

Using Inspiration from Synaptic Plasticity Rules to Optimize Traffic Flow in Distributed Engineered Networks

Jonathan Y. Suen

j.suen@duke.edu

*Duke University, Department of Electrical and Computer Engineering,
Durham, NC 27708, U.S.A.*

Saket Navlakha

navlakha@salk.edu

*Salk Institute for Biological Studies, Integrative Biology Laboratory,
La Jolla, CA 92037, U.S.A.*

Controlling the flow and routing of data is a fundamental problem in many distributed networks, including transportation systems, integrated circuits, and the Internet. In the brain, synaptic plasticity rules have been discovered that regulate network activity in response to environmental inputs, which enable circuits to be stable yet flexible. Here, we develop a new neuro-inspired model for network flow control that depends only on modifying edge weights in an activity-dependent manner. We show how two fundamental plasticity rules, long-term potentiation and long-term depression, can be cast as a distributed gradient descent algorithm for regulating traffic flow in engineered networks. We then characterize, both by simulation and analytically, how different forms of edge-weight-update rules affect network routing efficiency and robustness. We find a close correspondence between certain classes of synaptic weight-update rules derived experimentally in the brain and rules commonly used in engineering, suggesting common principles to both.

1 Introduction ---

In many engineered networks, a payload needs to be transported between nodes without central control. These systems are often represented as weighted, directed graphs. Each edge has a fixed *capacity*, which represents the maximum amount of traffic the edge can carry per unit time. Traffic in these networks consists of a series of *flows*, each of which contains some data that originate at a source node and attempt to reach a target node via a path in the network. For example, in vehicular transportation systems (nodes are intersections; edges are roads), cars travel from one location to another, and each road has a capacity that limits the number of cars that can traverse the road at once. In network-on-a-chip circuits (nodes are components such as

CPU and GPU cores equipped with extremely basic routers, and edges are circuit wiring between cores), tiny units of data called flits flow through the circuit, and each link can transport only a limited number of flits at a time (Cota, de Morais Amory, & Lubaszewski, 2012). On the Internet, packets navigate from one host to another, and each communication link has a capacity that constrains the rate of data flow through the link. In all these cases, the goal is to optimize performance metrics, including how long it takes for data to reach the target and how long data are queued or lost along the way due to exceeding capacities.

Two primary services are required in these networks: routing and flow control. Network routing refers to moving a payload from a source node in the network to a target node through some path (Royer & Toh, 1999), often maintained using distributed routing tables (Gavoille, 2001). The majority of networks always send traffic along the shortest path from source to target, where *shortest* refers to physical distance or transit time. In this article, we also assume routing occurs via the shortest source-target path and focus on algorithms for flow control.

Flow control determines when and how much data can be sent through the network. Engineered networks are often designed based on the concept of oversubscription or blocking, where there is not enough capacity to simultaneously service every possible flow at maximum capacity due to bottlenecks in the network. Bottlenecks can cause congestion, which results in data loss and reduction in useful throughput, since when demand for a certain link exceeds capacity, the excess data must be queued or dropped. Effective flow control simultaneously maximizes the utilization of link capacity while maintaining low loss and delay. In general, flow control is an NP-hard problem (Wang, Li, Low, & Doyle, 2003; Fortz & Thorup, 2000). It is even more challenging in online systems that serve many flows concurrently, where traffic can change unpredictably and optimization must happen in real time. Further, control logic must be implemented distributedly, with minimal communication between nodes.

To address these challenges, we developed a new neuro-inspired distributed computing model where flow control occurs by modulating (increasing or decreasing) edge weights using only 1-bit (binary) local feedback. Edge weights represent a control variable that denotes how many data should flow along the edge at the current time. For example, if 10 units of data want to travel from node u to v and if the weight of edge (u, v) is 6, then 6 units can be successfully transmitted and the other 4 units are either queued or dropped. There exists an optimal global weight distribution at every time step, which is dependent on the trade-off between maximizing data flow rates versus minimizing data drops and queueing. To inform the direction of how weights should change to approach this distribution, the network relies on 1-bit local feedback between neighboring nodes, which indicates whether the data were successfully transmitted (without being queued or dropped). This feedback is used to increase or

decrease edge weights to ensure that links are not underutilized or overloaded, respectively.

How does this relate to synaptic plasticity in the brain? Building off prior work, we argue that synaptic weight-update rules can be viewed as a distributed gradient descent process that attempts to find a weight distribution that optimizes a global objective (Bengio, Lee, Bornschein, & Lin, 2015; Bengio, Mesnard, Fischer, Zhang, & Wu, 2015). While the computational models and feedback mechanisms used to trigger weight changes in engineered networks are clearly different from those used in the brain, we consider the following the questions:

1. The direction and magnitude of weight updates—when and how much should they be increased or decreased.
2. How do these local decisions affect global performance objectives in engineering (bandwidth, drops, queueing delay).
3. Can general principles for modulating edge weights be found across engineered and neural systems.

Overall, this article makes four contributions:

1. A new neuro-inspired distributed computing model to optimize traffic flow based only on edge weight updates and 1-bit feedback between adjacent nodes.
2. A casting of long-term potentiation (LTP) and long-term depression (LTD) in terms of distributed gradient descent dynamics.
3. Simulations, using artificial and real networks, and theoretical analysis of five classes of weight-update rules.
4. Comparisons of the best-performing classes with experimental data detailing the functional forms of LTP and LTD in the brain.

1.1 Related Work. Congestion control in many distributed networks, such as the Internet, is performed on a global end-to-end basis per flow, meaning that the source of each flow regulates the rate at which it sends data into the network based on a binary feedback signal from the target node (Corless, King, Shorten, & Wirth, 2016). Our neuro-inspired per link model, described below, is stricter in its feedback constraints; each node can regulate traffic but based only on the congestion it observes on its incoming and outgoing links, independent of the source and target of the data. This model is more relevant to vehicular traffic networks (where it is impossible for a traffic control device, such as a traffic light, to know the ultimate destination of a vehicle) and network-on-a-chip circuits (where flits travel independently through the circuit, like vehicles). Traffic control algorithms have been analyzed in many forms but largely assume the problem is centralized or offline (Gayme & Topcu, 2011). Other approaches attempt to emulate a centralized algorithm by passing large-sized messages (Mateos-Nunez & Cortes, 2013). Distributed gradient descent algorithms have been

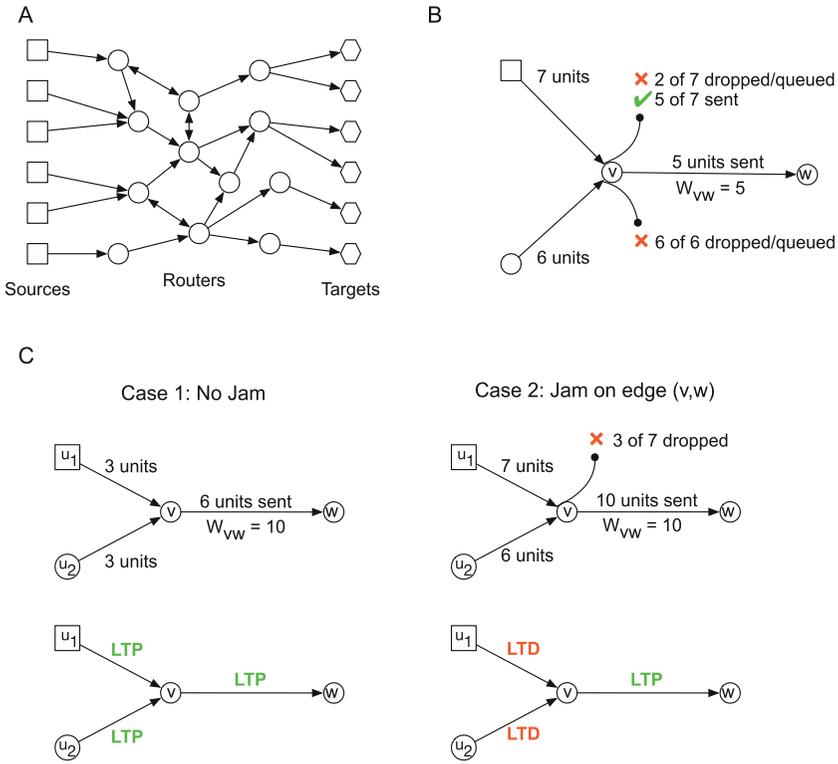


Figure 1: Model overview. (A) Input network, consisting of sources that transmit data to targets via a routing network. (B) Illustration of a congested link (v, w) , where the amount of incoming data to a node exceeds outgoing link weight, leading to dropped or queued data. Flows incoming to node v are processed in a random order. (C) If no jam occurs (case 1), all links LTP (i.e., increase their edge weight) in the next time step. If a jam does occur (case 2), the edges contributing data to the jam undergo LTD (i.e., they decrease their edge weight), and the jammed link LTPs, in an attempt to alleviate the jam in the next time step.

