

Design, Analysis and Performances of Chemical-Inspired Rate Controllers in Packet Networks

Massimo Monti, Manolis Sifalakis

Technical Report CS-2013-001

University of Basel

January, 2013

Abstract

In computer networks, a Distributed Rate Controller (DRC) must quickly propagate changes of inflow rates and let participating sites converge to their admissible rate. In this paper we introduce a family of DRCs where controllers can be easily customized whereas their performance and dynamics are strictly predictable. Borrowing from engineering methods in Chemistry, we show how to derive a deterministic mathematical model of the network flow that can be analyzed through standard tools of (linear) system theory. We also report on simulation and native experimental results that validate our theoretical approach.

Keywords: Distributed Rate Control, Chemical Networking, Token Bucket..



Introduction

Until today the Internet is a resource-interconnecting and access-sharing facility. Access sharing is enabled through the statistical multiplexing of flows. In this process, rate control is the main intrinsic function both for regulating the relative network utilization by the admitted flows, and for guaranteeing admissibility for all flows that demand access. Capacity allocations and resource reservation, congestion control and avoidance, DoS mitigation, service differentiation, and traffic shaping are examples of embedded rate control functions serving as a common means to different ends.

Given the prominent role of this function in network operations, we have developed a systematic way of designing, analyzing, and deploying rate controlling functions, which in its methodology resembles the engineering of electronic circuits, or playing with LEGO[®] mind-storm kits.

The design methodology borrows a paradigm and a structural representation from chemical engineering. The ease whereby a mathematical model can be thereafter generated allows (a) the direct derivation of a flow model and (b) the application of control-theory to study and formally verify their dynamic behavior, prior to deployment. So far, as we show in this paper, this approach has proven complete enough to alleviate the need for microscopic queue-theoretical modeling, which is cumbersome and largely bases on assumptions of the traffic patterns.

The resulting family of rate controllers designed in this way can be easily parametrized, extended, and customized, for various purposes and in different operational environments; for example (a) at the client side to enable service differentiation among user's flows, (b) among clients for a distributed coordination between aggregate flows of nodes, (c) in the network to manage the allocation of capacities to admitted flows, and (d) at the server side to regulate the access to a resource.

Related work

Rate control is a core function in the operation and engineering of computer networks. The topic has therefore received, and continues to receive, a lot of attention generating voluminous literature. However, we concentrate and discuss only a few works that bear close similarity to ours and we identify the key differences that make our work novel.

In [1], Raghavan *et al.* proposed a distributed rate-limiter for regulating access to cloud resources. In their Flow Proportional Sharing (FPS) approach, distributed instances of token-bucket rate-allocators exchange statistics and “shift” capacity slices to different flows. Although our work is not confined to cloud settings and applications only, the resulting characteristic behavior of our chemical rate controller comes very close to that of FPS. However, effective FPS operation relies on discrete rate estimates peeked from congestion aware flows making it more tailored for TCP traffic (and similar variates). The chemical rate controller is agnostic to intra-protocol rate estimators, and therefore suitable for usage with any type of flows. In addition the generic and formally analyzed design makes it more predictable and parametrizable for different applications and different objectives (flow cooperation or flow competition).

In order to address challenges on the opposite side of the cloud, i.e. controlled sharing and virtualization of resources within the cloud, Abu-Libdeh *et al.* in [2] proposed a rate control scheme (Ajil) for traffic differentiation across multicast channels. The system uses an out-of-band feedback channel for monitoring the utilization of the rate controlled resources, and a reactor component computes for each transmitter its estimated rate share and reinforces the rate control policy. Despite the conceptual similarity of the building blocks, Ajil applies coarse TCP-like rate-control policies (e.g. enforcing every TCP flow twice independently the same congestion avoidance loop!), depends on synchronized clocks between resource users, and the flow

rate estimates are discretely computed at periodic intervals. Our chemical rate controller does not suffer from these artifacts and operates smoother in continuous time.

On the analysis methodology, our work is very similar to a historical proposal by Keshav [3] that proposed the design of a flow control mechanism based on a control-theoretic approach, and the more recent work by Huibing *et al.* [4] that attempt to formalize the analysis of TCP-like flow control schemes. By contrast to both the work by Keshav [3] and Huibing [4], in our case we do not incrementally build the mathematical description of the system (and do not have to resort on approximating assumptions). Instead, the mathematical description of the system draws directly from its representation as a chemical reaction system. The model is the system itself and therefore accurate: the reaction model dictates the system behavior and is the direct estimate of its state at the same time.

