

FRANCIS: Fast Reaction Algorithms for Network Coordination In Switches

Paper #56, 12 pages body, 22 pages total

Abstract

Optimizing the reaction to network events, which is critical in tasks such as clock synchronization, multicast, and routing, becomes increasingly challenging as networks grow larger. To improve the reaction time compared to centralized solutions, the theory community has made significant progress in the design of *message-passing* algorithms that leverage all nodes for distributed computation, and the advent of programmable switches makes it now possible to materialize them.

We propose FRANCIS, a framework and associated libraries for running message-passing algorithms on programmable switches. It features primitives that allow easy integration of such algorithms for quickly reacting to network events while optimizing resource consumption. We use FRANCIS to implement event reaction solutions that improve clock synchronization, source-routed multicast, and routing and demonstrate up to $18\times$ reduction in reaction time.

1 Introduction

Quick reaction to network events is critical to ensure the performance of many applications such as distributed transactional databases [20, 62] and multicast-based machine learning [43]. For such reactions, these applications commonly rely on lower-level distributed networking protocols such as spanning tree [29], shortest path routing [30, 52], clock synchronization [50], and networked consensus protocols [65]. Such protocols must be optimized as networks grow larger (e.g., connecting 100K servers) and faster (e.g., with a few microseconds RTT [2]) and as applications require ultra low latency (e.g., distributed transactional database [20, 62], memcached [26], or RAMCloud [55, 56]). Compared to centralized SDN approaches [32, 36], distributed algorithms inherently adapt more quickly to transient network state changes such as link failures and short traffic bursts.

To leverage the distributed nature of networks, the theoretical computer science community has extensively studied the design of distributed message-passing algorithms for a variety of fundamental problems, including those mentioned above.

Here, by message-passing algorithms, we mean a computational approach where nodes exchange short messages with their neighbors to coordinate and solve problems in a decentralized manner. This model, in which the goal is to minimize the number of messaging rounds, is called the CONGEST model. Many core issues to the networking community have been explored within this framework, such as spanning tree protocols [23, 57], shortest paths computation [7, 18], consensus [58], and graph queries [25, 40]. However, the intricate nature of these solutions, combined with the constant emergence of new advancements, has made them challenging to implement on traditional hardware switches.

The recent emergence of programmable switches [1, 3] provides a natural platform for deploying message-passing algorithms for low-latency distributed systems. Today, the data plane of programmable switches can provide nanosecond-level packet processing time and terabit-level bandwidth. Moreover, switches are first to observe network events such as link failure or load imbalance and are thus best placed to provide fast reaction for network protocols [19, 31].

In this paper, we present a general framework for distributed reaction to network events, which we dub FRANCIS (Fast Reaction Algorithms for Network Coordination In Switches). FRANCIS enables easier conversion of current and future theoretical results in message-passing algorithms into viable real-world distributed implementations on networked switches. It introduces the following three-phase paradigm of distributed reaction. First, the switches detect a network event in the data plane, initiating the reaction protocol. Next, a fast recovery phase is triggered, in which we run distributed message-passing algorithms on the switches to restore connectivity and functionality. Finally, we distributedly optimize the solution to bring our system back to a stable state that provides better efficiency.

FRANCIS makes the following contributions towards switch implementations of message-passing algorithms. First, we establish a framework (Section 3) with algorithmic abstractions and a library of primitives to allow easier specification of reaction algorithms. Second, we provide an execution en-

gine (Section 4) that facilitates implementation on switches by overcoming hardware limitations, handling network unreliability, managing resource constraints, and introducing optimizations to speed up reactions.

We demonstrate the benefit of FRANCIS in the context of clock synchronization, source-routed multicast, and routing, by implementing Tofino [3] and simulation prototypes for each use case respectively. For clock synchronization, we compare with Sundial [45] and PTP [6], and demonstrate how FRANCIS prototype quickly recovers from multiple failures, improves the time uncertainty bound by up to $9.3\times$, and lowers the clock drift. For multicast, compared to Elmo [61] and Orca [22], our prototype uses packet headers that are smaller by up to 52.0% and achieves up to $18\times$ faster reaction to network events. For routing, we adapt Contra using FRANCIS, achieving approximately $10\times$ faster reaction speed than the original Contra protocol.

The rest of the paper is organized as follows. Section 2 describes FRANCIS’s objectives for network event reaction, surveys the most relevant works, and justifies our choice to focus on synchronous algorithms. Section 3 explains our three-phase reaction framework and how FRANCIS’s primitives allow modular and easy algorithmic specification of a reaction scheme. Section 4 introduces FRANCIS’s message-passing execution engine that facilitates implementation on programmable switches. Section 5 showcases how FRANCIS is deployed in real-world examples. Sections 6 and 7 cover implementation details and evaluation of FRANCIS. We then provide an extended description of related work in Section 8 and discuss future work in Section 9. Finally, many more details and discussions appear in the appendices.

2 Background and motivation

In modern data centers, with faster and larger networks [63] and greater demand for low latency for distributed user applications, it becomes increasingly crucial to design a system that reacts efficiently to network events such as (possibly multiple concurrent) link and switch failures, packet loss, sudden increases in link utilization, connections to new switches, etc.

2.1 Objectives for reacting to network events

We describe FRANCIS’s general objectives for reacting to network events: fast reaction, low resource overhead, and resiliency to concurrent events.

Fast reaction. Because many applications have little tolerance for even a short period of disturbance or disconnection caused by a network event, fast reaction and recovery are critical. Events such as link failures can happen as frequently as once every few minutes in a large-scale data center network [28, 51]. These and other more transient events, such as traffic bursts, are becoming more frequent as networks grow larger [63], increasing the need for fast reactions. We justify

the importance of fast reaction with the following examples (which are expanded in Section 5).

Clock synchronization. Clock synchronization [6, 27, 37, 45, 67] is a key component for applications such as distributed transactional databases [20, 62]. These applications require the synchronization system to provide accurate timestamps to minimize their transaction delays. Failures which is recovered slowly cause severe drifts Δt between clocks, and accordingly, as transactions must wait for Δt to guarantee consistency [20], such drifts lead to degraded transaction delays.

Multicast. Multicast [21, 22, 44, 61] is a common one-to-many transmission pattern that is widely deployed. As the number of multicast groups (rules) can be orders of magnitude larger than the number of links [61], a switch or link failure can affect the connectivity of many groups. Applications like distributed machine learning systems [43] running on top of these failed groups can suffer from degraded performance during the reaction time, which in turn leads to increased training costs and longer training time.

Low resource overhead. When reacting to network events, whether operating in the data plane or in the control plane, the message overhead generated by the distributed protocols should be minimized to reduce congestion. Further, as switches need to support other operations, it is also essential to optimize the amount of other resources such as the memory and compute operations.

Resiliency to concurrent events. Network events are sometimes coupled with each other, in which case a single incident takes down multiple devices or links [28]. This poses the challenge of being resilient to handling multiple concurrent events, ideally without adding assumptions to the pattern of such concurrent events and being applicable to any topology [60, 63, 64]. We note that concurrent events of an incident can happen asynchronously, *i.e.*, new events happen while FRANCIS reacts to earlier events of the same incident.

2.2 Existing network event reaction systems

We now summarize previous works on reacting to network events and point out their limitations on why they fail to fulfill our objectives in Section 2.1.

SDN-based methods. A common practice is to compute reaction configurations to network events in the centralized [24] SDN control plane. While centralized methods bring benefits such as separated management and optimized configuration solutions, the *reaction time* is orders of magnitude slower than distributed methods [33, 41, 45]. The bottlenecks as identified in [41] include time to gather and transmit relevant information for all switches through out-of-band networks, and to install new rules at the switch level. For example, for controller-based network event reaction solutions such as Sundial [45] in clock synchronization and Elmo [61] in source-routed multicast, it is reported to require 10s to 100s of milliseconds in data center networks, unacceptable for many applications.

Decentralized SDN-based solution [41] seems promising in reducing the reaction time in WAN. While FRANCIS’s framework could also be extended to running message-passing in switches’ control-plane (CP), we argue that data-plane (DP) message-passing achieves much lower latency. The high latency for distributed CP [68] mainly comes from low processing speed of switches’ CPU and extra delay in delivering packets from DP to CP (100Kops and $\sim 10\mu\text{s}$ for Tofino [3]).

Backup plan based solutions. Other works propose to handle network events via a backup (failover) plan computed by an SDN controller and installed on the switches. When an event happens, each impacted switch executes the backup plan directly at the data plane, thus optimizing the reaction time. However, such solutions are typically not resilient to multiple concurrent events, as the large number of event combinations requires an excessive amount of time to compute and excessive memory at the switches to encode each plan. For example, Sundial [45] for clock synchronization proposes to react to a *single* failure by using the backup plan to form a new synchronization tree that replaces the affected one. However, *when a second failure happens*, this may require falling back to a significantly slower controller-based solution to compute a new backup plan. Other solutions using backup plans, such as F10’s rerouting protocol [48], also require additional assumptions, such as having a Clos topology and falling back to the controller if there is *more than one failure*.

Periodically run solutions. Another line of work includes solutions that execute periodically, regardless of whether there is a network event or not. Essentially, these solutions are event-agnostic and thus commonly have *higher message overheads*. Contra [33] (which generalizes a distance-vector routing protocol for unicast to support custom performance metrics) and PTP [6] (which runs the BMCA algorithm for clock synchronization tree construction) are two examples. Such protocols are generally inefficient as running the algorithm periodically wastes bandwidth if there is no network event.

2.3 Synchronous versus asynchronous

Message-passing algorithms can be either synchronous or asynchronous. FRANCIS is mainly designed and optimized for synchronous algorithms; these are generally easier to design and optimize, as we next explain.

For synchronous message-passing algorithms, communication takes place over rounds. At the beginning of each round, each node receives all messages sent by its neighbors in the previous round. Next, each node can perform some computation and update its local state. At the end of the round, each node can send messages to its neighbors (potentially a different message for each neighbor), which are delivered at the next round. This framework has been formalized theoretically in, e.g., the CONGEST model [46, 59], where the packet size is restricted to be $O(\log n)$ bits where n is the number of network devices (e.g., switches). Syn-

chrony can be achieved by adding synchronizers [9] to an asynchronous network, as we elaborate on in Section 4.1.

In contrast, asynchronous algorithms act on a per-message basis: incoming messages trigger updates in the switch, which performs some computation and may send new messages to one or more neighbors. As mentioned, Contra [33] is essentially an asynchronous protocol where a switch’s shortest-path entries are updated upon receiving probe messages, and the resulting messages may be sent to neighbors for further updates.

We motivate FRANCIS’s focus on supporting synchronous protocols by comparing them in terms of convergence time, bandwidth usage, and other properties. A potential benefit of asynchronous algorithms is that they can send messages with updated information immediately after a message arrives, without waiting until the subsequent communication round. Synchronous protocols have to wait for all messages in a round to arrive, potentially leading to straggler issues that may increase convergence time e.g., due to packet loss or local computations. We expect this to be less of an issue in practice, as FRANCIS is meant for a data center setting where loss is minimal and uses lightweight local computations. On the other hand, asynchronous protocols typically use additional bandwidth by sending intermediate values that may not have been necessary to send because of later received information. With restrictions on the bandwidth usage, this effect can lead to longer convergence time. (Of course, one can reduce wasted bandwidth by introducing waiting to make an asynchronous algorithm more like a synchronous one.) For example, we adapt Contra’s asynchronous routing algorithm into a synchronous version, and our results show that our basic synchronous version converges $1.87\times$ faster (Table 6).

We additionally argue that synchrony allows for easier protocol design and optimization. Specifically, the round-based granularity of synchronous distributed algorithms allows us to easily build high-level algorithmic abstractions and primitives that simplify users’ design of their network event reaction algorithms (see Section 3). Similarly, round-based granularity appears more amenable to certain optimizations in our framework. One example is message packing, which takes advantage of the fact that we may have multiple algorithm instances running the same algorithm in parallel, so that instances can be synchronized together to pack many messages into a single packet per port, as we describe in detail in Section 4.3.