studied in many areas, but they also assume the ability to pass large and frequent messages, or they require significant data aggregation at individual nodes (Li, Wu, Xu, Shi, & Shi, 2014; Zinkevich, Weimer, Li, & Smola, 2010; Yuan, Ling, & Yin, 2016).

2 A Distributed Network Model for Traffic Flow Control

We are given a network $G = (V, E, W)$. The node set V is partitioned into three nonoverlapping subsets: sources $S \subset V$, targets $T \subset V$, and routers $R \subset V$. Data are transmitted from sources to targets via the routers (see Figure 1A). The edges E are directed; for simplicity, we assume each source

is connected to exactly one router (e.g., a feeder road), and each target has an incoming edge from exactly one router (e.g., a highway exit). The routers are connected with a uniform or scale-free degree topology.

The weight $W_{uv}(t)$ of each edge corresponds to the maximum amount of traffic the flow control algorithm allows to travel from $u \rightarrow v$ at time t . Each weight $W_{uv}(t) \in [1, C]$ for all $(u, v) \in E$ and for all t , where C is defined as a fixed maximum capacity for each edge. Edge weights from routers to targets are always set to C because there are no outgoing edges from the targets, and hence no possibility of downstream congestion. Weights of all other other edges will vary over time based on traffic flow and congestion.

Each source desires to send $L \gg C$ units of data to one random target; hence, there are $|S|$ flows actively competing in the network and sending data. Data for each flow are routed from the source to the target via the shortest path in the network. In practice, this is easily accomplished using a distributed routing algorithm (Gavoille, 2001).¹ In each time step t , each source injects new data into the network, the amount of which is equal to the outgoing edge weight from the source to the router it is connected to. Thus, if this weight changes, the amount of new data injected into the network by the source also change. Data arriving at a router are forwarded to the next node along the shortest path to the target. If data arrive at router u destined for node v , the data actually sent are upper-bounded by the edge weight $W_{uv}(t)$, which can change over time. If multiple flows desire to use edge $u \rightarrow v$ in the same time step, we process each flow in a random order. This means that a random incoming flow is chosen, and its data units are all sent (up to the link weight; the rest are dropped). Then another incoming flow is chosen at random, and its units are handled similarly. A flow is complete when its source has successfully transferred L data units to its target.

The only control variables of the algorithm to achieve flow control are the edge weights W at each time step. Our objective is to:

$$\begin{aligned} & \text{maximize} && \sum_t \mathcal{U}_{\mathcal{F}}(W(t)) \\ & \text{subject to} && \\ & && W_{uv}(t) \leq C \quad \text{for all } t \text{ and } (u, v) \in E. \end{aligned} \tag{2.1}$$

Here, $\mathcal{U}_{\mathcal{F}}$ is an objective function (described later) that measures how well the current edge weights perform when routing data from the flows \mathcal{F} currently in the network. To enforce the constraint, we propose two application-dependent models for penalizing excess traffic on an edge: a *drop* model

¹This model of routing is likely different from that occurring in the brain. Our work is not meant to derive a mapping between the two or suggest that their mechanisms are similar.

and a *queue* model (see Figure 1B). Let $D_{uv}(t)$ be the total data at router u that desire to be passed to node v at time t . If $D_{uv}(t) > W_{uv}(t)$, the excess data are either dropped (i.e., discarded) or queued at node u until at least the next time step. Loss models are consistent with many data networks and neural circuits (Branco & Staras, 2009). Queue models are consistent with transport networks and can help smoothe transient or bursty traffic. In each time step, data queued at a node are processed before nonqueued data (i.e., a first-in, first-out buffer). We assume a queue of infinite length, but we penalize solutions that produce long queues.

At every time step, there exists an optimal global weight distribution for a given objective function. This optimum will change as traffic demand varies over time. The goal is to track this distribution as closely as possible by applying gradient descent on the edge weights. We use 1-bit feedback between adjacent nodes, indicating whether data were successfully sent from one node to the other (i.e., no data were dropped or queued). If the transmission is successful, the data rate is increased (to probe whether higher bandwidth can be achieved). Conversely, if congestion is experienced, the data rate is decreased. This feedback thus serves as the direction of the local gradient (at the edge) for the global objective. By modifying the edge weight in accordance with this local gradient, we attempt to minimize the global objective. The magnitude of the edge weight change is based on both the direction of change and the current weight. Therefore, the algorithm seeks the point that locally maximizes traffic flow without triggering congestion.

Next, we describe how synaptic plasticity rules can serve as inspiration toward regulating these edge weights in an online, activity-dependent manner using simple distributed computation.

2.1 Synaptic Plasticity as Distributed Gradient Descent on Edge Weights. Bengio et al. and others have argued in recent work that many forms of Hebbian learning, such as spike-timing-dependent plasticity, may correspond to gradient descent toward neural activities that optimize a global objective (Bengio, Lee et al., 2015; Bengio, Mesnard et al., 2015; Osgami & Otsuka, 2015). They propose that neurons perform “inference” to try to better predict future activity given current and past data. To approach optimal activity levels, feedback signals between pre and postsynaptic neurons, such as those used to trigger LTP and LTD, cause firing rates to increase or decrease based on the gradient of the objective.

Under a connectionist assumption, where neural activity is a function of the synaptic weights coupling neurons together, the state of the network can be described by the edge weights over the population of synapses (along with other presumed constants, such as activation functions). The evolution of the system can be described by how these weights change in response to activity. In our case, activity-dependent feedback signals between adjacent nodes (described in detail below) provide a measurement

of the direction of the local gradient at the edge, which similarly triggers edge weights to increase or decrease. These weight changes thus correspond to a distributed gradient descent algorithm for finding a set of edge weights that maximizes a global objective. The movement toward the global optimal is complicated by the nonindependence of weight changes (network effects) and the uncertainty in future inputs (nonstationarity of traffic).

One critical question then is, how much should the weight increase (LTP) or decrease (LTD) following feedback? We next describe experimental data detailing what forms these rules might take in the brain.

2.2 Experimentally Derived Weight-Update Rules for LTP and LTD.

To inspire our search into different possible weight-update rules, we surveyed the recent literature for models based on experimental data (e.g., electrophysiology, imaging) that provided evidence of the functional forms of LTP and LTD (see Table 1). We categorized rules into four classes:

1. *Additive*. The change in edge weight is based on an additive constant.
2. *Multiplicative*. The change is based on a multiplicative constant.
3. *Weight dependent*. The change more generally is based on a function of the existing edge weight.
4. *Time dependent*. The change also depends on the history of recent edge weight changes.

In this article, we focus on the first three classes.

