: **for** each  $t_i \in T$  **do**
  - 10:     Apply Algorithm 3 to obtain the set of independent update block  $\hat{O}$  at  $t_i$
  - 11:     Apply Algorithm 4 to update each block in set  $\hat{O}$  and obtain the return value  $\sigma$
  - 12:     **if**  $\sigma = -1$  **then**
  - 13:         **return false**
  - 14:      $O = O \setminus \hat{O}$
  - 15:     **for** each  $o_j^f(t) \in O$  **do**
  - 16:          $o_j^f(t) = o_j^f(t + \sigma)$
  - 17:     Finding the rest of dependent update blocks  $O^*$  at  $t_i$
  - 18:     Apply Algorithm 4 to update each block in set  $O^*$  and obtain the return value  $\sigma$
  - 19:     **if**  $\sigma = -1$  **then**
  - 20:         **return false**
  - 21:      $O = O \setminus O^*$
  - 22: **if**  $G_o = \emptyset$  **then**
  - 23:     **return true**
  - 24: **return false**
- 

The complete process of our scheduling algorithm is shown in Algorithm 2. We first calculate the set of all update blocks  $\{o_j^f(t_0)\}$  starting at time step  $t_0$  (line 1). Then we add the set of update blocks  $\{o_j^f(t-x)\}$  at the past time steps whose update time points should be earlier than that in  $\{o_j^f(t_0)\}$  (line 2). After that, we construct the set  $O$  and adjust each update block such that all of them start at current or future time steps (lines 3-4). When the merge operation is done, we check that whether the equation  $o_j^f(t-\alpha) = o_j^f(t)$  can be established or not. If this condition holds, the algorithm stops because a

deadlock forms, and we cannot schedule the update block  $o_j^f$  at two different time points ( $t-\alpha$  and  $t$ ) simultaneously (lines 5-7). Next we apply Algorithm 1 to construct the resource dependency graph  $G_o$  and schedule each update block step by step (lines 8-20). In each time step, we apply Algorithm 3 to obtain the independent set  $\hat{O}$  and try to update them using Algorithm 4 (lines 10-11). If all updates are feasible, we add  $\sigma$  time steps for the rest of each update block, where  $\sigma$  is the maximum path delay obtained from Algorithm 4 (lines 15-16). For the rest of the dependent update block, we update using Algorithm 4 as well (lines 17-18). When the loop terminates and the set  $G_o$  is empty, our algorithm outputs a feasible update sequence. Otherwise, it indicates that a feasible solution does not exist (lines 22-24). For convenience, the main steps are illustrated in Table II.

TABLE II  
MAIN STEPS FOR THE EXAMPLE SHOWN IN FIG. 1(B).

1	Calculate all update blocks: $o_1^{fr}(t-3), o_1^{fr}(t_0), o_1^{fg}(t_0), o_2^{fg}(t_0)$
2	Establish the relation: $o_1^{fr}(t-3) \rightarrow o_2^{fg}(t_0), o_1^{fr}(t_0) \leftrightarrow o_1^{fg}(t_0)$
3	Adjust the starting time: $o_1^{fr}(t_0) \rightarrow o_2^{fg}(t_3), o_1^{fr}(t_0) \leftrightarrow o_1^{fg}(t_0)$
4	Merge the common update block: $o_2^{fg}(t_3) \leftarrow o_1^{fr}(t_0) \leftrightarrow o_1^{fg}(t_0)$
5	Construct the resource dependency graph (Fig. 6)
6	Break the dependency cycles: $o_2^{fg}(t_3), o_1^{fr}(t_0), o_1^{fg}(t_0)$
7	Assign the update time instant for each switch

---

**Algorithm 3:** Finding the set of independent blocks

---

**Input:** The resource dependency graph  $G_o$ ; the set of all update blocks  $O$ ; the time step  $t$ .

**Output:** The set of independent update blocks  $\hat{O}$ .

- 1: **for** each  $o_j^f(t) \in O$  **do**
  - 2:     **for** each  $\langle u, v \rangle \in p_{fin}^f$  **do**
  - 3:         **if**  $C'_{u,v} < d_f$  **then**
  - 4:             **continue**
  - 5:      $\hat{O} = \hat{O} \cup \{o_j^f(t)\}$
- 

Algorithm 3 describes the procedure of finding the set of independent update blocks. For each one, if it can directly move to the final path without link capacity violation, we add it into set  $\hat{O}$  and the algorithm enters into the next loop.