3 The FRANCIS Framework

This section presents the conceptual and algorithmic contributions of FRANCIS towards our proposed objectives in Section 2.1. Section 3.1 introduces the three-phase distributed reaction framework. Specifically, FRANCIS’s detector and its mechanism of reinvoking reaction address handling *concurrent events* with any temporal and topological patterns; the fast recovery phase prioritizes *fast reaction* using more efficient message-passing algorithms. Section 3.2 elaborates the

Phase	Clock synchronization	Source-routed multicast	Routing
Detection	Passive, based on the clock synchronization messages	Active	Active
Fast recovery	Construct a flooding spanning tree and initializing the synchronizer.	Flooding spanning tree construction and initialization. Then Synchronously build a multicast subtree for each group.	Run F10 [48]’s Pushback Flow Redirection algorithm to find new routes for affected flows efficiently.
Distributed optimization	Run BFS from multiple roots. For each tree, calculate its center and the depth if we select it as the root. Elect the min-depth tree.	Construct a new graph spanner, and build a shortest-path tree on the spanner for each group. Use greedy set-cover to minimize the number of edges in each tree. (Note that this is possible as the tree does not span the whole graph.)	Run a synchronous Contra version. In each round, wait for the neighbors in the product graph to send updates. Forward the updated entries to neighbors at the end of the round.

Table 1: The three reaction phases of our three use cases.

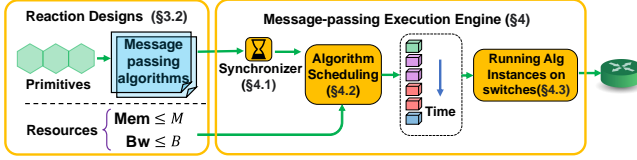


Figure 1: An overview of FRANCIS. Algorithm scheduling is a module that generates *at runtime* the schedule to perform the next computation or communication actions (shown as cubes).

algorithmic abstractions and primitives to facilitate specifying message-passing reactions. It also enables explicit specification of resource constraints to achieve low resource overhead.

3.1 Three phase distributed reaction

We now describe our detection phase and justify why we need both fast recovery and distributed optimization for reaction. The high-level description of our reaction phases for our use cases is summarized in Table 1.

Detection. Our detector listens to events such as failures and congested links through timer-based approaches, which periodically exchange messages between neighboring switches and report an event upon a timeout. FRANCIS works with different protocols [34, 35, 38, 45], including passive detection (which listens to some existing periodic traffic, as in [45]) and active detection (which generates probe packets, as in [38]).

Handling multiple concurrent events. One assumption that simplifies the design of distributed reaction algorithms is to ensure a static network, *i.e.*, no events or failures occur during the execution of the reaction algorithms. This does not hold in a real network, and we thus propose to *reinvoke* the reaction routine from scratch upon new events and invalidate previous runs. This way, the single validated run of reactions always operates on the static network state after the latest network event, where the assumption still holds. FRANCIS’s detector achieves this by listening to any new events and reinvoking a new run of reactions upon a new event.

Fast recovery. To achieve fast reaction, we prioritize restoring basic functionality or connectivity quickly, without necessarily finding optimized solutions or network configurations. The reaction algorithm could be an approximation or heuristic algorithm that converges quickly. For example, if a link that is a part of a clock synchronization tree goes down, we compute a new spanning tree for synchronization quickly through flooding, even though the flooding tree is usually suboptimal and can lead to a high time uncertainty bound.

Distributed optimization. After restoring functionality,

FRANCIS aims to further improve the quality of the solution for better long-term performance of applications. For example, as the clock drift is proportional to the synchronization tree depth, we run the algorithm listed in Table 1 to improve it. The distributed optimization needs not to yield an optimal solution; if desired, a later step may apply further centralized optimization of the solution.

3.2 Designing FRANCIS reactions

A synchronous message-passing reaction scheme includes two basic components: the specification of message-passing algorithms for reaction and their resource constraints. Here, we assume that synchrony is guaranteed by our synchronizer module (as detailed in Section 4.1).

The specification of a message-passing algorithm includes communications with neighbors and switch-local computations. For communication, the switch sends up to one message per neighbor for each round; for computation, the switch processes incoming messages either in a streaming or batched fashion. With *streaming*, the switch updates its state upon the arrival of each message. However, some algorithms are not amenable to streaming computation. In that case, we can use *batching*, where we store messages in memory and perform computations upon receiving all messages from neighbors. The downside of batching is the need for recirculation to process messages from all neighbors, which increases the total reaction time. We elaborate on the two options in Appendix B.

We note that a single message-passing algorithm may correspond to the simultaneous execution of multiple algorithm instances. For example, in multicast, we run the same (but many instances of) multicast tree recovery algorithm for each affected multicast group. As discussed in the following section, FRANCIS makes parallel execution of multiple instances feasible, which not only reduces the reaction time but also allows further optimizations such as message packing.

To allow simpler development of distributed reaction algorithms, we provide a library of some commonly used message-passing algorithm primitives, as summarized in Table 2. Each of them is useful for one or more FRANCIS’s use cases in Section 5. For example, `shortest path tree` serves as a single-source-multi-destination communication structure that minimizes the latency to every destination, which is useful for tasks such as clock synchronization and multicast. `Asynchronous global flooding` runs the communication of flooding to broadcast messages and can optionally bind with functions on top such as running the synchronizer’s initialization across the networks and constructing a (flooding)

Algorithmic primitives	Description
Asynchronous global flooding (Asynchronous extension)	Use flooding to 1) broadcast a short message to all switches 2) construct a flooding spanning tree, 3) initializing the synchronizer throughout the network.
Bottom-up tree aggregation	Aggregate values from the leaves of a spanning tree to the root.
Conditional broadcast and aggregation (as detailed in Algorithm 3)	Maintain a single value v for each switch. Aggregate v with those received from neighbors for each round. Then broadcast the updated value next round provided that the condition function is met.
Shortest path tree	Construct a spanning tree layer-by-layer in a BFS fashion. Each round the current leaf switches look for unvisited switches to be added as a new layer of leaf switches.
Set cover on a shortest path graph (as detailed in Appendix F.2)	In the shortest path graph, greedily select a minimal number of switches in layer i to cover layer $i+1$, i.e., every switch in layer $i+1$ has a selected neighboring switch in layer i .
Synchronous shortest path	Synchronously compute shortest path routes w.r.t. a routing policy.

Table 2: Descriptions of some FRANCIS’s algorithmic primitives. The full version of the table is shown in Table 8.

tree with minimum time; it is useful for all three use cases.

Switch memory and bandwidth are two scarce resources, and FRANCIS as a control program should not consume too much of them. Therefore, we allow users to define the resources dedicated to FRANCIS for network event reaction. FRANCIS’s execution engine then achieves algorithm scheduling (e.g., rate-limiting, batch execution) at runtime to meet the resource constraints (Section 4.2).

4 Message-passing execution engine

FRANCIS proposes to leverage hardware advances of programmable switches [1, 3] and Ethernet [5] and execute message-passing in the switch’s data plane to achieve towards fast reaction. Despite data plane limitations, we enable the efficient execution of many message-passing algorithms on switches, subject to resource constraints and network unreliability. Here we identify the following general challenges, provide platform-agnostic primitives to address them, and implement a proof of concept system on Tofino [3].

First, many message-passing algorithms do not consider network unreliability, such as heterogeneity in packet delay, packet order uncertainty, and transient packet losses.¹ Accordingly, our solutions use synchronization protocols (Section 4.1) and include a mechanism for handling packet loss (Section 4.3). Another challenge is scheduling the execution of multiple algorithm instances under the constraints on memory and bandwidth, minimizing the reaction time. We address this problem in Section 4.2. A further challenge is to introduce optimizations for improved performance on computationally restrictive programmable switches. We elaborate on the two techniques, namely dynamic multicasting and message packing for efficient communication², but also integrate arithmetic approximation and for-loop elimination in FRANCIS for achieving switch compatibility (Section 4.3).

4.1 Algorithm synchronization

FRANCIS’s synchronizer module is a critical component to allow message-passing algorithms to execute synchronously

¹Persistent losses are treated as a new failure and handled in Section 3.

²Message-packing improves the bandwidth utilization rate and is a key to our objective of low bandwidth overhead.

(allowing all switches to know when a round begins) in the presence of network unreliability (e.g., packet loss, uneven queues, etc.). Different synchronizers present different trade-offs between message overheads, convergence, and synchrony properties. Here we explain the α synchronizer we implemented, and defer discussion on other choices to Appendix D.

Before a switch starts its r th round, it has to confirm that all neighboring switches have finished their $(r-1)$ th round. Hence each switch u broadcasts a short synchronization message to all their neighbors, telling them that it has completed its $(r-1)$ st round. The neighboring switches perform similarly, and u listens to all such synchronization messages until it receives messages from all neighbors. This indicates the completion of round $r-1$ of all its neighbors, and switch u can then proceed to round r . Finally, at the beginning of the (synchronous) algorithm execution, we invoke the `Asynchronous global flooding` primitive (Table 2) to flood messages to all switches and bootstrap the synchronizer on each switch, which initializes its local round number to 1 and begins the message-passing execution.

4.2 Algorithm scheduling

As mentioned in Section 3.2, FRANCIS aims to limit both bandwidth and switch memory consumption. It also aims to support running multiple algorithm instances in parallel, which means that we have to split these resources between the instances. This is achieved by our scheduling module.

Memory constraints. Ideally, full parallelization of algorithm instances achieves the best reaction time, since different instances are bottlenecked by the bandwidth of different subsets of links. However, the allocated memory for FRANCIS can be insufficient to support this. To that end, at runtime, FRANCIS picks the maximal number of algorithm instances that have not yet been completed and can fit within the switch’s memory constraint. Such a set of instances is called a batch. When a batch terminates, our scheduler picks the next batch of instances to run. We call this batch execution. We note that most message-passing algorithms, including those in our use cases, have a fixed memory requirement known at compile time, and we assume this is the case for all instances here. As a result, there will never be a need to preempt an algorithm instance due to insufficient memory.

Bandwidth constraints. Next, we need to rate-limit the scheduled instances at runtime to ensure we adhere to the bandwidth constraint. This is done by storing messages in the switches’ memory whose transmission is postponed to avoid exceeding the link’s bandwidth usage constraint. We then use the packet generator to generate packets that retrieve stored messages at a frequency that adheres to this constraint. Here we do not need dynamic memory allocation since each synchronous message-passing algorithm (and the `Asynchronous global flooding` algorithm primitive) instance only needs




Figure 2: Running algorithm instances on programmable switches.

to cache up to a fixed number of messages (the number of ports of the switch). To share bandwidth among different algorithm instances, we choose which instances next transmit on each link in a fair-sharing way. We note that this is a design choice, and one can program other policies.

4.3 Switch compatibility and optimizations

We now discuss challenges in running FRANCIS algorithms in programmable switches. Figure 2 depicts the key designs, in which we highlight our main contributions.

Computation inside switches While not all message-passing algorithms are implementable in the data plane of programmable switches, the following primitives ease adapting such algorithms.

1. Approximating complex arithmetic operations. **Challenge.** Many distributed algorithms involve operations like multiplication and floating-point arithmetic, which may not be natively supported by the switch’s architecture.

Solution. We enable complex arithmetic via approximate arithmetic operations, following previous works [13, 66].

2. Eliminating for-loops. **Challenge.** Message-passing’s computation is specified as a subroutine that executes after all messages from neighbors arrive, but switches are generally not capable of looping over all of them in a single pipeline pass.

Solution. The computation of many of the algorithms of interest can be done in a streaming way, one packet at a time. While other methods can be implemented using recirculation, in all of our use cases, we did not require it.

Efficient communication across switches

1. Enabling dynamic multicast. **Challenge.** The theoretical models allow a switch to send different messages to different subsets of neighbors. To achieve this, we establish dynamic multicast rules *on the fly*, i.e., at runtime, we may send different packets to an arbitrary set of destination ports as determined by the algorithm. However, in programmable devices like Tofino [3], the switches’ CPU can only configure a limited number of *static* multicast rules (2^{16} for Tofino) before running the distributed algorithm. In an $\ell > 16$ port switch, there are $2^\ell > 2^{16}$ potential port subsets, which means that we cannot pre-configure all groups ahead of time. As using the CPU to modify the groups at runtime is costly (for Tofino, about $10\mu\text{s}$ to reach the CPU [54] and another $10\mu\text{s}$ to update [68]), we require a different approach.

Solution. To support general (any neighbor subset) and scalable (suitable for switches with many ports) messaging, we

introduce *dynamic* multicasting. In Tofino, we store Boolean filter vectors in egress SRAM registers, and modify them at runtime according to the algorithm. To multicast packets to neighbors, we *broadcast* it to all egress ports, and then drop those that are filtered. This approach inevitably generates wasted packets from the traffic manager to the egress pipeline. However, since our multicast is only used on a small portion of the traffic used by the distributed algorithms, this appears quite manageable. Thus, we implement dynamic multicasting for scalability and generality.