These rules will be used to derive a class of neuro-inspired distributed gradient descent algorithms for updating edge weights, as described in the next section. Table 1 is not meant to exhaustively list all forms of synaptic plasticity rules derived in the literature (see section 4) but rather to provide some basic structure into possible simple-to-implement rules and their parameters.

2.3 Distributed Algorithms for Updating Edge Weights. First, to inform the direction of the weight change (increase or decrease), in each time step, we allow 1-bit feedback between adjacent nodes. Let $\text{Jam}_{vw}(t)$ be an indicator variable equal to 1 if data at node v are lost or queued due to congestion on edge (v, w) at time t , and 0 otherwise. Assume data for a flow are traveling from node u to v to w . When considering how to modify the weight of edge (u, v) , we need to consider what happens on the adjacent downstream edge (v, w) where traffic is flowing. Intuitively, if $\text{Jam}_{vw}(t) == 1$, then the incoming edge weight $W_{uv}(t + 1)$ should LTD since it contributed data to the jam. Further, the edge weight $W_{vw}(t + 1)$ should LTP to attempt to alleviate the congestion. If neither edge (u, v) nor (v, w) is jammed, then both should LTP (see Figure 1C). Thus, the Jam term serves as a 1-bit measurement of the direction of the local gradient at edge (u, v) . Overall, the logic implemented at each edge (u, v) is shown in algorithm 1.

Table 1: Correspondence between Synaptic Plasticity Rules with Experimental Support and Network Control Algorithms Used in Engineering.

Plasticity	Rule	Equation	Brain	Engineering
LTP	Additive	$W(t + 1) = W(t) + k_i$	Song, Miller, and Abbott, 2000; Kopec, Li, Wei, Boehm, and Malinow, 2006	Chiu and Jain, 1989; Corless et al., 2016
	Multiplicative	$W(t + 1) = W(t) \times k_i$	—	Chiu and Jain, 1989
	Weight dependent	$W(t + 1) = W(t) \times f(W(t))$	Oja, 1982; Bi and Poo, 1998; van Rossum, Bi, and Turrigiano, 2000	Bansal and Balakrishnan, 2001; Kelly, 2003; Ha, Rhee, and Xu, 2008
	History dependent	$W(t + 1) = f(W(t), \dots, W(t - \Delta(t)))$	Bienenstock, Cooper, and Munro, 1982; Froemke, Debanne, and Bi, 2010; Cooper and Bear, 2012; Kramar et al., 2012	Jin, Guo, Matta, and Bestavros, 2002
LTD	Subtractive	$W(t + 1) = W(t) - k_d$	Song et al., 2000	—
	Multiplicative	$W(t + 1) = W(t) \times k_d$	Bi and Poo, 1998; van Rossum et al., 2000; Zhou, Homma, and Poo, 2004	Chiu and Jain, 1989; Corless et al., 2016
	Weight dependent	$W(t + 1) = W(t) \times f(W(t))$	Oja, 1982	Bansal and Balakrishnan, 2001; Corless and Shorten, 2012
	History dependent	$W(t + 1) = f(W(t), \dots, W(t - \Delta(t)))$	Bienenstock et al., 1982; Froemke et al., 2010; Cooper and Bear, 2012	Jin et al., 2002

Notes: The “—” indicates lack of support. Reviews of many of these rules are provided by Kepecs, van Rossum, Song, and Tegner (2002), Morrison, Diesmann, and Gerstner (2008), Watt and Desai (2010), and Froemke, Debanne, and Bi (2010). History-dependent rules may also depend on the history of recent activity instead of only the history of recent weights.

Algorithm 1: Logic for Applying LTP or LTD, Implemented at Each Edge (u, v) .

```

if  $\text{Jam}_{uv}(t) == 1$  then
    Apply LTP to  $W_{uv}(t + 1)$ 
else
    if  $\text{Jam}_{vw}(t) == 1$  and  $(u, v)$  contributed data that jammed  $(v, w)$ , for any  $w$  then
        Apply LTD to  $W_{uv}(t + 1)$ 
    else if  $\text{Jam}_{vw}(t) == 0$  for all  $w$  that received data from  $u$  then
        Apply LTP to  $W_{uv}(t + 1)$ 
    end if
end if

```

We assume that a node (in this case, v) can modify the edge weight of both its incoming and outgoing edges. In synapses, this may be achieved by modulating the presynaptic release probability or the number of post-synaptic receptors (e.g., Costa, Froemke, Sjöström, & van Rossum, 2015; Markram, Gerstner, & Sjöström, 2012; Yang & Calakos, 2013; Fitzsimonds, Song, & Poo, 1997) or via other gating mechanisms (Vogels & Abbott, 2009). In data networks, a node can pause incoming data by transmitting a signal and can pause outgoing data by simply not transmitting. If an edge gets both LTP and LTD signals in a time step, it defaults to LTD. In every time step, the weight of every edge that carries data will either LTP or LTD.

Second, to determine the magnitude of the weight change, we consider the following weight-update rules for LTP and LTD:

$$W_{uv}(t + 1) = \begin{cases} W_{uv}(t) + k_i, & k_i > 0 & \text{LTP: Additive increase (AI)} \\ W_{uv}(t) \times k_i, & k_i > 1 & \text{LTP: Multiplicative increase (MI)} \\ W_{uv}(t) - k_d, & k_d > 0 & \text{LTD: Subtract decrease (SD)} \\ W_{uv}(t) \times k_d, & 0 < k_d < 1 & \text{LTD: Multiplicative decrease (MD)} \end{cases}$$

We consider four combinations of LTP and LTD: AIMD, AISD, MIMD, and MISD. Each of these algorithms has some theoretical or experimental basis, or both (see Table 1). For example, AISD was proposed by Kopeck et al. (2006) and Song et al. (2000). AIMD was proposed by van Rossum et al. (2000) and Delgado, Gomez-Gonzalez, and Desai (2010). Multiplicative decrease rules have been proposed by Zhou et al. (2004), among others.

We also compare to an algorithm based on the classic Oja learning rule (Oja, 1982):

$$W_{uv}(t+1) = \begin{cases} W_{uv}(t) + k_i \left(1 - \frac{D_{uv}(t)^2}{W_{uv}(t)C} \right), & k_i > 0 \quad \text{LTP: Oja} \\ W_{uv}(t) - k_d \left(1 + \frac{D_{uv}(t)^2}{W_{uv}(t)C} \right), & k_d > 0 \quad \text{LTD: Oja} \end{cases}.$$

Unlike the previous rules, Oja uses the activity (traffic) of the edge as a variable, where $D_{uv}(t)$ is the amount of data traversing edge (u, v) at time t . This rule is slightly different from the typical Oja rule where the change in the weight for the i th input, $\Delta W_i = \alpha(x_i y - y^2 W_i)$ (learning weight α , synaptic input x , and output y). The squared term functions as a decay on the weight. We include a similar activity-dependent squared decay term ($D_{uv}(t)^2/W_{uv}(t)$), but we normalize it by C to lie within the required weight range. Our term decreases the effect of LTP and increases the effect of LTD as the link approaches capacity. This allows more aggressive k_i and k_d to quickly adjust traffic.

Another algorithm we compare is bang-bang control. This rule is often used in neural circuit design (Feng & Tuckwell, 2003; Zanutto & Staddon, 2007) and in engineering (Lazar, 1983) to control and stabilize activity:

$$W_{uv}(t+1) = \begin{cases} C & \text{LTP: Bang-bang} \\ 1 & \text{LTD: Bang-bang} \end{cases}.$$

Finally, we compare to a baseline rule, Max Send, which keeps all edge weights fixed at C in every time step. For all algorithms, if a weight equals C and is triggered to LTP, it stays at C . Likewise, if a weight equals 1 and is triggered to LTD, it stays at 1.