Deriving formally the control strategy, by starting from a certain model, is common to chemistry (the Metabolic Control Analysis [5, 6]) as well as to computer science. The main difficulty that researchers of both areas have to cope with is finding a correct flow model that is able to capture the macroscopic dynamics of the analyzed system [7, 8]. The added value of using the chemical paradigm to design algorithms for networking is that the flow model itself represents the executable model of the algorithm.

Past and recent works such as [9, 10, 11, 12] have shown the benefit in using control theory to design Active Queue Management (AQM) controllers, where the algorithm deployed in a server, having as a feedback the bottleneck-link queue level, computes the marking probability to drop the packets of the many TCP sessions competing for the link. Starting from the flow model, which approximates the behavior of the system that has to be controlled, they mathematically derive the transfer function of the control module and then find a way to implement the controller and to deploy it. Although we are inspired by analytical tools and techniques from Biology and Chemistry, we end up using a similar control-theoretical approach to that in [10]-[12]. However, our contribution differs from ones of [10, 11, 12] in both detailed and general terms: We propose a distributed solution that deals with rate-based signals and does not need a symbolic exchange of information (packets) but rather encodes the required information into rates. We build up the algorithm basing on simple patterns, rules, and concepts, though maintaining a theoretic methodology to justify design choices and calibrations. The chemical paradigm lets us work both intuitively and formally, and gives us the way to easily switch from implementation to modeling and vice-versa.

In order to act upon establishing fair use of the network, several proposals in recent and past literature [13, 14, 15, 16, 17, 18] advocate the use of network assisted signaling towards the end user, instead of taking rate control actions in the network (non-work conserving scheduling). We are aligned with this kind of thinking since actions for regulating the capacity sharing are made at the user-systems and we believe that our system is more suitable for this purpose than either of the above measures. This is because through our system the share-holders collectively regulate their relative utilization of the capacity rather than expecting from the network to act intrusively [16] or to compute and distribute individual capacity shares [15, 17].

Paper outline

In the remainder of this paper we start slowly setting the stage by revising the basic operation of a token-bucket rate-control scheme in Sect. 1, and establishing the analogy to chemical reactions in Sect. 2. We then start extending this basic analogy to introduce dynamism and feedback control, arriving at our core design of an adaptive chemical rate controller (Sect. 3). Thereafter, we present the analysis methodology and parametrization of the system in Sect. 4. Finally we evaluate our design both in simulation and with a real deployment in Sect. 5. We conclude this paper with a discussion of the insights developed through this work in Sect. 6, and our future orientation in Sect. 7.

1 Token-Bucket Based Rate Control

A *token bucket* or *leaky bucket* scheme (see Fig. 1) has probably been the most common building block in the design of more or less sophisticated (and protocol-independent) rate control for various purposes both in reservation-based and reservation-less networks. In this section, we briefly refresh the reader's memory on the principal operation of a token bucket scheme, so that in the next section we can introduce a chemical reaction as a valid metaphor for it.

In the simplest case (rate limiter), a token bucket is used to regulate the departure process of packets from a queue; tokens essentially authorize the consumption of a respective number of packets or bytes. The bucket is filled with tokens at a constant rate, with excess tokens being tail-dropped. In this arrangement, at any moment in time, the output of packets from the queue cannot exceed a burst size equal to the size of the bucket, and cannot sustain a constant rate higher than the refill rate of tokens.

In a more dynamic arrangement (rate controller), the rate limit imposed by the token bucket may vary in time. This can be achieved by adapting the refill rate of tokens in response to feedback from the network, or to the differentiation between departure processes of two or more regulated queues. To analyze the behavior of such a system, without loss of generality one may regard the token bucket as yet another queue. The most interesting aspect of such an arrangement is that the departure process depends on the state of both queues: The process blocks if either the packet or the token queue is empty.

2 Rate Control using Chemical Reactions

The author of [19] introduced the chemical reaction as a metaphor for modeling information processing of packet streams and thereby designing network protocols. Leaving aside a complete exploration of this metaphor, in this section we concentrate on the aspects that enable us to describe through this model the operation of a simple fixed rate token-bucket controller. This will serve as an entry point to the design of our adaptive rate controller family in the next section.