2. Handling packet loss. **Challenge.** Most message-passing algorithms assume reliable networks but, in practice, we have packet loss. We cannot fully implement reliable transmission protocols, such as TCP, on switches’ data plane, as switches are not capable of supporting dynamic memory management required for software TCP, and they may not intrinsically support timeouts and complex retransmission primitives.

Solution. We propose a simple retransmission protocol that works on programmable switches. Following our approach for handling memory and bandwidth constraints (Section 4.2), the switch keeps a copy of packets (the total number is bounded thanks to our algorithm scheduling) in memory until receiving an ACK message, at which point it can free the memory. In Tofino, we leverage the packet generator to trigger probes to the egress pipeline that actively check (in a round-robin fashion) if a packet timed out and performs retransmissions.

3. Message packing. **Challenge.** Generally, messages in CONGEST algorithms are small in size, ranging from a single bit to several bytes. However, Ethernet frames carrying the messages have a minimum size of 64 bytes. Dedicating an Ethernet frame to a single message can therefore yield a low bandwidth utilization rate, with only a small portion of the frame being the actual message and a large portion being paddings and headers. In our use cases of multicast and Contra routing [33], we launch thousands of algorithm instances to react to a network event. In such cases, a low bandwidth utilization would result in a longer reaction time.

Solution. Our message packing technique works by concatenating multiple messages from different algorithm instances destined for the same neighbor into a single packet, so that the receiver can process multiple messages at a time. Further, we would ideally like to perform as few arithmetic operations as possible, which minimizes the number of packets that traverse the switch’s pipeline, given that the switches’ arithmetic and pipeline stages are limited. We therefore aim to pack values in such a way that we can perform operations on multiple values with a single arithmetic operation. For example, we could pack two 16-bit values a_1 and a_2 into a single 32-bit word, so that incrementing both a_1 and a_2 by 1 can be achieved by adding $2^{16} + 1$ to $a_1 * 2^{16} + a_2$ (assuming no overflows).

Additionally, the receiver needs to know which algorithm instances of the messages are packed in the current packet. The sender therefore adds values ℓ, r and a bitmap to each packet indicating that it carries messages from instances IDs

Topology	FatTree	Jupiter	Jellyfish
Number of Affected Groups	1344	2147	1175

Table 3: The maximal number of multicast groups affected by a link failure over different datasets.

$\subseteq \{\ell, \ell + 1, \dots, r\}$. Again, message packing is designed specifically for synchronous algorithms, as it would be harder and take more bandwidth to apply it to asynchronous algorithms that act on each message arrival, whose patterns are irregular.

5 Use cases

We next show how FRANCIS can be used to easily implement distributed algorithms using three real-world use cases of our framework: clock synchronization, source-routed multicast, and routing. For the first two use cases, we propose our own data-plane reaction protocols to handle network events. For the routing case, we adapt and optimize Contra [33], a state-of-the-art data-plane protocol, to use FRANCIS. For these applications, we focus on link failures and switch failures as the important relevant network events.

5.1 Clock synchronization

Background. Clock synchronization is achieved by periodically exchanging timestamp information between adjacent devices. In state-of-the-art designs (e.g., [6, 45]), the information is carried by *synchronization packets* through the *synchronization spanning tree*. The tree’s root acts as the ‘grandmaster clock’ that serves as the time reference to other clocks. The precision of the clock of device k can be measured by its *time uncertainty* $\epsilon^{(k)}(t)$ [45] at time t , calculated as

$$\epsilon^{(k)}(t) = \epsilon_0 \cdot \text{depth}^{(k)} + (t - \tau_{\text{sync}}^{(k)}) \cdot \text{max_drift_rate}. \quad (1)$$

Here, $\tau_{\text{sync}}^{(k)}$ is the *last* time device k received a synchronization message, max_drift_rate is the maximum drift rate of a clock, ϵ_0 is the noise inherent in hardware, and $\text{depth}^{(k)}$ is the number of hops between device k and *root*. Finally, the overall *time uncertainty bound* is defined as $\epsilon(t) = \max_k \epsilon^{(k)}(t)$, which measures the maximal discrepancy between any pair of clocks. During normal operation,³ ϵ is primarily determined by the *tree depth*; when failures break connectivity of the synchronization tree, ϵ grows proportionally to the failure time $t - \tau_{\text{sync}}^{(k)}$ until a new tree is reconstructed.

Main objectives. As suggested by Sundial [45], to achieve fault-tolerant clock synchronization, we aim to minimize the peak ϵ in case of failures, and keep ϵ consistently small during normal times (without failures). We therefore think of ϵ in two phases: the peak ϵ obtained during event reaction, and ϵ after the distributed optimization. During fast recovery, we aim to construct a clock synchronization tree as quickly as possible to prevent large peak ϵ . We then further reduce ϵ by optimizing the depth of the synchronization tree in the distributed optimization phase.

³We use ϵ for $\epsilon(t)$ when convenient where t is implied.

Algorithm design. We now describe our proposed algorithms, following the key ideas summarized in Table 1.

(1) The detection phase. We adopt Sundial’s [45] passive detection protocol where we listen to synchronization messages and detect a failure on timeout. The original protocol, however, does not find which specific devices have failed, but only concludes that some upstream links or switches have failed. To determine the specific failed devices, we propose a simple ping-on-timeout technique, as discussed in Appendix G.

(2) The fast recovery phase. Upon detecting a failure, we restore connectivity by constructing a fast recovery spanning tree using asynchronous flooding. This spanning tree can then deliver synchronization messages until we complete the distributed optimization phase. The algorithm can be easily implemented leveraging FRANCIS’s `Asynchronous global flooding` algorithmic primitive (Table 2). This primitive also initializes the synchronizer for running synchronous algorithms in the next phase. Concurrent failures are handled according to the approach specified in Section 3.1.

(3) The distributed optimization phase. In the fast recovery phase, we also randomly sample several roots and broadcast them via `Asynchronous global flooding` to every switch. Then we run the `BFS-based shortest path tree` algorithmic primitive (Table 2) to construct the shortest path spanning tree for each root. We then run the `Bottom-up tree aggregation` primitive to calculate the depth for each subtree. Following this, we move the root to the tree’s center. Finally, we pick the tree with minimal tree depth via leader election as our final synchronization tree, leveraging the `Conditional broadcast and aggregation` primitive. We give the pseudo code as Algorithm 4 in Appendix C.

The shortest path tree algorithm only returns an optimal-depth spanning tree for a given root. In practice, one may choose which and how many roots to use. We observe from our experiments (Appendix H.5) that sampling a few candidate roots often suffices for finding an optimal-depth tree with $> 99\%$ probability in data center topologies. We note that running shortest path tree on all nodes creates excessive network overhead and takes significant time. Optionally, after the distributed optimization, the controller could compute an optimal depth synchronization tree and inform the switches.

5.2 Source-routed multicast

Background. Given a multicast group $\text{group}_i = (\text{src}_i, \text{dst}_i)$, where src_i is the single sender host and dst_i is the set of receiver hosts, the goal is to deliver the packet to dst_i through a *multicast tree*, T_i , duplicating the packet at tree branches as needed. In this section, we base our design on *source-routed* multicast. In this setting, source hosts *encode* the multicast tree as rules that are added to the headers of multicast packets. The switches then forward packets according to the rules in the packets. A network failure can affect thousands of

multicast groups by breaking the connectivity of their corresponding multicast trees (Table 3), requiring reconstructing a multicast tree and rules for each affected group.

Main objectives. We classify the goals for a fault-tolerant source-routed multicast system as follows. In the fast recovery phase, the system should ideally restore functionality for thousands of affected groups in 100s of microseconds. In the distributed optimization phase, we wish to minimize packet delay. The delay is determined mainly by the depth of the multicast tree, so we prefer shallow trees. Finally, we aim to minimize traffic overhead that is determined by both the size of the multicast header and the size of the multicast tree.

Algorithm design. We now describe our approach, following the key ideas summarized in Table 1.

(1) The detection phase. We use active detection of failures.

(2) The fast recovery phase. With limited bandwidth, fixing thousands of broken multicast trees in parallel incurs a large message overhead and requires significant time. Instead, we build a single fast recovery spanning tree T_0 and reconstruct multicast trees on top of it. This is achieved via flooding (see Table 2). At the same time, switches adjacent to the failure(s) retrieve the affected groups stored in their memory and inform the Top-of-Rack switches of the affected groups.

Once T_0 is constructed, connectivity is guaranteed. The remaining task is to compute multicast trees that are subtrees of T_0 , such that the subtree connects all destinations for each group. We propose the following synchronous algorithm that works in a bottom-up manner on T_0 . Starting from the destination switches, we determine whether switch i is on the multicast tree of group j (it is if one of i 's children on T_0 belongs to group j 's tree) and record switch i 's children for group j . However, to limit the bandwidth, we do not compute each such subtree individually. To that end, we instead propagate participation bitmaps for all groups simultaneously, where each switch i marks whether it belongs to the multicast tree of group j . For a switch that receives participation bitmaps from its children in T_0 , it updates its own bitmap accordingly using bitwise operations. To further reduce overhead, we leverage the message packing technique to place multiple such bitmaps into each packet. We defer further details to Appendix F.1.

(3) The distributed optimization phase. In state-of-the-art source-routed multicast solutions like Elmo [61], senders encode propagation rules on each multicast packet, thus requiring a header overhead driven by topology size and distribution of destinations. To reduce these overheads, FRANCIS leverages *spanners* to encode rules over a small sub-topology. A spanner is a sub-topology $G' \subset G$ of the original topology G such that the stretch $= \frac{dist_{G'}(u,v)}{dist_G(u,v)}$ is bounded.

We use the Baswana-Sen algorithm [11], that given a user-defined stretch parameter, computes a spanner with this stretch in $O((\text{stretch})^2)$ rounds. The benefit of higher stretch values is that they allow more pruning of the topology, thereby decreasing the needed multicast packet header

size, the overall latency of multicast traffic, and the total consumed network bandwidth.

Accordingly, after computing the spanner, we reconstruct a shortest path tree on top of it for each affected group. To that end, we use the algorithmic primitive `Shortest path tree`. Using this tree, we compute the distance of each destination from the source of the group. Ideally, the multicast tree for the group should then be the one with the minimal number of nodes such that each destination is connected to the source via a shortest path. However, finding such a tree is NP-hard, as we show in Appendix E. Instead, we leverage the primitive `Set cover on a shortest path tree`, which is an approximation algorithm to construct a small multicast tree for each group. Optionally, the controller can then spend more time optimizing the spanner and set covers.

5.3 Contra for Routing

We also adapt Contra's routing protocol [33] using FRANCIS, which we later show can reduce the reaction time to network events by up to 90% (Section 7.2.3). Motivated by the limitations of Contra mentioned in Section 2.3, we design a wait-and-merge approach to convert Contra's asynchronous protocol to an equivalent synchronous message-passing protocol that can be implemented in FRANCIS. This conversion also allows us to apply the message packing technique. Both of these approaches save bandwidth, and under bandwidth constraints can lead to significantly faster reaction.

Contra's adaptation to FRANCIS. We have designed a synchronous version of Contra's shortest path algorithm, as detailed in Algorithm 5 in Appendix C. This implementation is part of FRANCIS's library, and we provide an algorithmic primitive for it. Our modified Contra algorithm *waits-and-merges* information from different neighbors that correspond to the same routing entry during a round. Intuitively, the modified protocol prevents the intermediate update results from being sent early, saving bandwidth. We empirically verify this in Section 7.2.3. Using message packing on top of the synchronous message-passing algorithm offers further optimization over the original asynchronous Contra implementation. We regard the algorithm that reconstructs the routes of a given destination dst as one algorithm instance. FRANCIS packs 8 messages with consecutive instance IDs and the same transmitted state into a packet, and FRANCIS processes such a packet in a single pass.

6 Prototype implementation

We have implemented FRANCIS's prototypes both in Tofino P4 [3, 16] and in simulation on top of our own discrete event simulator. The Tofino implementation validates FRANCIS's design in small-scale network topologies and better indicates the real-world performance of FRANCIS, while simulation is ideal for efficiently executing large-scale experiments.

Application	LoC in P4	Stages	SRAM usage	TCAM
Clock Synchronization	1731	12	3.07%	2.43%
Multicast	1616	19	1.48%	1.88%
Contra	1150	18	1.78%	1.30%

Table 4: Lines of code of FRANCIS’s three use cases and their resource usage on Tofino switches for a 64-ary FatTree topology. Note that results for clock synchronization are for Tofino architecture while results for other applications are for Tofino 2.