We consider only integer units of data, so for all update rules, W is rounded to an integer. To prevent a link from getting stuck at 0 weight (e.g., for MISD), every weight was required to have a minimum of 1. The link capacity $C \gg 1$, so the integer rounding and minimum value were negligible in terms of overall performance.

2.4 Objective Functions, Simulations, and Data. Next, we describe network performance measures to quantify how well these rules behave toward optimizing global objectives ($\mathcal{U}_{\mathcal{F}}$). The objective functions we selected are typically used to evaluate performance in engineered networks (Pande, Grecu, Jones, Ivanov, & Saleh, 2005; Ahn, Danzig, Liu, & Yan, 1995).

Let \mathcal{F} define the set of $|\mathcal{S}|$ competing flows and L be the load of each flow. The three objective functions are:

- *Bandwidth*: Amount of data successfully transferred per time step, averaged over all flows: $|\mathcal{F}|^{-1} \sum_i L/\text{time}(\mathcal{F}_i)$, where $\text{time}(\mathcal{F}_i)$ is the number of time steps for flow i to complete.
- *Drop penalty*: Percentage of data lost, averaged over all flows: $|\mathcal{F}|^{-1} \sum_i \text{lost}(\mathcal{F}_i)/L$, where $\text{lost}(\mathcal{F}_i)$ is the number of data lost by flow i over all time steps until completion. The drop penalty may go above 100% if more data sent by the source are lost than delivered.
- *Queue penalty*: The percentage of data queued per hop, averaged over all flows: $|\mathcal{F}|^{-1} \sum_i \text{queued}(\mathcal{F}_i)/(L \times \text{path}(\mathcal{F}_i))$, where $\text{queued}(\mathcal{F}_i)$ is the total number of data inserted into queues and $\text{path}(\mathcal{F}_i)$ is the path length of flow i .
- *Parameter robustness*: One critical component of these algorithms is that they must work well in general. Traffic is highly dynamic, and thus optimizing parameters for one particular traffic regime or network topology will not be sufficient for real-world use. We thus focus on the robustness or sensitivity of each algorithm by testing the variability in their performance across a broad range of parameters.

2.4.1 Simulation Framework. We created a directed network with N sources, N targets, and N routers, where $N = 100$ or 1000 . Each source is connected to exactly one random router (the same router can have an edge from multiple different sources), and each target has an incoming edge from one random router. The router-router network was defined using a uniform or scale-free degree topology with each router connected to six other routers (Corless et al., 2016). We defined N concurrent flows, one starting from each source and each selecting a random target (the same target may be selected multiple times across flows). Each flow contained $L = 100 \times C$ data units, as 100 is roughly the average number of round-trip times taken for a data transfer on the Internet (Corless et al., 2016). The weight of each edge was initialized to the maximum capacity, C , in order to immediately experience contention. The mean path length of the artificial network was 4.7, large enough to provide several links of potential contention. Each performance measure was averaged over 25 repeat simulations.

2.4.2 Parameter Variation. For the weight-update rules, we varied AI ($k_i \in [1.0, 9.0]$), SD ($k_d \in [1.0, 9.0]$), MI ($k_i \in [1.1, 1.9]$), MD ($k_d \in [0.1, 0.9]$).

2.4.3 Real-World Data. We used the CAIDA autonomous system (AS) relationship data to generate a graph based on the Internet routing network (www.caida.org/data/as-relationships). Each AS represents the highest-level routing subnetwork on the Internet. The CAIDA data contained connectivity between 53,195 AS subnetworks. We treat each AS as a single routing node in our model. We created the same number of sources and targets (53,195 each); each source was connected to one random AS, and each target had one incoming edge from a random AS. The network

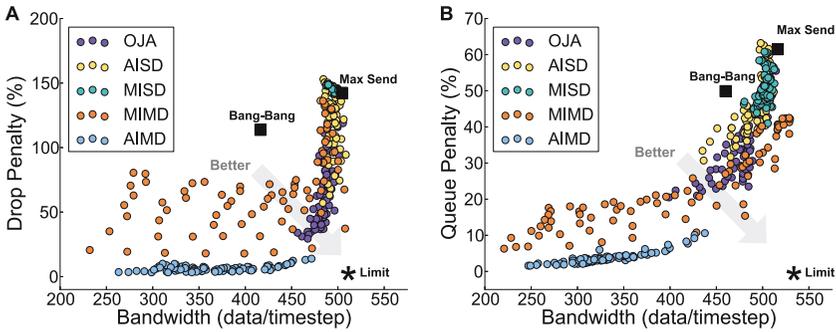


Figure 2: Comparison of seven activity-dependent weight-update rules. Bandwidth versus (A) drop penalty and (B) queue penalty. The lower right of the plot corresponds to an empirical upper bound, which occurs when the bandwidth is that of the highest observed algorithm but with zero drop or queue penalty. Each dot corresponds to an algorithm run using different values of the increase and decrease parameters (k_i, k_d). AIMD, MIMD, and Oja perform the best.

contained 537,582 total directed edges. To simulate flows, we selected about 1% of the sources (500) and paired each source with a random target. Each performance measure was averaged over 10 repeat simulations. We set the capacity, $C = 1000$.

3 Results

First, we describe the performance of each weight-update algorithm against the global objectives (bandwidth, drops, and queueing) using both simulated and real-world networks. Second, we support these results by analytically deriving the performance of each algorithm as it adapts to changing traffic demands. Third, we describe how the best-performing rules compare to those commonly used in engineering.

3.1 Observations from Simulations and Real-World Network Flows.

We first compared the performance of the seven edge-weight update algorithms (see section 2.3) by simulation. Each algorithm was evaluated according to the bandwidth offered and the amount of data dropped or queued.

The two strongest-performing algorithms were AIMD and MIMD, with Oja lying in between (see Figure 2). AIMD, for some parameters, achieved a bandwidth comparable to other algorithms, but its main strength was in reducing the drop penalty by at least one order of magnitude (averages: AIMD = 6.4%, OJA = 69.2%, MIMD = 69.3%, AISD = 124.4%, MISD = 145.2%; see Figure 2A). The former was due to its conservative

rule for increasing edge weights for LTP (additive) and the latter to its aggressive edge weight decrease for LTD (multiplicative). AISD and MISD showed a very high drop penalty primarily because upon contention, edge weights were only decreased subtractively; this led to slow adaptation though higher bandwidth because few links were ever underutilized. In general, the Oja rule (a variant of AISD) improved over AISD but still achieved a much higher drop penalty compared to AIMD, also due to its conservative decrease. MIMD showed great sensitivity to algorithm parameters, some of which performed well. While keeping all edge weights at maximum capacity may be intuitively appealing (Max Send), this is in general not a good solution because any downstream bottleneck will result in massive data drops. We also observed similar trends for all algorithms using the queue model (see Figure 2B).

Overall, for both models, the multiplicative decrease algorithms (AIMD, MIMD) and the Oja algorithm demonstrated a better trade-off between bandwidth and drop or queue penalties than other algorithms, indicating their ability to more closely approach the optimal edge weights. We also tested these observations for larger networks with different router connectivity topologies and observed similar overall trends (see the appendix and Figure 4).

Next, we tested how well and how quickly each algorithm could adapt to new traffic, simulated on a real Internet backbone routing network. The simulation was for 3000 time steps. During $t < 1000$ and $t > 2000$, 500 flows concurrently competed. During $1000 \leq t \leq 2000$, 500 additional flows were temporarily added (“rush hour”). All algorithms demonstrated some reduction in bandwidth when rush hour begins due to the additional number of competing flows. However, AIMD and MIMD incurred the least transient drop penalty (i.e., the drop penalty incurred immediately after rush hour starts; see Figure 3B). Overall, AIMD yielded significantly lower transient drop penalty than all other algorithms ($p < 0.01$; two-sample t -tests) and significantly lower overall drop penalty than MIMD. AIMD and MIMD also produced only a 4% to 5% difference in bandwidth compared to the other algorithms (average data per time step: AIMD = 776 ± 6.82 , MIMD = 786 ± 6.94 , OJA = 814 ± 19.38 , AISD = 819 ± 6.11 , MISD = 820 ± 5.93 at $t = 1500$; see Figure 3A). These results suggest that AIMD and MIMD adapt faster to changing traffic, yielding fewer transient drops with comparable bandwidth.