How to update a resource dependency graph is shown in Algorithm 4. A possible schedule for a specific update block is that the residual capacity  $C'_{u,v}$  of all links in the final path can accommodate the flow demand. If so, the delay  $\sigma_{max}$  is returned and will be used in Algorithm 2. Otherwise, the integer  $-1$  is returned, indicating that the update is infeasible (lines 11-12). When the update is done, the residual capacity  $C'_{u,v}$  of all links on the initial (final) path will be increased (decreased) by a flow demand (lines 4-10).

Based on the above, we have the following theorem.

**Theorem IV.1.** The timed update sequence obtained from Algorithm 2 is congestion-free.

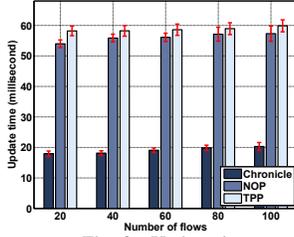


Fig. 8. Update time

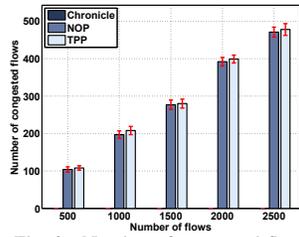


Fig. 9. Number of congested flows

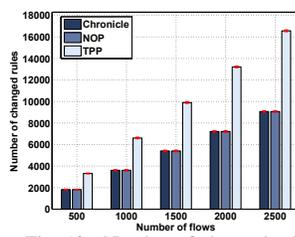


Fig. 10. Number of changed rules

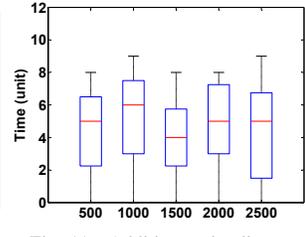


Fig. 11. Additive optimality gap

#### Algorithm 4: Updating the resource dependency graph

**Input:** The set of update blocks  $O$ ; the initial path  $p_{init}^f$  and the final path  $p_{fin}^f$  for each flow  $f \in F$ .

**Output:** The resource dependency graph  $G_o$  and an indicator variance that indicates whether the update is feasible or not.

```

1:  $\sigma_{max} = \sigma = 0$ 
2: for each  $o_j^f(t) \in O$  do
3:   Update  $o_j^f$  at  $t$ 
4:   for each  $\langle u, v \rangle \in p_{init}^f$  do
5:      $G_o = G_o \setminus \{\langle u, v \rangle \rightarrow o_j^f(t)\}$ 
6:      $C'_{u,v} = C'_{u,v} + d_f$ 
7:      $\sigma = \sigma + \sigma_{u,v}$ 
8:   for each  $\langle u, v \rangle \in p_{fin}^f$  do
9:      $G_o = G_o \cup \{\langle u, v \rangle \rightarrow o_j^f(t)\}$ 
10:     $C'_{u,v} = C'_{u,v} - d_f$ 
11:    if  $C'_{u,v} < 0$  then
12:      return  $-1$ 
13:    if  $\sigma > \sigma_{max}$  then
14:       $\sigma_{max} = \sigma$ 
15:     $\sigma = 0$ 
16: return  $\sigma_{max}$ 

```

## V. EXPERIMENTAL EVALUATION

We evaluate our scheduling algorithm using both prototype implementation and large-scale simulation.

**Benchmark schemes:** We compare the following schemes with our algorithm.

- **TPP:** The two-phase update protocol [28] that we use VLAN ID as version number in our experiments.
- **NOP:** The node ordering protocol [23] that avoids black hole and forwarding loops [17].
- **Chronicle:** Our scheduling algorithm in Algorithm 2.
- **OPT:** The optimal solution of the MUTP integer program obtained using branch and bound.

The traffic used in our evaluation is generated in [1], and we change the flow demand to simulate traffic variations. Given the demand, we calculate the initial and final routing to maximize the network utilization [10].

### A. Implementation and Testbed Emulations

**Implementation:** We develop a prototype of our algorithm using OFSoftSwitch and Dpctl [5] as Openflow switches and the controller. Now we describe how to perform accurate timing in our algorithm. We first obtain a solution to MUTP using Algorithm 2. Next we send update messages to each switch. We first send an `OFPBCT_OPEN_REQUEST` message to open a bundle, and then send a sequence of

`OFPT_BUNDLE_ADD_MESSAGE` messages in order to modify the rules. Modifications are stored in a temporary staging area without taking effect. Next we close the bundle. Finally when a bundle is committed, the modifications will be applied to the switch at a specific time point.