Tofino emulation. Clock synchronization for Tofino using FRANCIS is deployed onto a testbed with a Wedge 100BF-32X Tofino switch. Inspired by prior works [13, 37], our implementation is able to virtualize multiple logical switches on the single physical switch that form a 3-4-ary FatTree topology (as specified in Table 5). The virtualization is achieved directly in the data plane and thus incurs no additional overhead. Specifically, each logical switch is assigned dedicated ports for each corresponding logical link, and sending packets to a neighboring switch is achieved via looping back traffic to its neighbor’s port.

Resource consumption. We evaluate the resource consumption of FRANCIS’s Tofino prototypes, as shown in Table 4. We note that the SRAM usage is not explicitly restricted and the algorithm scheduling chooses full parallelization accordingly. We fit our clock synchronization prototype into the Tofino architecture and our multicast and Contra prototypes into Tofino 2⁴. We see that all use cases incur minimal consumption of $\leq 3.07\%$ of Tofino’s SRAM budget even without memory restrictions, and $\leq 2.43\%$ of the TCAM budget.

7 Evaluation

We evaluate FRANCIS on the three use cases detailed in Section 5. Additionally, we implement and evaluate F10 routing [48] using FRANCIS and discuss results in Appendix H.4. We first elaborate on the testbed evaluation (Section 7.1) comparing FRANCIS Tofino against state-of-the-art systems. We then proceed to large-scale simulation experiments in Section 7.2 to study FRANCIS’s performance at scale. In Section 7.3, we evaluate running the three applications concurrently in FRANCIS. Finally, we show in Section 7.4 how different parameter configurations affect performance.

7.1 Testbed Evaluation

We first evaluate FRANCIS’s performance in clock synchronization over the Tofino testbed. We compare against PTP’s [6] best master clock algorithm (BMCA), a state-of-the-art approach that automatically rebuilds a clock synchronization topology upon a network event. We choose 3-4-ary FatTree (Table 5, explained in Appendix H.1) as the logical topology that we could fit into our Tofino switch.

Testbed deployment of PTP. State-of-the-art PTP systems generally run BMCA in software, and we follow the setup to

⁴The physical switch of our testbed is in Tofino; multicast and Contra that cannot fit in Tofino are examined with Intel-modeled Tofino 2 architecture.

	3-4-ary FatTree	64-ary FatTree	Jupiter	Jellyfish
#Switches (n)	15	5120	22528	4000
#Edges (m)	21	131072	163840	126950
Diameter (D)	4	4	8	3

Table 5: Characteristics of the evaluated topologies.

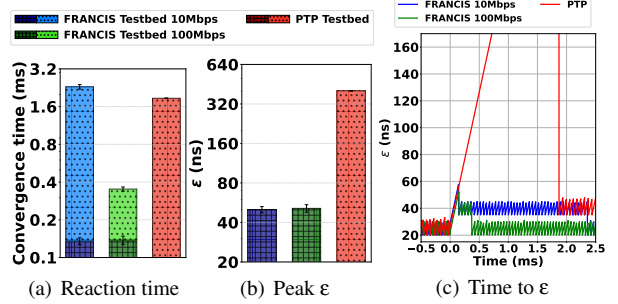


Figure 3: Performance of FRANCIS’s Tofino prototype in failure recovery for clock synchronization, compared with PTP [6]’s best master clock algorithm to handle failures. The dark, cross-patched bars represent fast recovery and the light, dotted bars correspond to distributed optimization. A single link failure happens at 0ms.

run LinuxPTP [4] with the same topology being virtualized on our server. The typical announce interval for BMCA is in order of seconds, but for testing purposes, we push it to the limit of $122\mu\text{s}$, the minimum we could achieve. Setting a shorter interval reduces the reaction time and thus peak ϵ (according to Eq. 1), but incurs more computation and communication costs - an $122\mu\text{s}$ interval corresponds to a total of $\sim 100\text{K}$ packets processed per second over all neighbors.

Other setups. We generate a single link failure at time 0. For FRANCIS, we limit via algorithm scheduling (Section 4.2) per-link bandwidth usage to be 100Mbps and 10Mbps respectively. Other settings are discussed further in Appendix H.1.

Comparisons with PTP. As illustrated in Figure 3(a), FRANCIS achieves prompt reaction to the link failure and successfully restores connectivity of the clock synchronization tree within $170\mu\text{s}$ in the fast recovery phase. The results do not depend on the bandwidth limitation, as the Asynchronous global flooding for fast recovery generates only one prioritized packet per link, for which no extra delay is incurred to meet the bandwidth constraint (Section 4.2). The subsequent distributed optimization phase further optimizes the longer-term time uncertainty bound ϵ by reducing the tree depth (as shown in Figure 3(c)). The overall reaction time as shown is $352\mu\text{s}$ for 100Mbps which is $> 5\times$ faster than PTP with the aforementioned, already optimized setups.

As suggested by Eq. 1, the peak ϵ is proportional to the time when connectivity is first restored. Accordingly, we observe in Figure 3(b) that FRANCIS reaches an order of magnitude lower peak ϵ than PTP. This can be attributed to the reduced time for fast recovery in FRANCIS from re-establishing connectivity with a suboptimal flooding tree.

We validate our simulator by finding the simulation results match with testbed results, as illustrated in Appendix H.2.





Figure 4: Performance of clock synchronization compared with Sundial [45] and PTP [6] with two switch failures.



(a) Time uncertainty bound (ϵ). (b) Message overhead per link.

Figure 5: Time to ϵ and per-link message overhead for clock synchronization with one switch failure on the FatTree topology.

7.2 Large-scale simulation on individual tasks

Next, we evaluate FRANCIS’s performance for each use case via simulation in large-scale data center topologies (Table 5), including $k = 64$ FatTree, Jupiter [63], and Jellyfish [64].

7.2.1 Clock synchronization


Baseline solutions. We compare FRANCIS with the state-of-the-art solutions PTP [6] and Sundial [45].

Experiment setup. We generate one or two switch failures while limiting the per-link bandwidth of FRANCIS to be either 100Mbps or 10Mbps. For PTP, we set the announce interval to be $100\mu\text{s}$. See further details in Appendix H.1.

Experimental results. We first evaluate how fast FRANCIS reacts to switch failures, as depicted in Figure 4 for two switch failures and Figure 14 in Appendix for one failure. We observe that FRANCIS restores connectivity in the fast recovery phase within $200\mu\text{s}$ for both 100Mbps and 10Mbps bandwidth. For Jupiter, the distributed optimization phase lasts for $506\mu\text{s}$ and 2.72ms with 100Mbps and 10Mbps bandwidth respectively. This is orders of magnitude faster than Sundial as its backup plan recovery could not handle multiple failures in different domains, in which case it falls back to the controller.


As shown in Figure 4(b), FRANCIS’s has lower peak ϵ in all topologies when two failures happen. PTP’s peak ϵ is $5.3\text{--}9.3\times$ higher than FRANCIS-10Mbps. Sundial incurs a large peak ϵ ($2 \cdot 10^4$) value in this setting, giving unfavorable performance.

For one switch failure, figure 5 depicts how ϵ and message overhead change over time. We observe that PTP is markedly slower in restoring connectivity than FRANCIS and has higher ϵ and message overhead throughout. Sundial achieves an optimal $150\mu\text{s}$ recovery by switching to the backup plan immediately upon detecting the failure. However,



(a) 1 link failure. (b) 2 link failures.

Figure 6: Reaction time (log-scale) of FRANCIS for source-routed multicast compared with the FatTree-specific Elmo [61].



(a) 1 link failure. (b) 2 link failures.

Figure 7: Reaction time comparison between asynchronous approach and synchronous approach for the fast recovery phase of source-routed multicast.

its ϵ remains suboptimal until the control plane steps in, as the tree of the backup plan is not optimal.

7.2.2 Multicast

Baseline solutions. We compare with state-of-the-art source routed multicast methods Elmo [61] and Orca [22].

Experiment setup. We generate 1 or 2 link failures for our experiments and allocate 200Mbps bandwidth for FRANCIS. We note that Elmo is designed specifically for FatTree topologies, and we are *unable* to adapt it to other topologies. Other settings are shown in Appendix H.1.


Experimental results. As depicted in Figure 6, FRANCIS’s fast recovery and distributed optimization are $11\times$ and $18 - 24\times$ faster than Elmo’s fast and (control-plane based) full recovery, respectively.

Next, to evaluate the spanner algorithm for multicast, we compare additional metrics such as the header size and message overheads in Appendix H.3. As shown, our fast recovery solution has a traffic overhead is $2 - 3.9\times$ lower than Elmo’s unicast-based fast recovery solution, while the distributed optimization solution is on par with both Elmo and Orca [22] while reducing the header size by $1.4 - 2\times$.

Finally, we also implement the asynchronous version of bottom-up tree aggregation for building fast recovery multicast trees. This is achieved by waiting for all of the children in the flooding tree to complete the aggregation and then executing the bitmap aggregation. Figure 7 shows that the reaction time difference between the asynchronous version and the synchronous version is negligible ($< 5\%$). On the other hand, as mentioned earlier, the benefits for FRANCIS to focus on synchronous algorithms include easier protocol design and better system-level optimization.

Methods	Category	Reaction time	Bandwidth Usage
Contra-FRANCIS Optimized	DP Sync	1.720ms	200Mbps
Contra-FRANCIS Unoptimized	DP Sync	9.186ms	
Original Contra	DP Async	17.15ms	
ONOS [14]	CP	145.8ms (expected)	/

Table 6: We compare Contra-FRANCIS with the original Contra protocol and ONOS. Here optimized and unoptimized refer to if we apply message packing or not. DP and CP refer to data plane and control plane respectively.



(a) Comparing FRANCIS-individual (b) Varying total allocated memory with FRANCIS-combined.

Figure 8: Evaluation on running concurrent applications. Here FRANCIS denotes FRANCIS-individual.


7.2.3 Contra Routing

Experiment setup. We compare Contra’s FRANCIS adaptation against both the original Contra [33] and ONOS [14], a control-plane routing system. We also run Contra-FRANCIS without message packing to study effects of our wait-and-merge and message packing techniques separately. We use the FatTree topology and set the bandwidth limit as 200Mbps. **Experimental results.** Table 6 shows that Contra-FRANCIS achieves an order of magnitude faster reaction speed than original Contra which is faster than the control-plane-based ONOS. Further, FRANCIS’s wait-and-merge technique reduces the reaction time by $1.8\times$ and the message packing contributes to a further $5.3\times$ reduction. In general, this suggests that deploying existing asynchronous data-plane protocols into FRANCIS may significantly reduce the recovery time.

7.3 Running Concurrent Applications

Experiment setup. We next evaluate running multiple concurrent applications. The baseline solution, “FRANCIS-individual”, simply launches multiple FRANCIS instances, one for each application. Each application is given a fixed amount of bandwidth and memory resources, and different applications do not share algorithm modules. Our alternative is “FRANCIS-combined”, where one instance runs all applications while sharing modules and resources across applications. To implement the combined version, we share the functionally identical instances of spanning tree and Shortest path tree for both clock synchronization and multicast (and remove the redundant instances from the clock synchronization). Other details can be found in Appendix H.1.

Experimental results. As Figure 8(a) shows, the reaction



(a) Header size. (b) Average latency (μ s). (c) Multicast overhead.

Figure 9: Evaluation on the spanner algorithm for multicast. Specifically, we vary the stretch s from 7 to 83.

time for all applications in the FRANCIS-combined is improved compared with FRANCIS-individual. For clock synchronization, its reaction time is improved by $1.35\times$ as the redundant algorithm components of clock synchronization are removed. For both multicast and Contra routing, FRANCIS-combined improvements arise from bandwidth sharing that improves the bandwidth utilization efficiency. For instance, if Contra routing is idle on some links, its bandwidth is allocated to multicast, speeding up the reaction.

As Figure 8(b) shows, because of Tofino’s [3] architecture, FRANCIS’s scheduler cannot adjust the memory usage of the clock synchronization. For other applications, algorithm scheduling (Section 4.3) enables some algorithm instances to execute in parallel to fit into different memory constraint settings. Thus, the reaction time of the clock synchronization application remains the same, while the reaction time of multicast and Contra increases as we shrink the allocated switch memory. Finally, we note that the gains are consistent for the whole memory range.