Next, to support these observations, we formally analyze the adaptive behavior of each algorithm.

3.2 Analyzing Transient Response Times for AIMD, MIMD, AISD, MISD, and Oja. An important aspect of algorithm performance is its nonstationary or transient behavior—how well it adapts when traffic suddenly increases and the available bandwidth per flow decreases (see Figure 3). Real networks are never static; all flows experience some level

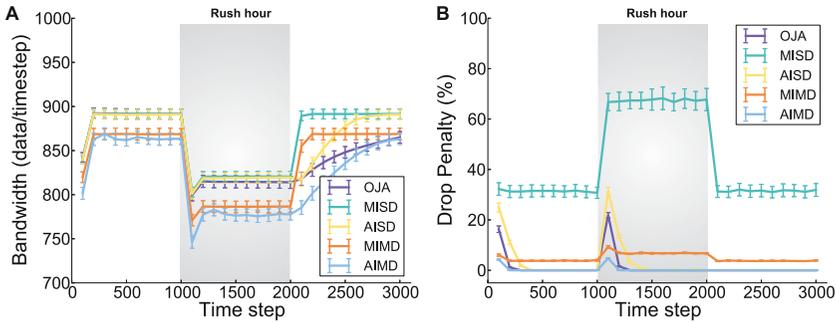


Figure 3: Adaptation to changing traffic demands. Traffic simulated on a real Internet backbone routing network for 3000 time steps. Performance measures were averaged every 100 time steps. We selected the parameters for each algorithm that had the highest bandwidth in Figure 2 while being within 1% of the minimum drop penalty for the algorithm. AIMD and MIMD showed the least additional penalty (B) due to rush hour, suggesting quick adaptation, while also yielding a similar bandwidth (A) as MISD, AISD, and Oja.

of perturbation due to varying traffic. Thus, as opposed to analyzing convergence properties (as is typically done for gradient descent algorithms), we analyzed a simple but informative scenario: the performance of each algorithm when a second flow is added to a link that is initially serving only a single flow at maximum capacity.

Assume the link (u, v) under consideration has fixed weight C and that there are two flows starting from s_1 and s_2 that both need to use (u, v) to reach their downstream target (and no other link in the network is limiting). Let $W_{s_1 u}(t) = C$; then when the second flow begins, assume $W_{s_2 u}(t) = 1$, where $C \gg 1$. Assume a single LTD operation on (u, v) will lower the total traffic sent by both flows along (u, v) to be below C , hence alleviating the congestion. After the initial congestion event, there will be ≥ 1 time steps of LTP before congestion reoccurs. Let n be the number of time steps before congestion reoccurs. We wish to find the amount of data dropped or queued (called *overshoot*) at time n , defined as the difference between the amount of data desired by both flows along edge (u, v) and C at the moment LTD is reactivated.

Theorem 1 (*overshoot of AIMD*). *The two-flow transient response of AIMD has an overshoot that increases linearly with n but is reduced by a term proportional to the capacity, C .*

Proof. At the first time step (-1 to simplify notation), the amount desired by both flows on (u, v) is:

Time t	Flow 1	Flow 2	Jam $_{uv}(t)$
-1	C	1	True

Since the sum of the flows is greater than C, congestion occurs and the jam indicator variable Jam $_{uv}(t)$ is true. In the next step, LTD via multiplicative decrease is applied:

Time t	Flow 1	Flow 2	Jam $_{uv}(t)$
0	Ck_d	k_d	False

This brings the total desired traffic along (u, v) under C. Additive increase then occurs for n steps:

Time t	Flow 1	Flow 2	Jam $_{uv}(t)$
1	$Ck_d + k_i$	$k_d + k_i$	False
2	$Ck_d + 2k_i$	$k_d + 2k_i$	False
\vdots	\vdots	\vdots	\vdots
n	$Ck_d + nk_i$	$k_d + nk_i$	True

At time step n , the overshoot along (u, v) , which is the excess traffic over C, is:

$$\begin{aligned}
 (\text{Flow 1} + \text{Flow 2}) - C &= (Ck_d + nk_i + k_d + nk_i) - C, \\
 &\approx 2nk_i + C(k_d - 1),
 \end{aligned}
 \tag{3.1}$$

assuming the single k_d term is negligibly small. Note that since $k_d \in (0, 1)$, the second term is negative. □

We similarly derived the overshoot of MIMD, AISD, and MISD (see the appendix and summary in Table 2). The theoretical performances of the algorithms on the two-flow case correlate well with the simulated performance using hundreds of concurrent flows in a larger network. Both AI algorithms have an overshoot that has a $2nk_i$ term. For AIMD, this term is reduced by $C(k_d - 1)$. AISD, however, is hurt by a slow subtractive decrease, and hence only reduces the $2nk_i$ term by $2k_d$. Since $|C(k_d - 1)| > |2k_d|$ for large values of C, AIMD typically performs better (compare the blue dots versus yellow dots in Figure 2).

The MIMD overshoot shows a complex dependence of k_i , k_d , and C, which makes performance highly parameter dependent, as we also observe by simulation (see the variability of orange dots in Figure 2). MISD has uniformly high overshoot, since $k_i^n > 1$ is multiplied by a large constant C and thus performs poorly. Finally, Oja uses an AISD rule with a weight-dependent squared decay; this decay cancels out the additive increase term

Table 2: Analysis of Overshoot of Weight-Update Algorithms.

Rule	Transient Overshoot	Balanced (n)	Aggressive Increase (n)	Aggressive Decrease (n)
AIMD	$2nk_i + C(k_d - 1)$	0.5 (250)	101 (3)	0.1 (450)
AISD	$2(nk_i - k_d)$	1.0 (5)	196 (1)	1.0 (100)
MIMD	$C(k_d k_i^n - 1)$	72.8 (8)	126 (2)	84.6 (25)
MISD	$C(k_i^n - 1)$	95.6 (1)	494 (1)	90.2 (2)
OJA	Undershoots	-1 (∞)	-1 (∞)	-1 (∞)
OPT	0 (∞)	0 (∞)	0 (∞)	0 (∞)

Notes: Analytic (left) and simulated (right) overshoot values of each algorithm using $C = 1000$. The variable n corresponds to the number of time steps before congestion re-occurs. The Oja-based algorithm always undershoots for the simple two-flow case; however, a closed-form solution appears difficult to derive because the algorithm is dependent on edge traffic. The optimal solution (OPT) has zero overshoot. For simulations of other algorithms, we selected three sets of parameters: Balanced: additive ($k_i = 1.0, k_d = 5.0$), multiplicative ($k_i = 1.1, k_d = 0.5$). Aggressive increase: additive ($k_i = 100.0, k_d = 5.0$), multiplicative ($k_i = 1.5, k_d = 0.5$). Aggressive decrease: additive ($k_i = 1.0, k_d = 100.0$), multiplicative ($k_i = 1.1, k_d = 0.1$).

as the traffic of flow 1 and 2 approaches C . This leads to performance that always undershoots for this simple two-flow case, improving drop and queue penalty over AISD; however, in practice, when many flows concurrently compete and overshoot does occur (as in Figures 2 and 3), it is limited by the weak decrease term, as the decay only doubles k_d at best.

We also simulated the transient overshoot under the assumptions of theorems 1 to 4 using different parameter settings (see Table 2, right). These simulations further validate the theorems, showing that AIMD overshoots the least, while MIMD varies from the second best to second worst, depending on k_i and k_d . Further analysis of each algorithm is provided in the appendix.