**Testbed setup:** Our experiments are performed using a 5-server testbed, equipped with two Intel E5-2650 CPUs with 12 cores and 64 GB memory. Each server runs a software-based Openflow switch [5]. We adopt a small scale topology with 5 switches and seven 1 Gbps links as illustrated in Fig. 2. We use the Network Time Protocol (NTP) to synchronize the clocks of all the switches. The `scheduled` bundles feature [3] is used to guarantee accurate timing. We use `pktgen` to generate different numbers of UDP flows in each run. The aggregate flow rate is 1 Gbps in total. The forwarding rules are installed and updated via Dpctl API [5].

**Experiment results:** Fig. 8 shows the total update time for different schemes. As Openflow `barrier` feature cannot provide accurate acknowledgments [19] to indicate the completion of update operation, we use `tcpdump`—a powerful packet analyzer—to confirm when the new rules take effect. We can observe that the update time of TPP and NOP is around 55 ms on average, while Chronicle is around 20 ms. Chronicle can hence reduce the update time by 63% compared to NOP and TPP. This demonstrates that Chronicle can leverage the benefits of accurate timing to accelerate the update process, reducing the time overhead resulting from the wait-invoke pattern.

### B. Simulation

We also conduct extensive simulations to evaluate our algorithm at scale.

**Setup.** In addition to the small-scale topology used in our testbed, here we use a large-scale synthetic scale-free topology that is randomly produced by the `scale_free_graph` function [2]. There are 100 switches and 586 10 Gbps links in total. We randomly generate different numbers of flows with demand ranging from 100 Mbps to 500 Mbps for unique source-destination switch pairs. Accordingly, we adjust the link capacity in order to ensure that the congestion-free condition holds in both initial and final stages. We run the algorithms on a server with Intel Xeon CPU and 15 GB memory. Each data point is an average of ten runs.

**Experiment results:** We first investigate the number of congested flows during the entire update process. We can see that in Fig. 9, as the number of flows increases, NOP and TPP yield significantly more congested flows, while that of Chronicle is

zero all the time. Specifically, the number of congested flows for NOP and TPP is 471 and 478 respectively, when the number of flows is 2500. The congested flows for NOP and TPP account for around 20% of the total flows in the network. This demonstrates that Chronicle takes full advantage of accurate timing and completely avoids congestion by assigning different update points for each flow.

Fig. 10 shows the number of changed rules during update. We define the number of changed rules as the number of rules that needs to be added, modified or deleted during the update. Essentially this measures the number of operations, as well as the number of flow table entries required to perform the update. We observe that TPP induces more changed rules than NOP and Chronicle. When the number of flows is 2000, the changed rules of TPP, NOP and Chronicle is 16569, 9069 and 9069 respectively. TPP is almost twice as that of NOP and Chronicle. This is because TPP relies on different version numbers to indicate two stages during the update. This process involves more update (add/remove) operations compared with NOP and Chronicle.

Finally, we show the additive optimality gap. Fig. 11 shows the box plot of additive optimality gap between Chronicle and OPT as the number of flows increases. We can see that the additive optimality gap in the worst-case is 10 time units and that in the average-case is 5 time units. In general, the update time of Chronicle is near optimal compared to OPT.

## VI. CONCLUSION

We studied the problem of minimizing the makespan in timed SDNs and proved its hardness. We proposed Chronicle to find a feasible update sequence in polynomial time. Our evaluation results show that Chronicle can significantly reduce the update makespan.

## ACKNOWLEDGMENT

The authors would like to thank anonymous reviewers for their valuable comments. The work is partly supported by China 973 projects under Grant Number 2014CB340303, the National Science and Technology Major Project of China under Grant Number 2017ZX03001013-003, the Fundamental Research Funds for the Central Universities under Grant Number 0202-14380037, the National Natural Science Foundation of China under Grant Numbers 61772265, 61602194, 61672353, 61502229, 61672276, and 61321491, the Collaborative Innovation Center of Novel Software Technology and Industrialization, and the Jiangsu Innovation and Entrepreneurship (Shuangchuang) Program, U.S. Natural Science Foundation under Grant Numbers CNS 1757533, CNS1629746, CNS 1564128, CNS 1449860, CNS 1461932, CNS1460971, and IIP 1439672.

## REFERENCES

[1] Fnss. <http://fnss.github.io/>.  
 [2] Networkx. <https://networkx.github.io/>.  
 [3] Openflow switch specification. <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf>.