7.4 Flexibility in achieving various tradeoffs

In this section, we demonstrate how FRANCIS allows users to flexibly customize their distributed algorithms and FRANCIS’s resource usage, achieving different tradeoffs. We also present our evaluation for bandwidth allocation and reliably handling packet loss in Appendix H.5.

Tradeoffs between reaction time and performance.

Experiment setup. In our multicast use case, we vary the stretch parameter from 7 (the smallest allowed in [11]’s algorithm) to 83 and follow the previous setup parameters.

Experimental results. In Figure 9(a), increasing the stretch value decreases the header size by up to 38%. In Figure 9(b), larger stretch adds multicast latency in Jellyfish and Jupiter but slightly reduces the multicast latency in the FatTree. To explain this, the multicast latency is determined by both the packet size and the path length between the root of the multicast tree (source) and the leaves (destinations). Empirically, sparser graphs with higher stretch tend to result in longer path length but lower packet size. Finally, in Figure 9(c), the stretch has little impact on the total network bandwidth consumption. Accordingly, users can set a suitable stretch parameter to obtain the best tradeoffs.

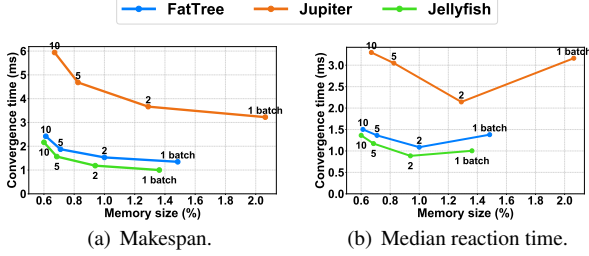


Figure 10: Evaluation for memory limitation in multicast use case. Here our scheduling policy is to optimize Makespan.




Figure 11: Time to performance with limited memory. MCT denotes multicast tree construction.

Tradeoffs between reaction time and memory.

Experiment setup. As detailed in Section 4.2, FRANCIS’s execution scheduler flexibly adapts to different memory constraints via batch-by-batch execution, which schedules more but smaller batches for lower memory size. In our multicast use case, we adjust the memory size so that the resulting number of batches ranges from 1 to 10.

Experimental results. Figure 10 shows that FRANCIS achieves flexible memory adjustment. As expected, the Makespan (the time at which the last algorithm terminates) increases monotonically when shrinking memory sizes. In the Jupiter topology, the reaction time is nearly doubled as we reduce the memory from 2.06% to 0.67%. However, for the median reaction time over failed multicast groups, we observe that the value is minimized with 2 batches. This is because we optimize Makespan, which does not take into account that batches that are scheduled early will be recovered sooner. We leave it as future work to optimize for different metrics.

Indeed, such a batch execution progressive recovery schedule can be observed in other performance metrics, as demonstrated in Figure 11. These results show how the fast recovery algorithm restores multicast service with suboptimal performance and that with batch execution, the performance metrics are further optimized progressively.

8 Related work

Clock synchronization. The last decade has witnessed research efforts such as DTP [42], Huygens [27], and DPTP [37] optimizing against the synchronization precision. These works, however, do not take into account network failures which, if not handled timely, lead to magnitudes higher time

uncertainty bound. While Precision Time Protocol (PTP) [6] and Sundial [45] design heuristic algorithms to handle failures (which FRANCIS outperforms), FRANCIS is not tied to any specific algorithm but enables incorporation of more advanced protocols. Recently, Nafaji and Wei proposed an orthogonal technique that minimizes the local drift rate [53]. Indeed, such a method reduces \max_drift_rate (see Eq. 1) and improves the uncertainty bound for FRANCIS as well.

Source routed multicast. Source-routed multicast achieves good scalability for multicast without requiring storing multicast rules inside switches’ SRAM. Elmo [61] and Orca [22] encode multicast trees as a set of multicast rules in each packet at the sender. Both incur considerable traffic overhead and do not support real-time failure recovery. IP multicast [21] is often even slower, as its detection is reported to require up to 3 minutes to recover from a failure.

Routing. Liu et al. [47] propose to let user packets explore new routes, but this explore-with-detour increases the latency of user traffic. F10 [48] designs heuristics to recover from a network event, but their algorithm is specific to the ab-FatTree topology. Contra [33] proposes an asynchronous protocol to support different routing policies but is less efficient in maintaining low bandwidth consumption than FRANCIS.

Fault tolerant switch state applications. The recently introduced RedPlane [39] makes running applications that rely on switch state robust to network failures. Their solution allows replicating switch state to other switches, thereby ensuring functionality in case of a failure.

9 Discussion

The message-passing algorithms we employ assume a static network topology, but we utilize them in a *dynamic* network where changes could happen any time. Our solution (Section 3) is to reinvoke reactions upon every event. An alternative is *dynamic* message-passing algorithms [10, 17] that leverage previous solutions to derive a new one on a topology change, which may reduce round complexity. However, these algorithms are often more complex and may require more switch resources, and we found our solution sufficiently fast.

Possible future work includes defining a language of FRANCIS and implementing a compiler to automatically translate user code into switch implementation. Some components of the language would be an execution graph that specifies the dependency of different algorithms, and an objective function, in terms of convergence time or other performance metrics. The compiler would compute an execution plan to satisfy all constraints and minimize the objective function.

We remark that after completing FRANCIS’s reaction, users may optionally launch the controller to perform some management tasks, such as running a centralized algorithm to obtain a fully optimal solution, or offloading the switches’ states back to end hosts to free up switch memory.

References

- [1] High-capacity strataxgs trident4 ethernet switch series. <https://www.broadcom.com/products/ethernet-connectivity/switching/strataxgs/bcm56880-series>.
- [2] Hp and mellanox benchmarking report for ultra low latency 10 and 40gb/s ethernet interconnect. https://www.mellanox.com/related-docs/whitepapers/HP_Mellanox_FSI%20Benchmarking%20Report%20for%2010%20%26%2040GbE.pdf.
- [3] Intel tofino programmable ethernet switch ASIC. <https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-series.html>.
- [4] The linux ptp project. <https://linuxptp.sourceforge.net/>.
- [5] Ieee standard for ethernet - amendment 10: Media access control parameters, physical layers, and management parameters for 200 gb/s and 400 gb/s operation. *IEEE Std 802.3bs-2017 (Amendment to IEEE 802.3-2015 as amended by IEEE's 802.3bw-2015, 802.3by-2016, 802.3bq-2016, 802.3bp-2016, 802.3br-2016, 802.3bn-2016, 802.3bz-2016, 802.3bu-2016, 802.3bv-2017, and IEEE 802.3-2015/Cor1-2017)*, pages 1–372, 2017.
- [6] Ieee standard for a precision clock synchronization protocol for networked measurement and control systems. *IEEE Std 1588-2019 (Revision of IEEE Std 1588-2008)*, pages 1–499, 2020.
- [7] U. Agarwal and V. Ramachandran. Faster deterministic all pairs shortest paths in congest model. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*, pages 11–21, 2020.
- [8] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986.
- [9] B. Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985.
- [10] S. Baswana, S. Khurana, and S. Sarkar. Fully dynamic randomized algorithms for graph spanners. *ACM Transactions on Algorithms (TALG)*, 8(4):1–51, 2012.
- [11] S. Baswana and S. Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, July 2007.
- [12] R. Ben-Basat, G. Even, K.-i. Kawarabayashi, and G. Schwartzman. Optimal distributed covering algorithms. *Distributed Computing*, pages 1–11, 2021.
- [13] R. Ben Basat, S. Ramanathan, Y. Li, G. Antichi, M. Yu, and M. Mitzenmacher. Pint: probabilistic in-band network telemetry. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 662–680, 2020.
- [14] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, et al. Onos: towards an open, distributed sdn os. In *Proceedings of the third workshop on Hot topics in software defined networking*, pages 1–6, 2014.
- [15] O. Borůvka. O jistém problému minimálním. 1926.
- [16] M. Budiú and C. Dodd. The p416 programming language. *SIGOPS Oper. Syst. Rev.*, 51(1):5–14, Sept. 2017.
- [17] K. Censor-Hillel, N. Dafni, V. I. Kolobov, A. Paz, and G. Schwartzman. Fast Deterministic Algorithms for Highly-Dynamic Networks. In Q. Bramas, R. Oshman, and P. Romano, editors, *24th International Conference on Principles of Distributed Systems (OPODIS 2020)*, volume 184 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:16, Dagstuhl, Germany, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- [18] S. Chechik and D. Mukhtar. Single-source shortest paths in the congest model with improved bound. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 464–473, 2020.
- [19] M. Chiesa, R. Sedar, G. Antichi, M. Borokhovich, A. Kaminski, G. Nikolaidis, and S. Schmid. Fast reroute on programmable switches. *IEEE/ACM Trans. Netw.*, 29(2):637–650, 2021.
- [20] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, et al. Spanner: Google's globally distributed database. *ACM Transactions on Computer Systems (TOCS)*, 31(3):1–22, 2013.
- [21] D. S. E. Deering. Host extensions for IP multicasting. RFC 1112, Aug. 1989.
- [22] K. Diab, P. Yassini, and M. Hefeeda. Orca: Server-assisted multicast for datacenter networks. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 1075–1091, Renton, WA, Apr. 2022. USENIX Association.
- [23] M. Elkin. A simple deterministic distributed MST algorithm with near-optimal time and message complexities. *J. ACM*, 67(2), Apr. 2020.
- [24] A. D. Ferguson, S. Gribble, C.-Y. Hong, C. Killian, W. Mohsin, H. Muehe, J. Ong, L. Poutievski, A. Singh, L. Vicisano, R. Alimi, S. S. Chen, M. Conley, S. Mandal, K. Nagaraj, K. N. Bollineni, A. Sabaa, S. Zhang, M. Zhu, and A. Vahdat. Orion: Google's Software-Defined networking control plane. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 83–98. USENIX Association, Apr. 2021.
- [25] O. Fischer, T. Gonen, F. Kuhn, and R. Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*, pages 153–162, 2018.
- [26] B. Fitzpatrick. Distributed caching with memcached. *Linux journal*, 124, 2004.
- [27] Y. Geng, S. Liu, Z. Yin, A. Naik, B. Prabhakar, M. Rosenblum, and A. Vahdat. Exploiting a natural network effect for scalable, fine-grained clock synchronization. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 81–94, 2018.