3.3 Comparing Distributed Gradient Descent Algorithms in the Brain and the Internet. Simulations and theory both suggest that AIMD achieves a robust and well-balanced trade-off between bandwidth and drop or queue penalties, with MIMD and Oja also performing comparably depending on the parameters selected. This implies that these algorithms approach the optimal global edge-weight distribution more quickly and accurately than other distributed gradient descent algorithms, including MISD, AISD, bang-bang, and max-send.

In the brain, the additive increase and multiplicative decrease rule (AIMD) for LTP and LTD, respectively, has strong theoretical and experimental support (see Table 1), particularly over MI and SD models. The

AIMD algorithm is also very similar to the rule proposed by van Rossum et al. (2000) and has been commonly referred to as the mixed-weight update rule (Delgado et al., 2010). This rule, in neural network simulations, has been shown to produce stable Hebbian learning compared to many other rules (van Rossum et al., 2000; Billings & van Rossum, 2009).

While Oja was one type of weight-dependent rule, other rules have also been proposed where a weak (i.e., low-weight) synapse that undergoes LTP is strengthened by a greater amount than a strong synapse that undergoes LTP, and vice versa for LTD (see Table 1). Prior work has also suggested that individual synapses may have a “memory” that allows for more sophisticated update rules to be implemented, including history-dependent updates (Lahiri & Ganguli, 2013). One such form of update, integral control in engineering, uses the time integral of a variable from $t = 0$ to the present. We show the general forms of these rules in Table 1 but do not explore them further here.

Interestingly, in engineering, AIMD also lies at the heart of the most popular congestion control algorithm used on the Internet today: the transmission control protocol (TCP; Corless et al., 2016). In contrast to our per-link model, congestion control on the Internet is performed on a global end-to-end basis per flow, meaning that the source of each flow regulates its transmission rate based on a binary feedback signal (called an ACK, or acknowledgment) sent by the target. If there is a long delay before the ACK is received (or if the ACK is never received at all), congestion is assumed, and the source decreases its rate of sending data by a multiplicative constant (often 0.5). Otherwise the source increases its rate by an additive constant (often 1.0). TCP was also designed with the goal of converging to steady-state flow rates over time in a gradient-descent-like manner (Shakkottai & Srikant, 2007; Jose, Yan, Alizadeh, McKeown, & Katti, 2015). Thus, despite different models and objectives, both the brain and the Internet may have discovered a similar distributed algorithm for optimizing network activity using sparse, local feedback.

4 Discussion

Our work connects distributed traffic flow control algorithms in engineered networks to synaptic plasticity rules used to regulate activity in neural circuits. While we do not claim that there is a one-to-one mapping between mechanisms of synaptic plasticity and flow control problems, we showed how both problems can be viewed abstractly in terms of gradient descent dynamics on global objectives, with simple local feedback. We performed simulations and theoretical analyses of several edge-weight update rules and found that the additive-increase multiplicative-decrease (AIMD) algorithm performed the best in terms minimizing drops or queueing with comparable bandwidth as other algorithms. This algorithm also matched experimental data detailing the functional forms of LTP and LTD in the

brain and on the Internet, suggesting a similar design principle used in biology and engineering. Further, these weight rules use limited (1-bit) local communication and hence may be useful for implementing energy-efficient and scalable flow control in other applications, including integrated circuits, wireless networks, or neuromorphic computing.

There are many avenues for future work. First, other plasticity rules, such as short-term plasticity, may be explored within our framework. Second, in cases where source or receiver rates are fixed, the payload needs to be routed over alternative paths (i.e., routes may change over time basic on traffic; Isa, Mohamed, & Yusoff, 2015). This requires that heavily used edges become down-weighted and unused edges become more attractive, which effectively performs load balancing over all resources (edges) in the network. Biologically, similar behavior is observed due to homeostatic plasticity mechanisms, which may inspire algorithms for this problem. Third, these distributed gradient descent updates rules may be useful in machine learning applications for nonstationary learning. Fourth, more sophisticated weight and history-dependent update rules already explored in engineering may provide insight into their form and function in the brain. Overall, we hope our work inspires closer collaborations between distributed computing theorists and neuroscientists (Navlakha et al., 2015).

Appendix: Analyzing Transient Response Times for AIMD, MIMD, AISD, MISD, and Oja

In this appendix, we derive the transient overshoot of AIMD, MIMD, AISD, and MISD.

Theorem 1 (*overshoot of AIMD*). *The two-flow transient response of AIMD has an overshoot that increases linearly with n but is reduced by a term proportional to the capacity, C .*

Proof. See the main text for the derivation. The overshoot is:

$$\begin{aligned} (\text{Flow 1} + \text{Flow 2}) - C &= (Ck_d + nk_i + k_d + nk_i) - C, \\ &\approx C(k_d - 1) + 2nk_i, \end{aligned} \tag{A.1}$$

assuming the single k_d term is negligibly small. □

Theorem 2 (*overshoot of MIMD*). *The two-flow transient response of MIMD is highly parameter dependent. Based on k_d and k_i , the scaling of overshoot can be dominated by C or by a power of n .*

Proof. Following the proof of theorem 1 in the main text, the time evolution will be:

Time t	Flow 1	Flow 2	Jam _{uv} (t)
-1	C	1	True
0	Ck_d	k_d	False
1	Ck_dk_i	k_dk_i	False
2	$Ck_dk_i^2$	$k_dk_i^2$	False
\vdots	\vdots	\vdots	\vdots
n	$Ck_dk_i^n$	$k_dk_i^n$	True

The overshoot for MIMD is

$$\begin{aligned}
 (\text{Flow 1} + \text{Flow 2}) - C &= k_d k_i^n (C + 1) - C. \\
 &\approx C(k_d k_i^n - 1),
 \end{aligned}
 \tag{A.2}$$

if $C \gg 1$. Thus, the overshoot shows a positive dependence on C (unlike AIMD, which has a negative dependence on C) and a power dependence on n . □

Theorem 3 (overshoot of AISD). *The two-flow transient response of AISD increases linearly with n but does not scale relative to the link capacity, C .*

Proof. The time evolution will be:

Time t	Flow 1	Flow 2	Jam _{uv} (t)
-1	C	1	True
0	$C - k_d$	$1 - k_d$	False
1	$C - k_d + k_i$	$1 - k_d + k_i$	False
2	$C - k_d + 2k_i$	$1 - k_d + 2k_i$	False
\vdots	\vdots	\vdots	\vdots
n	$C - k_d + nk_i$	$1 - k_d + nk_i$	True

The overshoot for AISD is thus:

$$(\text{Flow 1} + \text{Flow 2}) - C = 2(nk_i - k_d).
 \tag{A.3}$$

□

In our theorems, we ignore the constraint that $W \geq 1$. This limit in our implementation was due to integer rounding to ensure MISC links do not get stuck at 0 weight; hence, the theorems we present here are more general. If we apply this limit, it affects the weight of flow 2 at $t = 0$, which should equal 1. In all cases except AISD, the difference is removed by the

approximation at the end. For AISD, equation A.3 becomes $2nk_i - k_d$, a negligible change when $n \gg 1$.