[4] Broadcom trident. <http://www.broadcom.com/docs/features/StrataXGSTRidentIIpresentation.pdf>, 2012.  
 [5] Cpqd ofsoftswitch. <https://github.com/CPqD/ofsoftswitch13>, 2014.  
 [6] S. A. Amiri, S. Dudycz, S. Schmid, and S. Wiederrecht. Congestion-free rerouting of flows on dags. In *ICALP*, 2018.  
 [7] R. Ben-Basat, G. Einziger, R. Friedman, M. C. Luizelli, and E. Waisbard. Constant time updates in hierarchical heavy hitters. In *SIGCOMM*, pages 127–140, 2017.  
 [8] S. Brandt, K.-T. Foerster, and R. Wattenhofer. On consistent migration of flows in SDNs. In *INFOCOM*, 2016.  
 [9] S. Chopra and M. R. Rao. The partition problem. *Math. Program.*, 59:87–115, 1993.  
 [10] R. Cohen, L. Lewin-Eytan, J. Naor, and D. Raz. On the effect of forwarding table size on sdn network utilization. In *INFOCOM*, 2014.  
 [11] N. Feamster, J. Rexford, and E. W. Zegura. The road to SDN: an intellectual history of programmable networks. *Computer Communication Review*, 44(2):87–98, 2014.  
 [12] K. Foerster, A. Ludwig, J. Marcinkowski, and S. Schmid. Loop-free route updates for software-defined networks. *IEEE/ACM Transactions on Networking*, 26(1):328–341, 2018.  
 [13] K.-T. Foerster, S. Schmid, and S. Vissicchio. Survey of consistent network updates. In *ArXiv Technical Report*, 2016.  
 [14] L. R. Ford and D. R. Fulkerson. Construct maximal dynamic flows from static flow. *Operation Research.*, 6:419–433, 1958.  
 [15] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving high utilization with software-driven wan. In *SIGCOMM*, pages 15–26, 2013.  
 [16] V. Jalaparti, I. Bliznets, S. Kandula, B. Lucier, and I. Menache. Dynamic pricing and traffic engineering for timely inter-datacenter transfers. In *SIGCOMM*, pages 73–86, 2016.  
 [17] X. Jin, H. H. Liu, X. Wu, R. Gandhi, S. Kandula, R. Mahajan, M. Zhang, J. Rexford, and R. Wattenhofer. Dynamic scheduling of network updates. In *SIGCOMM*, pages 539–550, 2014.  
 [18] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972.  
 [19] M. Kuzniar, P. Peresini, and D. Kostic. Providing reliable FIB update acknowledgments in SDN. In *CoNEXT*, pages 415–422, 2014.  
 [20] H. H. Liu, S. Kandula, R. Mahajan, M. Zhang, and D. Gelernter. Traffic engineering with forward fault correction. In *SIGCOMM*, 2014.  
 [21] H. H. Liu, X. Wu, M. Zhang, L. Yuan, R. Wattenhofer, and D. A. Maltz. zupdate: updating data center networks with zero loss. In *SIGCOMM*, pages 411–422, 2013.  
 [22] A. Ludwig, S. Dudycz, M. Rost, and S. Schmid. Transiently secure network updates. In *SIGMETRICS*, pages 273–284, 2016.  
 [23] A. Ludwig, J. Marcinkowski, and S. Schmid. Scheduling loop-free network updates: It’s good to relax! In *PODC*, pages 13–22, 2015.  
 [24] T. Mizrahi and Y. Moses. Software defined networks: It’s about time. In *INFOCOM*, pages 1–9, 2016.  
 [25] T. Mizrahi, O. Rottenstreich, and Y. Moses. Timeflip: Scheduling network updates with timestamp-based TCAM ranges. In *INFOCOM*, 2015.  
 [26] T. Mizrahi, E. Saat, and Y. Moses. Timed consistent network updates. In *SOSR*, pages 21:1–21:14, 2015.  
 [27] Z. A. Qazi, C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu. SIMPLE-fying middlebox policy enforcement using SDN. In *SIGCOMM*, 2013.  
 [28] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker. Abstractions for network update. In *SIGCOMM*, pages 323–334, 2012.  
 [29] S. Vissicchio and L. Cittadini. Flip the (flow) table: Fast lightweight policy-preserving sdn updates. In *INFOCOM*, pages 1–9, 2016.  
 [30] J. Zheng, G. Chen, S. Schmid, H. Dai, and J. Wu. Chronus: Consistent data plane updates in timed sdn. In *ICDCS*, 2017.  
 [31] J. Zheng, H. Xu, X. Zhu, G. Chen, and Y. Geng. We’ve got you covered: Failure recovery with backup tunnels in traffic engineering. In *ICNP*, 2016.