- [28] P. Gill, N. Jain, and N. Nagappan. Understanding network failures in data centers: measurement, analysis, and implications. In *Proceedings of the ACM SIGCOMM 2011 Conference*, pages 350–361, 2011.
- [29] D. Harrington and D. Levi. Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol. RFC 4318, Dec. 2005.
- [30] C. L. Hedrick. Rfc1058: Routing information protocol, 1988.
- [31] T. Holterbach, E. C. Molero, M. Apostolaki, A. Dainotti, S. Vissicchio, and L. Vanbever. Blink: Fast connectivity recovery entirely in the data plane. In J. R. Lorch and M. Yu, editors, *16th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2019, Boston, MA, February 26-28, 2019*, pages 161–176. USENIX Association, 2019.
- [32] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. Alimi, C. Bhagat, S. Jain, J. Kaimal, S. Liang, K. Mendeleev, et al. B4 and after: managing hierarchy, partitioning, and asymmetry for availability and scale in google’s software-defined wan. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 74–87, 2018.
- [33] K.-F. Hsu, R. Beckett, A. Chen, J. Rexford, P. Tamma, and D. Walker. Contra: A programmable system for performance-aware routing. In *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020*, 2020.
- [34] IEEE Standards Association. IEEE Standard for Local and Metropolitan Area Networks Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management., 2007.
- [35] T. Instruments. DP83867IR/CR Robust, High Immunity 10/100/1000 Ethernet Physical Layer Transceiver. <https://www.ti.com/lit/ds/symlink/dp83867ir.pdf>, 2015.
- [36] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al. B4: Experience with a globally-deployed software defined wan. *ACM SIGCOMM Computer Communication Review*, 43(4):3–14, 2013.
- [37] P. G. Kannan, R. Joshi, and M. C. Chan. Precise time-synchronization in the data-plane using programmable switching asics. In *Proceedings of the 2019 ACM Symposium on SDN Research*, pages 8–20, 2019.
- [38] D. Katz and D. Ward. Bidirectional Forwarding Detection, 2010. RFC 5880 (Proposed Standard).
- [39] D. Kim, J. Nelson, D. R. K. Ports, V. Sekar, and S. Seshan. Redplane: Enabling fault-tolerant stateful in-switch applications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM ’21*, page 223–244, New York, NY, USA, 2021. Association for Computing Machinery.
- [40] J. Korhonen, J. Rybicki, et al. Deterministic subgraph detection in broadcast congest. In *21st International Conference on Principles of Distributed Systems (OPODIS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2018.
- [41] A. Krentsel, N. Saran, B. Koley, S. Mandal, A. Narayanan, S. Ratnasamy, A. Al-Shabibi, A. Shaikh, R. Shakir, A. Singla, and H. Weatherspoon. A decentralized sdn architecture for the wan. In *Proceedings of the ACM SIGCOMM 2024 Conference, ACM SIGCOMM ’24*, page 938–953, New York, NY, USA, 2024. Association for Computing Machinery.
- [42] K. S. Lee, H. Wang, V. Shrivastav, and H. Weatherspoon. Globally synchronized time via datacenter networks. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 454–467, 2016.
- [43] W. Li, J. Zhang, G. Zeng, Y. Liu, Z. Wang, C. Zeng, P. Zhou, Q. Wang, and K. Chen. Gleam: An rdma-accelerated multicast protocol for datacenter networks. *arXiv preprint arXiv:2307.14074*, 2023.
- [44] X. Li and M. J. Freedman. Scaling ip multicast on datacenter topologies. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT ’13*, page 61–72, New York, NY, USA, 2013. Association for Computing Machinery.
- [45] Y. Li, G. Kumar, H. Hariharan, H. Wassel, P. Hochschild, D. Platt, S. Sabato, M. Yu, N. Dukkipati, P. Chandra, and A. Vahdat. Sundial: Fault-tolerant clock synchronization for datacenters. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 1171–1186. USENIX Association, Nov. 2020.
- [46] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [47] J. Liu, A. Panda, A. Singla, B. Godfrey, M. Schapira, and S. Shenker. Ensuring connectivity via data plane mechanisms. In *NSDI*, pages 113–126, 2013.
- [48] V. Liu, D. Halperin, A. Krishnamurthy, and T. Anderson. F10: A fault-tolerant engineered network. In *Presented as part of the 10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*, pages 399–412, 2013.
- [49] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986.
- [50] J. Martin, J. Burbank, W. Kasch, and P. D. L. Mills. Network Time Protocol Version 4: Protocol and Algorithms Specification. RFC 5905, June 2010.
- [51] J. Meza, T. Xu, K. Veeraraghavan, and O. Mutlu. A large scale study of data center network reliability. In *Proceedings of the Internet Measurement Conference 2018, IMC ’18*, page 393–407, New York, NY, USA, 2018. Association for Computing Machinery.
- [52] J. Moy. OSPF Version 2. RFC 2328, Apr. 1998.
- [53] A. Najafi and M. Wei. Graham: Synchronizing clocks by leveraging local clock properties. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 453–466, Renton, WA, Apr. 2022. USENIX Association.
- [54] R. Neugebauer, G. Antichi, J. F. Zazo, Y. Audzevich, S. López-Buedo, and A. W. Moore. Understanding pcie performance for end host networking. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM ’18*, page 327–341, New York, NY, USA, 2018. Association for Computing Machinery.
- [55] J. Ousterhout, P. Agrawal, D. Erickson, C. Kozyrakis, J. Leverich, D. Mazières, S. Mitra, A. Narayanan, G. Parulkar, M. Rosenblum, et al. The case for ramclouds: scalable high-performance storage entirely in dram. *ACM SIGOPS Operating Systems Review*, 43(4):92–105, 2010.

- [56] J. Ousterhout, A. Gopalan, A. Gupta, A. Kejriwal, C. Lee, B. Montazeri, D. Ongaro, S. J. Park, H. Qin, M. Rosenblum, et al. The ramcloud storage system. *ACM Transactions on Computer Systems (TOCS)*, 33(3):1–55, 2015.
- [57] G. Pandurangan, P. Robinson, and M. Scquizzato. A time- and message-optimal distributed algorithm for minimum spanning trees. *ACM Transactions on Algorithms (TALG)*, 16(1):1–27, 2019.
- [58] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)*, 27(2):228–234, 1980.
- [59] D. Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
- [60] F. Petrini and M. Vanneschi. k-ary n-trees: high performance networks for massively parallel architectures. In *Proceedings 11th International Parallel Processing Symposium*, pages 87–93, 1997.
- [61] M. Shahbaz, L. Suresh, J. Rexford, N. Feamster, O. Rottenstreich, and M. Hira. Elmo: Source routed multicast for public clouds. *IEEE/ACM Trans. Netw.*, 28(6):2587–2600, Dec. 2020.
- [62] A. Shamis, M. Renzelmann, S. Novakovic, G. Chatzopoulos, A. Dragojević, D. Narayanan, and M. Castro. Fast general distributed transactions with opacity. In *Proceedings of the 2019 International Conference on Management of Data*, pages 433–448, 2019.
- [63] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart, and A. Vahdat. Jupiter rising: A decade of clos topologies and centralized control in google’s datacenter network. *SIGCOMM Comput. Commun. Rev.*, 45(4):183–197, Aug. 2015.
- [64] A. Singla, C.-Y. Hong, L. Popa, and P. B. Godfrey. Jellyfish: Networking data centers randomly. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pages 225–238, San Jose, CA, Apr. 2012. USENIX Association.
- [65] M. Stenberg and S. Barth. Distributed Node Consensus Protocol. RFC 7787, Apr. 2016.
- [66] M. Tirmazi, R. Ben Basat, J. Gao, and M. Yu. Cheetah: Accelerating database queries with switch pruning. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data, SIGMOD ’20*, page 2407–2422, New York, NY, USA, 2020. Association for Computing Machinery.
- [67] L. Yu, J. Sonchack, and V. Liu. {OrbWeaver}: Using {IDLE} cycles in programmable networks for opportunistic coordination. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 1195–1212, 2022.
- [68] C. Zeng, L. Luo, T. Zhang, Z. Wang, L. Li, W. Han, N. Chen, L. Wan, L. Liu, Z. Ding, et al. Tiara: A scalable and efficient hardware acceleration architecture for stateful layer-4 load balancing. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 1345–1358, 2022.

A Theoretical problems and algorithms

The CONGEST model is one of the main models studied in the distributed graph algorithms community, and most practical message-passing algorithms follow this model [46, 59]. The network is modeled as a graph, $G(V, E)$, where nodes are computational units and edges are communication links. We denote the number of vertices by $n = |V|$, the neighbors of v in G by $N(v)$, and the diameter of G by D . We assume that every node in the network has a unique $O(\log n)$ bit identifier. In the CONGEST model, communication takes place in synchronous rounds, where at the beginning of the round a node receives all messages sent by its neighbors. Next, the node performs some computation and updates its local state. At the end of the round, the node sends (potentially different) messages to all of its neighbors, which are delivered at the beginning of the next round. The size of each message is bounded by $O(\log n)$ bits. One performance metric used in this model is the *round complexity* of the algorithm, which is the number of synchronous communication rounds required for completing a task. Another metric is the *message complexity* of the algorithm – the number of total messages sent during the execution. The focus on communication cost over local computation is meant to model the fact that the main bottleneck in the performance of real-world distributed systems is often the communication cost.

Alternatively, **asynchronous event-based algorithms** provide another framework for distributed algorithms that has been widely studied. Unlike the CONGEST model, which executes on a per-round basis, the asynchronous event-based algorithms proceed on a per-packet basis: for each incoming packet, the switch may decide to update its stateful memory and multicast messages to some neighbors. As discussed in Section 2.3, since asynchronous algorithms can typically be converted to synchronous CONGEST algorithms and they are typically more bandwidth-consuming than CONGEST algorithms, CONGEST algorithms are a better design choice for the distributed optimization phase. However, for the fast recovery phase, asynchronous algorithms can be a better choice because they enable immediate reactions to incoming packets without waiting for all packets in a round, which can reduce the convergence time.

Graph problems. While we have presented three use cases to show the effectiveness of our framework in improving the performance of networking applications, we also list several other theoretical problems and algorithms (e.g., maximal independent set, minimum vertex cover, etc.) that are potentially useful for networking. The relevant theory problems are given in Table 7.

Problem name	Description	Complexity	Deterministic	Example application	Reference
$(2k-1)$ -Spanner	An edge set $S \subseteq E$ of size $ S = O(kn^{1+1/k})$ such that for all u, v : $d_{(V,S)}(u, v) \leq (2k-1)d_{(V,E)}(u, v)$.	$O(k^2)$	N	Source-routing multicast	Baswana and Sen [11]
Shallowest spanning tree 2-approx (shortest path tree).	A rooted spanning tree of minimum depth.	$O(D)$	Y	Clock synchronizaiton	Section 5.1
Leader election	Elect one vertex as "leader".	$O(D)$	Y	Clock synchronization	Folklore
Min spanning tree	A tree with minimum weight.	$O(n \log n)$	Y	Loop prevention	Boruvka [15]
Maximal Independent Set	A set of vertices without edges between them such that any other vertex has neighbor in the set.	$O(\log n)$	N	Concurrent wireless transmission	Luby [49], Alon et. al [8]
Min vertex cover $(2+\epsilon)$ -approx.	A minimal weight vertex set such that each edge has an endpoint in the set.	$O\left(\frac{\log \Delta}{\log \log \Delta}\right)$	Y	Assign edge-monitoring responsibilities	Ben-Basat et al. [12]

Table 7: A summary of the theory problems either applied to our use cases or potentially being useful for other networking problems.

Algorithm 1 Workflow of a message-passing algorithm module. Users only need to implement the highlighted parts, and the blue-color code is function calls of our algorithmic primitives.

```

1: procedure SYNCHRONOUS CONGEST ALGORITHMS
2: Configs:
3:   Synchronizer  $\in \{\alpha, \beta, \dots\}$ 
4:   Priority: Integer
5:   Round_number (rcnt): Integer expression
6:   Bandwidth and memory for this module.
7:   Parallel instance IDs  $[l, r]$ :  $l, r \in$  Integer expression
8:   Algorithm_parameters: Numbers and functions
9: for parallel instance_id  $\in [l, r]$  do
10:  Option 1:       $\triangleright$  algorithmic primitives provided by FRANCIS
11:    Primitive_calls(params)
12:  Option 2:       $\triangleright$  Streaming Computation
13:    for round_id  $\in \{0, \dots, rcnt - 1\}$  do
14:      Synchronization()
15:      Upon receiving a packet  $p$ :
16:        Stream-processing( $p$ )
17:      Until Completion() of the current round
18:        if Termination() then
19:          Term_notify_and_exit
20:        else Send_to_neighbors()
21:  Option 3:       $\triangleright$  Computation at the end of the round
22:    for round_id  $\in \{0 \dots rcnt-1\}$  do
23:      Synchronization()
24:      Receive and stash packets from neighbors
25:      Until Completion() of the round:
26:        Computation(pkts)
27:      if Termination() then
28:        Term_notify_and_exit
29:      else Send_to_neighbors()

```

B Pseudo code for message-passing algorithm specification

The distributed algorithm modules. Implementation of a distributed algorithm module begins with a wide range of user configurations. As Algorithm 1 show, the configurations consist of 6 components. Synchronizer specifies which synchronizer to choose for the CONGEST algorithm. Different synchronizers achieve different performance tradeoffs and we elaborate them in Section 4.1. Priority is an integer to determine the priority of allocating resources to algorithm instances in this module and scheduling their execution. Round_number refers to the number of rounds in this CONGEST algorithm module. Bandwidth and memory for this module impose further resource limitations to this particular module. A single algorithm module may correspond to multiple algorithm instances, and users may manually decide to spread out the execution of these instances by running only a batch of instances at a time and another batch later on. Users thus may optionally specify Parallel instance ID, which is the algorithm instance IDs to run in parallel in this module. Finally, users can also add other Algorithm_parameters flexibly to further configure their algorithms.

Then comes the body of distributed algorithms, which consists of instances of algorithms to be run in parallel, following the user implementation. To implement, one option is to directly call the primitives in FRANCIS's library. Table 2 lists some of our algorithmic primitives as well as their potential applications in this paper. We note that users can supply their own parameters and even their own functions to be executed in addition to what the primitive executes. In Section 5, the use cases provide more contexts and explanations for these algorithmic primitives. Another option is to implement custom user algorithms. For CONGEST algorithm modules, users can choose the streaming computation scheme, corresponding to option 2. In this case, FRANCIS aggregates incoming packets until Completion() of the round is met, when the module may terminate or continue to send_to_neighbors next round, as implemented by users. Alternatively, users may opt for doing computation at the end of the round, where users

implement their own computation functions.