Theorem 4 (*overshoot of MISD*). *The two-flow transient response of MISD increases to the power of n and as a product with C .*

Proof. The time evolution will be:

Time t	Flow 1	Flow 2	Jam _{uv} (t)
-1	C	1	True
0	$C - k_d$	$1 - k_d$	False
1	$(C - k_d)k_i$	$(1 - k_d)k_i$	False
2	$(C - k_d)k_i^2$	$(1 - k_d)k_i^2$	False
\vdots	\vdots	\vdots	\vdots
n	$(C - k_d)k_i^n$	$(1 - k_d)k_i^n$	True

The overshoot for MISD is thus

$$\begin{aligned}
 (\text{Flow 1} + \text{Flow 2}) - C &= k_i^n (C - 2k_d + 1) - C \\
 &\approx C (k_i^n - 1), \tag{A.4}
 \end{aligned}$$

since C is large relative to $2k_d + 1$. □

General conclusions can be drawn from these relations that correlate well with the simulation results. We can bound the overshoot of both AI algorithms (AIMD and AISD) by $|\mathcal{F}|k_i$, where $|\mathcal{F}|$ is the number of flows that must share a link because each individual flow will not overshoot more than one k_i . The more precise overshoot of AIMD, $C(k_d - 1) + 2nk_i$ (theorem 1 main text), shows a significant capacity-dependent stabilizing factor, as $C(k_d - 1)$, which is negative, counteracts the factor of n . When we repeated the AIMD flow simulation (see Table 2) with $C = 50$ (as opposed to $C = 1000$), the overshoot of AIMD increased, as expected.

AISD has a conservative increase similar to AIMD but suffers because subtractive decrease slowly adjusts when available bandwidth decreases. While some parameter settings may overcome this, it is difficult in practice to tune the SD constant to be effective in all scenarios. Simulation results support this observation, showing a large drop penalty during rush hour and a slow recovery after traffic subsides (see Figure 3). For AISD, if multiple flows share a link, transient overshoot is approximately $|\mathcal{F}|nk_i$. When a large number of flows share a link (e.g., millions of flows on an Internet backbone), overshoot will be very large.

The Oja-inspired algorithm is based on AISD but subtracts a normalized traffic-dependent quadratic term. This attempts to correct for the slow LTD decrease of AISD, especially at high weights, and we do observe

improvement over AISD in our network simulations. However, we also observe larger drop and queue penalties for Oja compared to AIMD due to its still rather conservative decrease following congestion. The decay term can also completely counteract the AI contribution when $D_{uv} \rightarrow \sqrt{WC}$ (i.e., when edge utilization is high), thus potentially leading to edge underutilization. This convergence to a weight less than C causes LTD never to be triggered (hence, the ∞ term in Table 2 for this simple two-flow case). The lack of periodic overshooting appears to be key for the drop and queue penalty improvements over AISD.

The overshoot and performance of MIMD are highly dependent on C relative to the k_i and k_d parameters, explaining the scattered performance of MIMD in Figure 2. With optimal parameter tuning, MIMD can be made to operate well under transient behavior, which is seen when certain MIMD points reach the AIMD region.

MISD can easily be seen as worse than AISD in terms of drop or queue penalties in our simulations (see Figures 2 and 3). Intuitively, this is because MISD reduces slowly (subtractively) but increases aggressively (multiplicatively). The transient overshoot analysis for MISD shows that it increases as a product of C and a constant to the power of n . Since the decrease term is weak, n will be small, meaning that poor performance will be especially seen for high-capacity links.

A.1 Algorithm Performance with Additional Topologies. We performed simulations with 10-fold larger networks ($N = 1000$), with both uniform and scale-free degree distributions. The scale-free topology was derived using the Barabasi-Albert preferential attachment model (Barabasi & Albert, 1999). When generating graphs using this model, new nodes are connected to existing nodes with probability proportional to their existing degree. This mechanism has been shown to produce a power-law degree distribution. We found no qualitative change in our conclusions here (see Figure 4) compared to the results discussed in the main text. Thus, our results are applicable to at least two classes of network topologies: uniform (inspired by grid-like road networks) and power law (Internet). This invariance is likely due to the distributed nature of the flow control algorithms. The overshoot theorems exemplify this, having no assumptions on network size and topology.

A.2 Changes in Edge Weights (W) under Dynamic Traffic. To study how W (edge weights) changed under the rush hour protocol of Figure 3, we plotted the average W (computed in 10 time-step bins) for all used source-to-router links in the network (see Figure 5). We focused on source-router links as these primarily control the amount of new data injected into the network in each time step. Error bars in Figure 5 correspond to the standard error of the average over 10 trials. Both MIMD and AIMD reduce

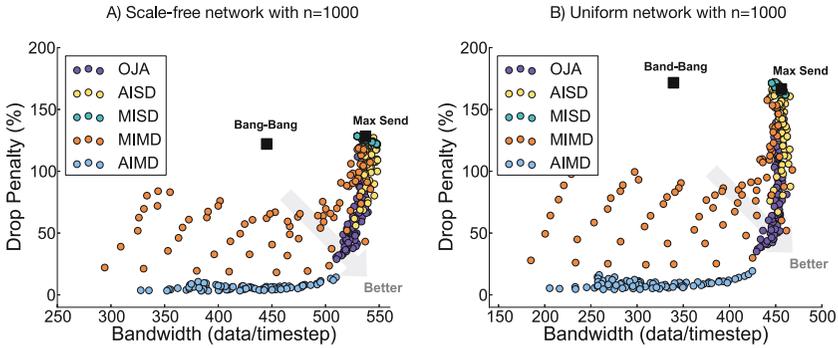


Figure 4: Performance using additional topologies. Comparison of all the algorithms using a scale-free degree distribution (A) and a uniform degree distribution (B) with 1000 nodes.

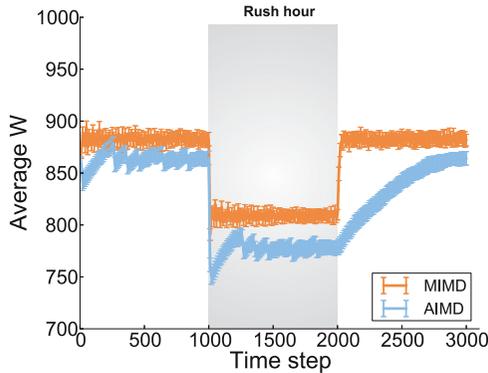


Figure 5: Changes in the edge weight under the dynamic traffic protocol. Error bars correspond to the standard error of the average over 10 trials.

edge weights in response to rush hour, showing that they handle excess traffic by reducing bandwidth instead of dropping data. The characteristic oscillatory probing nature of both algorithms is also apparent. In both cases, the range of W in each time step is narrowly bounded, indicating network stability.

Acknowledgments

J.Y.S. and S.N. thank Alex Lang and Ben Regner for helpful comments on the article. S.N. thanks the Army Research Office (grant DOD W911NF-17-1-0045) for funding support.