C Pseudocode of some algorithmic primitives and use cases

Algorithm 2 Set cover algorithm on a shortest path tree.

```

1: procedure SET COVER
2: Round  $3k$  for layer  $D - k$ :  $\triangleright D$ : depth of the SPT
3:   Send msg() to all the neighbors on layer  $l$ 
4: Round  $3k + 1$  for layer  $D - k - 1$ :
5:   self.count = # messages received
6:   Multicast msg(.count = self.count) back to neighbors
   that sent a packet at round  $3k$ 
7: Round  $3k + 2$  for layer  $D - k$ :
8:   self.parent = arg maxnb msg[nb].count
9:   Send msg(.father = "yes") to self.parent

```

Algorithm 3 Conditional broadcast and aggregation, and its application to leader election.

```

1: procedure CONDITIONAL BROADCAST AND AGGREGATION
2:    $\triangleright$  Each switch initially owns a value self.value
3:    $\triangleright$  The goal is to aggregate the value and
4:    $\triangleright$  broadcast them across the network.
5: Stream-processing(p):  $\triangleright$  Line 16 in Algorithm 1
6:   self.value = arithmetic(self.value, p.value)
7: Completion:  $\triangleright$  Line 17 in Algorithm 1
8:   The current switch has received one message per each
   neighbor
9: Termination:  $\triangleright$  Line 18 in Algorithm 1
10:   A total of  $D$  rounds have passed, where  $D$  is the esti-
   mated diameter of the graph.
11: Send_to_neighbors:  $\triangleright$  Line 20 in Algorithm 1
12:   if A condition() is met. then
13:     Broadcast self.value to all neighbors.
14: procedure LEADER ELECTION
15: arithmetic(): min(self.value, p.value)
16: condition(): always true.

```

Algorithm 4 Reaction algorithms for clock synchronization.

```

1: procedure FAST RECOVERY
2: Use the Asynchronous global flooding primitive
3: procedure DISTRIBUTED OPTIMIZATION
4: First round:
5:   Randomly sample  $r$  switches as candidate roots.
6: Flooding(record_candidate_roots)
7: for each candidate root do
8:   Use the Shortest path tree primitive
9:   Use the Bottom up tree aggregation primitive
10: Move the root to the center of the tree.
11: Pick the tree with the minimal depth among these.

```

Algorithm 5 The modified Contra algorithm

```

1: procedure SYNCHRONOUS SHORTEST PATH (CONTRA)
2: At each round:
3:   Receive and stash packets from all neighbors in PG.
4:   Until all the updates have been received:
5:     Merge packets corresponding to the same PG
6:     Update the routing table.
7:   Send:
8:     for Merged packet that contributes to an update do
9:       Find the downstream states and multicast.

```

D The Choice of Synchronizers

The CONGEST model assumes a synchronous network, but real-world networks are usually asynchronous. Various theoretical synchronization algorithms have been proposed, such as the α -synchronizer and β -synchronizer [9]. The α -synchronizer is essentially a local synchronizer that only ensures that a node does not begin round $t + 1$ until its neighbors have completed round t , while β -synchronizer is a global synchronizer that guarantees full synchronization, so round $t + 1$ does not begin until all nodes have completed round t . As the CONGEST model operates locally, a local synchronizer suffices. However, global synchrony may be needed for scenarios such as launching a CONGEST algorithm after an asynchronous algorithm. We have implemented both α and β synchronizers and leave the implementations of other synchronizers (such as γ synchronizer [9]) as future works.

We now briefly present the message-passing process of the α and the β synchronizer. In the α -synchronizer, every message is tagged with the round number. For each round, every node sends the message to each neighbor (we send a "no-msg" if there is no message-passing to a neighbor). When a node receives messages from all its neighbors, it performs local computation and proceeds to the next round. In the β -synchronizer, we construct a spanning tree before running the CONGEST algorithm. Before the execution of round t , the root of the spanning tree must confirm that all nodes complete their execution of round $t - 1$. This is achieved by sending notification messages from every node to the root in a bottom-

Algorithmic primitives	Category	Description	Applications
Asynchronous global flooding	Asynchronous	Use flooding to 1) broadcast a short message to all switches 2) construct a flooding spanning tree, 3) initializing the synchronizer throughout the network.	Fast recovery for clock sync and multicast; bootstrapping for Contra routing.
Bottom-up tree aggregation	Asynchronous	Aggregate values from the leaves of a spanning tree to the root.	Distributed optimization for clock sync, fast recovery for multicast.
Conditional broadcast and aggregation (as detailed in Algorithm 3)	CONGEST	Maintain a single value v for each switch. Aggregate v with those received from neighbors for each round. Then broadcast the updated value next round provided that the condition function is met.	Leader election for clock sync, Spanner for multicast.
Shortest path tree	CONGEST	Construct a spanning tree layer-by-layer in a BFS fashion. Each round the current leaf switches look for unvisited switches to be added as a new layer of leaf switches.	Distributed optimization for clock sync, multicast tree for multicast.
Set cover on a shortest path tree (as detailed in Appendix E.2)	CONGEST	In the shortest path tree, greedily select a minimal number of switches in layer i to cover layer $i + 1$, <i>i.e.</i> , every switch in layer $i + 1$ has a selected neighboring switch in layer i .	Multicast tree for multicast.
Synchronous shortest path	CONGEST	Synchronously compute shortest path routes w.r.t. a routing policy.	Contra routing.

Table 8: Descriptions of some FRANCIS’s algorithmic primitives.

up manner. Similarly, to launch round t , the root replies to every node through the tree that round t has begun in a top-down manner. One can think of the α - and β -synchronizers as two extremes, where the first sends more messages ($O(n^2)$ on each link) but the messages only travel for one hop, while the latter is more message-efficient ($O(n)$ messages through the tree) but the latency is potentially longer because messages travel for multiple hops.

E Hardness of min-size shortest path tree

In this appendix, we show that finding the minimal shortest path tree is at least as hard as set cover, thus implying its NP-hardness.

Formally, the set cover problem is defined as follows: Let \mathcal{S} be a set of subsets of the universe $U = \{1, \dots, n\}$ (*i.e.*, each $S \in \mathcal{S}$ satisfies $S \subseteq U$). The set cover problem asks what is the smallest set $\mathcal{C} \subseteq \mathcal{S}$ such that $\bigcup_{S \in \mathcal{C}} S = U$.

Recall that our min-cast shortest path tree problem is defined as: Given a graph $G = (V, E)$, a sender $s \in V$ and receivers $D \subseteq V$, find a shortest path tree T (*i.e.*, $dist_G(s, d) = dist_T(s, d)$) of minimal size. Here, we show that an algorithm to the min-cast shortest path tree problem implies an algorithm for set-cover, and therefore the above hardness result holds for our problem as well.

Given a set cover instance (U, \mathcal{S}) , construct the following graph:

$$G = (\{s\} \cup \mathcal{S} \cup U, \{\{s, S\} \mid S \in \mathcal{S}\} \cup \{\{S, u\} \mid u \in S, S \in \mathcal{S}\}).$$

Further, let s be the source and $D = U$ be the destinations. That is, in the graph, the source s is connected to all sets $S \in \mathcal{S}$, and each element $u \in U$ is connected to the sets that it is a member of.

We have that the optimal solution must include s and all elements $u \in U$, in addition to some cover $\mathcal{C} \subseteq \mathcal{S}$. Thus, the solution’s size is $n + 1 + |\mathcal{C}|$, and finding the optimal shortest path tree reveals the optimal cover to the underlying problem.

F Multicast Algorithms

F.1 Constructing Fast Recovery Multicast Trees on top of a Spanning Tree

We mentioned in Section 5.2 that we need to construct fast recovery multicast trees that are subtrees of T_0 , the spanning tree we constructed earlier. The fast recovery multicast tree should connect together the source and the destinations. Instead of directly constructing the trees independently, we propose to build the trees altogether in a bottom up manner, thanks to our message packing technique. We define the bitmap $is_son[child][j]$ as whether the multicast tree of group j contains an edge linking itself to the child. We also define a bitmap in_tree , the j ’th entry of which denotes whether the multicast tree of group j contains the current switch. Apparently, $in_tree[j]$ indicates if a switch is in a multicast tree of group j or not, and $is_son[\cdot][j]$ indicates the ports corresponding to the set of children in the tree. Both are boolean arrays that can be compressed into a single bitmap. We have the following equation:

$$in_tree(v) = \bigvee_{child} is_son(v)[child]$$

$$is_son(parent(v))[v] = in_tree(v)$$

, where v is the current switch. Using this equation, we carry the in_tree bitmap (with < 2000 -bit) into a single < 2000 -bit packet, compute is_son and in_tree bitmaps with bitwise operations and send the bitmaps in a bottom-up manner from the leaf switches to the root switch. Apparently, the message packing reduces the bandwidth consumption and thus the convergence time, compared with the approach where we dedicate one upstream packet per multicast group, resulting in thousands of 64-byte Ethernet packets to be sent per hop.

F.2 Efficient Set Cover for Multicast Tree Construction

Once we have executed the shortest path tree algorithm for finding shortest path length between the sender (source) switch and all the switches in the network, we start to build

the multicast tree from leaf switches to the source switch from the leaf layer to the source progressively. We would like to minimize the number of nodes in the multicast tree so as to minimize the multicast traffic. We have abstracted this algorithm as one of our algorithmic primitives, and we present its pseudocode in Algorithm 2.

Now we elaborate on the details of the algorithm. Initially, all the receivers are marked as active switches. For all the active switches in layer l we run the distributed set cover to find as few parent switches as possible to cover them. These parent switches are marked as active switches in the next layer, and we repeat the process. We propose a greedy approach for set cover within three rounds:

1. Nodes on layer l send one message to each neighbor on layer $l - 1$, and each node v on layer $l - 1$ counts the total number of received messages as $count[v]$.
2. Each node v on layer $l - 1$ multicasts $count[v]$ back to nodes that sent messages in the previous round in layer l , and each node in layer l picks up the maximum value of $count[v]$ and identifies the corresponding node on layer $l - 1$ as its parent.
3. Each Node v in layer l sends a message to its parent to confirm, and its parent add v as one of its children.

Specifically for our multicast application, once the trees has been constructed, our last step is to transmit the encoding rules (the bitmaps) into the source tor switch and store the rules on it. We continue to use the encoding information from the tor switch for multicast of the affected groups until the controller later fully installs the multicast rules on end hosts so that we switch back to the source-routed multicast again.

E.3 Installing Multicast Rules into Tor Switches

During the period between the completion of the distributed optimization phase (when we construct multicast trees) and the time when the controller computes and installs multicast rules on end hosts, we propose to temporarily cache *all* the source-routed multicast rules of the affected groups in the tor switches. When a multicast packet from an affected group reaches the tor switch, it appends the multicast rules stored on the switch. For later hops, switches look up the multicast rules on the header. We note that such a source-routed multicast solution guarantees a bounded memory overhead per switch, compared with the solution that we store the multicast rules separately on all the switches. This is because the number of multicast groups such that their sources are within a certain tor t is generally a bounded quantity.

G Implementation details

Our detection phase for clock synchronization. As a starting point, Sundial [45] detects failures by passively listening to the synchronization messages, so that if several consecutive




Figure 12: The 3-4-ary FatTree.

losses of synchronization packets occur, the switch concludes that some upstream links or switches have failed. However, such an approach cannot detect exactly which devices have failed. Therefore, in FRANCIS’s implementation, we add a technique called ping-on-timeout. In this approach, if a switch does not hear the synchronization messages for two timeouts, it actively pings its parent, and if the parent does not reply with an ACK message then the switch can conclude that the failure stems directly from its parent rather than from its parent’s ancestor.

H Supplemental experimental settings and results

H.1 Evaluation setups

Testbed experiment setups in Section 7.1. To configure FRANCIS’s reaction algorithm for clock synchronization, we reconstruct 4 candidate trees and select the best one in terms of the tree depth in the distributed optimization phase. The synchronization interval is set as $50\mu s$ for FRANCIS prototype and $61\mu s$ for PTP, the lowest possible interval we could achieve for PTP. To measure the time uncertainty bound, we follow Sundial [45]’s setup and configure $\epsilon_0 = 5ns$ and $max_drift_rate = 0.0002$ or equivalently 200ppm.

The 3-4-ary FatTree. We here define the 3-4-ary FatTree topology used in our testbed experiments. As depicted in Figure 12, the 3-4-ary FatTree is created from a standard 4-ary FatTree but with the last core switch and the last pod (containing 2 aggregation switches and 2 ToR switches) removed. The number of switches remaining is therefore $n = 15$. Links connected to the removed switches are also deleted accordingly.