References

- Ahn, J. S., Danzig, P. B., Liu, Z., & Yan, L. (1995). Evaluation of TCP Vegas: Emulation and experiment. *SIGCOMM Comput. Commun. Rev.*, 25(4), 185–195.
- Bansal, D., & Balakrishnan, H. (2001). Binomial congestion control algorithms. In *Proceedings of IEEE Infocom 2001*. Piscataway, NJ: IEEE.
- Barabasi, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512.
- Bengio, Y., Lee, D., Bornschein, J., & Lin, Z. (2015). *Towards biologically plausible deep learning*. arXiv:abs/1502.04156
- Bengio, Y., Mesnard, T., Fischer, A., Zhang, S., & Wu, Y. (2015). *STDP as presynaptic activity times rate of change of postsynaptic activity*. arXiv:abs/1509.05936
- Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, 18(24), 10464–10472.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, 2(1), 32–48.
- Billings, G., & van Rossum, M. C. (2009). Memory retention and spike-timing-dependent plasticity. *J. Neurophysiol.*, 101(6), 2775–2788.
- Branco, T., & Staras, K. (2009). The probability of neurotransmitter release: Variability and feedback control at single synapses. *Nat. Rev. Neurosci.*, 10(5), 373–383.
- CAIDA AS Relationships Dataset (Feb. 1, 2016), <http://www.caida.org/data/as-relationships>
- Chiu, D.-M., & Jain, R. (1989). Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Comput. Netw. ISDN Syst.*, 17(1), 1–14.
- Cooper, L. N., & Bear, M. F. (2012). The BCM theory of synapse modification at 30: Interaction of theory with experiment. *Nat. Rev. Neurosci.*, 13(11), 798–810.
- Corless, M., King, C., Shorten, R., & Wirth, F. (2016). *AIMD dynamics and distributed resource allocation*. Philadelphia: Society for Industrial & Applied Mathematics.
- Corless, M., & Shorten, R. (2012). Deterministic and stochastic convergence properties of AIMD algorithms with nonlinear back-off functions. *Automatica*, 48(7), 1291–1299.
- Costa, R. P., Froemke, R. C., Sjöström, P. J., & van Rossum, M. C. (2015). Unified pre- and postsynaptic long-term plasticity enables reliable and flexible learning. *Elife*, 4.
- Cota, É., de Moraes Amory, A., & Lubaszewski, M. (2012). *Reliability, availability and serviceability of networks-on-chip*. New York: Springer Science.
- Delgado, J. Y., Gomez-Gonzalez, J. F., & Desai, N. S. (2010). Pyramidal neuron conductance state gates spike-timing-dependent plasticity. *J. Neurosci.*, 30(47), 15713–15725.
- Feng, J., & Tuckwell, H. C. (2003). Optimal control of neuronal activity. *Phys. Rev. Lett.*, 91(1), 018101.
- Fitzsimonds, R. M., Song, H. J., & Poo, M. M. (1997). Propagation of activity-dependent synaptic depression in simple neural networks. *Nature*, 388(6641), 439–448.

- Fortz, B., & Thorup, M. (2000). Internet traffic engineering by optimizing OSPF weights. In *Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies* (vol. 2, pp. 519–528). Piscataway, NJ: IEEE.
- Froemke, R. C., Debanne, D., & Bi, G. Q. (2010). Temporal modulation of spike-timing-dependent plasticity. *Front. Synaptic Neurosci.*, 2, 19.
- Gavoille, C. (2001). Routing in distributed networks: Overview and open problems. *SIGACT News*, 32(1), 36–52.
- Gayme, D., & Topcu, U. (2011). Optimal power flow with distributed energy storage dynamics. In *Proceedings of the 2011 American Control Conference* (pp. 1536–1542). Piscataway, NJ: IEEE.
- Ha, S., Rhee, I., & Xu, L. (2008). Cubic: A new TCP-friendly high-speed TCP variant. *SIGOPS Oper. Syst. Rev.*, 42(5), 64–74.
- Isa, N., Mohamed, A., & Yusoff, M. (2015). Implementation of dynamic traffic routing for traffic congestion: A review. In M. Berry, A. Mohamed, & B. W. Yap (Eds.), *Soft computing in data science*. New York: Springer-Verlag.
- Jin, S., Guo, L., Matta, I., & Bestavros, A. (2002). *A spectrum of TCP-friendly window-based congestion control algorithms* (Technical Report, Computer Science Department). Boston: Boston University.
- Jose, L., Yan, L., Alizadeh, M., McKeown, G. V. N., & Katti, S. (2015). High speed networks need proactive congestion control. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks* (pp. 14:1–14:7). New York: ACM.
- Kelly, T. (2003). Scalable TCP: Improving performance in highspeed wide area networks. *SIGCOMM Comput. Commun. Rev.*, 33(2), 83–91.
- Kepecs, A., van Rossum, M. C., Song, S., & Tegner, J. (2002). Spike-timing-dependent plasticity: Common themes and divergent vistas. *Biol. Cybern.*, 87(5–6), 446–458.
- Kopec, C. D., Li, B., Wei, W., Boehm, J., & Malinow, R. (2006). Glutamate receptor exocytosis and spine enlargement during chemically induced long-term potentiation. *J. Neurosci.*, 26(7), 2000–2009.
- Kramar, E. A., Babayan, A. H., Gavin, C. F., Cox, C. D., Jafari, M., Gall, C. M., . . . Lynch, G. (2012). Synaptic evidence for the efficacy of spaced learning. *Proc. Natl. Acad. Sci. U.S.A.*, 109(13), 5121–5126.
- Lahiri, S., & Ganguli, S. (2013). A memory frontier for complex synapses. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 26 (pp. 1034–1042). Red Hook, NY: Curran.
- Lazar, A. (1983). Optimal flow control of a class of queueing networks in equilibrium. *IEEE Transactions on Automatic Control*, 28(11), 1001–1007.
- Li, F., Wu, B., Xu, L., Shi, C., & Shi, J. (2014). A fast distributed stochastic gradient descent algorithm for matrix factorization. In *Proc. of the Intl. Workshop on Big Data, Streams and Heterogeneous Source Mining* (pp. 77–87).
- Markram, H., Gerstner, W., & Sjostrom, P. J. (2012). Spike-timing-dependent plasticity: A comprehensive overview. *Front. Synaptic Neurosci.*, 4, 2.
- Mateos-Nunez, D., & Cortes, J. (2013). Noise-to-state exponentially stable distributed convex optimization on weight-balanced digraphs. In *Proceedings of the IEEE Conf. on Decision and Control* (pp. 2781–2786). Piscataway, NJ: IEEE.
- Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biol. Cybern.*, 98(6), 459–478.

- Navlakha, S., Barth, A. L., & Bar-Joseph, Z. (2015). Decreasing-rate pruning optimizes the construction of efficient and robust distributed networks. *PLoS Comput. Biol.*, *11*(7), e1004347.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*(3), 267–273.
- Osogami, T., & Otsuka, M. (2015). Seven neurons memorizing sequences of alphabetical images via spike-timing dependent plasticity. *Sci. Rep.*, *5*, 14149.
- Pande, P. P., Grecu, C., Jones, M., Ivanov, A., & Saleh, R. (2005). Performance evaluation and design trade-offs for network-on-chip interconnect architectures. *IEEE Trans. Comput.*, *54*(8), 1025–1040.
- Royer, E. M., & Toh, C.-K. (1999). A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications*, *6*(2), 46–55.
- Shakkottai, S., & Srikant, R. (2007). Network optimization and control. *Found. Trends Netw.*, *2*(3), 271–379.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.*, *3*(9), 919–926.
- van Rossum, M. C., Bi, G. Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.*, *20*(23), 8812–8821.
- Vogels, T. P., & Abbott, L. F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat. Neurosci.*, *12*(4), 483–491.
- Wang, J., Li, L., Low, S. H., & Doyle, J. C. (2003). Can shortest-path routing and TCP maximize utility? In *Proceedings of the Twenty-Second Annual Joint Conference of the IEEE Computer and Communications* (vol. 3, pp. 2049–2056). Piscataway, NJ: IEEE.
- Watt, A. J., & Desai, N. S. (2010). Homeostatic plasticity and STDP: Keeping a neuron's cool in a fluctuating world. *Front. Synaptic Neurosci.*, *2*, 5.
- Yang, Y., & Calakos, N. (2013). Presynaptic long-term plasticity. *Front. Synaptic Neurosci.*, *5*, 8.
- Yuan, K., Ling, Q., & Yin, W. (2016). On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, *26*(3), 1835–1854.
- Zanutto, B. S., & Staddon, J. E. (2007). Bang-bang control of feeding: Role of hypothalamic and satiety signals. *PLoS Comput. Biol.*, *3*(5), e97.
- Zhou, Q., Homma, K. J., & Poo, M. M. (2004). Shrinkage of dendritic spines associated with long-term depression of hippocampal synapses. *Neuron*, *44*(5), 749–757.
- Zinkevich, M., Weimer, M., Li, L., & Smola, A. (2010). Parallelized stochastic gradient descent. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, & A. Culotta (Eds.), *Advances in neural information processing systems*, *23* (pp. 2595–2603). Red Hook, NY: Curran.