Simulation setups. We configure Ethernet links to have a propagation delay randomly within $100 \pm 10ns$. The bandwidth of links is set to be 100Gbps; regardless of the bandwidth capacity of links, we always allocate a $< 1Gbps$ bandwidth for FRANCIS (as detailed in each application) identically for all links, which is far below the bandwidth capacity of modern Ethernet links. The packet loss rate for each link is set to be 0.1% independently at random for each packet, except in Section H.5 where we study how the packet loss rate affects the convergence time.

Experimental setups in Section 7.2.1. We generate one or

two switch failures for our experiments. With multiple failures, we deliberately set the failed switches to be in different domains so that Sundial switches back to the control-plane solution. We set the interval between two consecutive synchronization messages to be $50\mu\text{s}$ for both our solution and Sundial, and detect a timeout on 3 loss of synchronization messages. This corresponds to $< 1\text{Mbps}$ per-link bandwidth consumption (Figure 5(b)). To configure FRANCIS, we set the number of sampled candidate roots in the distributed optimization phase to be 4. We add the synchronous messaging proposed by Sundial [45] technique to PTP [6] to further reduce its ϵ value.

Experimental setups in Section 7.2.2. Following the setup from Elmo [61] and Orca [22], we generate 10^6 multicast groups; for each multicast group, we generate several VMs and randomly place them on end hosts. One of the VMs serves as the sender and the rest of them serve as the receivers. The number of VMs per group (*i.e.*, group size) follows the IBM WebSphere Virtual Enterprise (WVE) [44] distribution, and the average group size is set to be $0.05n$, where n is the number of switches in the network. To configure FRANCIS, we allocate 200Mbps bandwidth (Figure 16(b)) and 1.5% memory of the Tofino’s total switch memory, which is enough to enable failure recovery algorithms for all the failed groups to run in parallel. We also set $k = 13$ for the spanner algorithm [11]. For Elmo, we choose the redundancy value $R = 3$. Since Orca does not adapt to Jupiter and Jellyfish topologies and does not report its failure recovery time, we only evaluate its packet header size and total multicast traffic on FatTree.

Experiment setups of running concurrent applications in Section 7.3. We choose the FatTree as our topology, and inject one link failure for our experiments. The bandwidth for FRANCIS-individual is allocated to be 200Mbps, 200Mbps, 10Mbps for multicast, Contra, and clock synchronization respectively. Correspondingly we allocate the sum of FRANCIS-individual’s bandwidth to FRANCIS-combined. For Figure 8(a), we do not limit the memory usage, and the FRANCIS-combined with Contra consumes 3.75% memory. For Figure 8(b) we gradually shrink the memory size from 3.75% to 1.58% for FRANCIS-combined, and proportionally adjust the memory for each application in FRANCIS-individual.

H.2 Validating the simulator

In order to evaluate how accurately our simulator could model behaviors of a real-world distributed programmable-switch networks, we compare results conducted on the simulator against testbed results against testbed results. As shown in Figure 13, the testbed results regarding the convergence time and time uncertainty bound overall match with simulation results, which validates the utility of our simulation model. FRANCIS Tofino’s convergence time is less than 10% slower than that derived from simulation, primarily due to us conser-




Figure 13: Performance of FRANCIS’s Tofino prototype for clock synchronization comparing against simulation results. The dark, cross-patched bars represent the fast recovery phase while the light, dotted bars correspond to the distributed optimization phase.


vatively implementing the bandwidth limitation component on Tofino to meet with the hardware restrictions (Section 4.2). That is, the allocated bandwidth may not be fully utilized, and accordingly extra delays are incurred so that the convergence time becomes longer.

H.3 Additional experiments for comparisons with state-of-the-arts

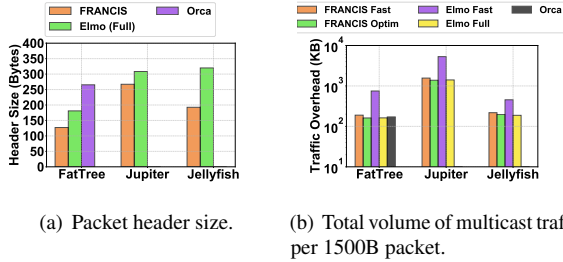
Clock synchronization. In this section, we provide more experimental results for clock synchronization. Figure 14 shows the convergence time and peak ϵ results with one switch failure. Recall that Sundial [45] handles one failure by directly switching to the pre-computed backup plan within each switch. Therefore, Sundial’s fast recovery plan achieves the best convergence time $150\mu\text{s}$, outperforms FRANCIS-100Mb’s $186.7\mu\text{s}$ fast recovery time and $359.5\mu\text{s}$ distributed optimization time for the FatTree topology. However, as shown in Figure 14(b) and Figure 5(a), both the peak ϵ and the back-to-normal ϵ value (ϵ after the fast recovery of Sundial or the distributed optimization phase of FRANCIS) of Sundial are worse than FRANCIS-100Mb. In fact, the suboptimality of Sundial’s backup tree is to be blamed for. For the peak ϵ , although the reaction time of FRANCIS is slightly faster than FRANCIS-100Mb, its resulting backup tree is 3 hops deeper than FRANCIS’s, and, according to Equation 1, the peak ϵ becomes larger. The same explanation also applies to the back-to-normal ϵ value. As previously argued, due to Sundial’s slow control-plane based full recovery, the suboptimal ϵ continues to have effect on the clock synchronization application for the duration of $\approx 100\text{ms}$ full recovery time. Finally, to compare with PTP, although it achieves a similar level of convergence time than FRANCIS-10Mb (which is several times longer than FRANCIS-100Mb), both the peak ϵ and the back-to-normal ϵ are worse than our solution.

Multicast. In Section 7.2.2, we take the FatTree as an example to show the results of packet header size and the total volume of multicast traffic. Figure 15 illustrates all of our


results in the three topologies. For packet header size, we can derive a straightforward conclusion that FRANCIS’s Spanner algorithm reduces the source-routed packet header size by a significant amount. For evaluating the multicast traffic overhead, we measure how much traffic it generates in the whole topology to forward a 1500B multicast packet. As described in Section 7.2.2, the results clearly show how FRANCIS’s full recovery achieves comparable traffic overhead than Elmo’s final solution and Orca, and Elmo’s fast recovery attempt results in an overwhelmingly high traffic overhead. After all, Elmo’s proposes to switch to the highly inefficient unicast as its fast recovery approach. Finally, we provide a decomposition of the convergence time of each FRANCIS’s recovery algorithm in Figure 16(a). We also illustrate FRANCIS’s time-to-message-overhead results in Figure 16(b) to provide more intuitions of FRANCIS’s multicast reaction algorithms.



(a) Convergence time. (b) Peak ϵ .
Figure 14: Performance for clock synchronization compared with Sundial [45] and PTP [6] with one switch failure.



(a) Packet header size. (b) Total volume of multicast traffic per 1500B packet.
Figure 15: Evaluation on the Spanner algorithm for FRANCIS’s multicast. Here "Optim" refers to results after the distributed optimization phase, and "Full" refers to full recovery for Elmo.



(a) Anatomy of failure recovery time (b) Time to message overhead
Figure 16: The recovery time and message overhead for multicast. SPT refers to the shortest path tree algorithm. Note that we build multicast trees by first running the shortest path tree and then the set cover algorithm.

H.4 Deploying F10 into FRANCIS

To demonstrate more use cases, we have also implemented F10 [48]’s re-routing protocol into FRANCIS. F10 proposes a local-rerouting protocol as a fast recovery protocol to re-establish an ad-hoc route in place. When the local rerouting protocol fails to find a new route, it falls back to the push-back flow redirection protocol, which is essentially an asynchronous algorithm to redirect the affected upstream flows to new routes. Both protocols are designed specifically for their AB FatTree topology. We deploy their pushback flow redirection protocol to FRANCIS, and study on their local-rerouting protocol is out-of-scope of our paper.

Following the experimental setups elaborated in Section 7.3, we evaluate the combined version of FRANCIS with our clock synchronization and multicast use cases. Again we compare the combined version against the solution that runs each application individually. The bandwidth for FRANCIS-individual is 200Mbps, 10Mbps, 10Mbps for clock synchronization, multicast, and F10 respectively.

Figure 17 shows the results comparing both versions. One direct observation is that FRANCIS-combined achieves faster convergence for all applications than FRANCIS-individual. This is because our bandwidth sharing technique utilizes bandwidth resources of idle applications to the full, and we merge the same algorithm components from different applications together to avoid duplicated calculations. Another finding is that, regarding Figure 8(a), FRANCIS-combined with Contra achieves higher improvement for multicast application in convergence time than FRANCIS-combined with F10. Specifically, FRANCIS-combined with Contra reduces the convergence time of multicast to 1.187ms, while FRANCIS-combined with F10 only reduces it to 1.298ms. We note that Contra is higher in bandwidth complexity than F10. Therefore, the results suggest that combining multiple bandwidth-intensive protocols together achieves a better bandwidth utilization rate than combining one bandwidth-intensive protocol with non-intensive protocols and thus exhibits a faster convergence time.





Figure 17: We integrate F10 [48] into FRANCIS as the routing protocol and combine it with reaction algorithms of both clock synchronization and multicast. The figure compares the combined FRANCIS with FRANCIS-individual to run each application individually.



(a) Clock Synchronization

(b) Multicast

Figure 18: The convergence time with respect to different bandwidths for both clock synchronization and multicast.

H.5 Evaluations on FRANCIS’s flexibility and reliability


Number of trees	1	2	3	4	5
Convergence time (ms)	1.687	1.743	1.805	1.859	1.973
Expected ϵ (ns)	26.054	25.131	25.018	25.003	25.0005

Table 9: Convergence time and ϵ varying the number of clock synchronization trees to be constructed in the FatTree topology. The bandwidth is limited as 10Mbps.

For the clock synchronization application, one critical parameter is the number of shortest path trees (SPTs) to be constructed. Constructing more SPTs takes longer convergence time but allows the leader election algorithm to yield better clock synchronization tree and thus lower time uncertainty bound (ϵ). In Table 9, we vary the number of SPTs from 1 to 5 evaluate both the convergence time and the converged ϵ value on the FatTree topology. We observe that with more trees constructed, the converged ϵ value reduces from expected 26.054ns to an optimal 25ns, at the cost of a slight increase of convergence time from 1.687ms to 1.973ms.

Convergence time with respect to different bandwidth settings. Next, we vary the bandwidth configuration to evaluate its impact on the convergence time. We range the bandwidth from 10Mbps to 100Mbps for the clock synchronization application and 100Mbps-1000Mbps for multicast. The results in Figure 18 verify that the convergence speed ($1/\text{Convergence-time}$) grows proportional with respect to the bandwidth allocated to FRANCIS.

FRANCIS ensures that the convergence time is hardly affected by packet losses. To demonstrate FRANCIS’s reliability with respect to packet losses, we evaluate if a large loss rate may lead to much slower convergence. The loss rate in data center networks is typically less than 0.1% , but in our simulations, we further raise the loss rate to be 1%. As shown in Figure 19, compared with an ideal lossless network, the convergence time with 1% loss rate for both tasks inflates within 15% for fast recovery and within 12% for full recovery. We argue that a 1% loss rate is rare in data centers, and such inflation is even more negligible for $< 0.1\%$ loss rates.



(a) Clock Synchronization

(b) Multicast

Figure 19: The recovery time with respect to different loss rates. Optim denotes the convergence time after the completion of distributed optimization.

I Discussion: Controller Participation

After the completion of the distributed optimization phase, we have obtained either an optimal or a near-optimal solution to recover user applications. The remaining tasks left are to clean up the memory and launch a centralized algorithm to obtain a better solution. Clean-ups of some temporary stateful objects can be done entirely in the data plane. Some other data structures in the data plane, however, should be kept in the switches’ memory to keep functioning after the distributed optimization until they are offloaded into the end hosts. Therefore, we propose to periodically launch the controller to do the management tasks such as offloading switches’ stateful objects to hosts to prevent accumulation of memory consumption and executing a centralized algorithm to further optimize the system performance. For example, the controller may offload the multicast rules stored in tor switches back to end hosts; for clock synchronization, with a negligible 0.0025% (Table 9) probability that FRANCIS fails to find an optimal clock synchronization tree, the controller is responsible to compute an optimal one. For both tasks, we argue that the need of launching the controller happens much more infrequently than occurrences of network events, adding minimal overhead.