Consider the simple reaction $S + E \xrightarrow{k_1} ES$ schematically illustrated in Fig. 1. An S-molecule reacts with an E-molecule to produce an ES-molecule. By associating packets to virtual molecules, we can refer to the fill-level of a queue as a molecular concentration. For example, the concentration c_S and c_E of reactants can be seen as the number of queued packets and available tokens respectively. If one of the two reactant concentrations is depleted, the reaction becomes *inert*. This condition essentially captures the dependency between token and packet availability for scheduling transmissions out of a token-bucket regulated queue.

The reaction between the reactant species takes place at an average rate v . If we choose the inter-reaction times to be drawn from an exponential distribution, then this representation models the interaction of a $\cdot/M/1$ queue with the token bucket (subject to token availability).

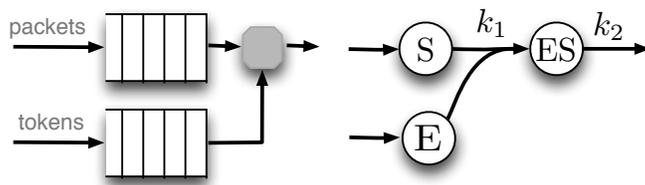


Figure 1: Token-Bucket scheme and simple equivalent chemical reaction.

Practically in this model, the reaction merely defines a rule (dependency) that relates the departure process of two (or more) queues. Also, the rate of the reaction describes the probability distribution of this rule being successfully applied.

2.1 Law of Mass Action Scheduling

At this point, to be consistent with the chemical model in [19], we need to introduce the laws that underpin the dynamic operation of a reaction network. The algorithm has to decide how the reactions in the network are scheduled. A work-conserving round-robin approach seems at first plausible for deterministic fairness or a random selection approach for statistical fairness. However on closer consideration, a non-work-conserving, fill-level weighted scheduling approach is actually more appropriate for traffic shaping purposes. This is because ideally the regulation process must be a response towards balancing the pressures applied to the system, which can only be achieved by afflicting packets of certain queues with a small additional delay. This will be understood more intuitively in the next section, where our design needs to make rate controlling decisions that balance multiple flows.

Our non-work-conserving scheduling algorithm that implicitly assigns weights to reactions bases on the chemical *Law of Mass Action* (LoMA):

L1: *The reaction rate is proportional to the concentration (quantity) c of all reactant molecules.*

LoMA essentially relates the reaction rates with the reactant concentrations through the equation

$$v_r = k_r \times \prod_{i=1}^{|\mathcal{S}|} c_i^{\alpha_i}$$

where c_i is the concentration of the i -th species, k_r is a “speed” coefficient for reaction r , $|\mathcal{S}|$ is the set of all molecule species, and α_i is the number of molecules from species i that are consumed by reaction r . For example for reaction $S + E \xrightarrow{k_1} ES$ in Fig. 1 its rate based on the LoMA is given by $v_{r_1} = k_1 \cdot c_S^1 \cdot c_E^1$. (Further details on the LoMA as a non-work conserving scheduling discipline can be found in [20].)

A direct consequence of this fill-level-proportional scheduling algorithm is the *flow conservation* principle:

L2: *At equilibrium state, for each molecular concentration, the total production rate equals the total consumption rate.*

Another consequence is that of mass conservation in a chemical reaction along closed loops (*moieties* in chemistry):

L3: *The total sum of molecule concentrations along a loop remains constant, if (a) the total number of molecules consumed by reactions along the loop is equal to the total number of molecules produced, and if (b) all concentrations along the loop are altered only by reactions involved in this or another loop.*

2.2 Enzymatic rate control

We now have to determine the rate at which the token bucket is filled in the chemical model. If we assume that the concentration of E-molecules is the one that represents the number of tokens in the bucket, we care that this concentration is refreshed at a fixed rate until a maximum size. A simple possibility is to externally inject new E-molecules at a certain rate, and constantly monitor the concentration size such that it does not exceed the set limit (which equals the bucket length). However, an elegant design pattern (or *motif* according to [19]) comes directly from chemistry to capture this situation with a control loop instead. This enzymatic reaction pattern, illustrated in Fig. 2, introduces a second reaction, $ES \xrightarrow{k_2} E$.

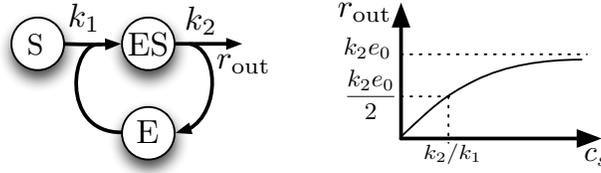


Figure 2: A chemical fixed-rate token-bucket.

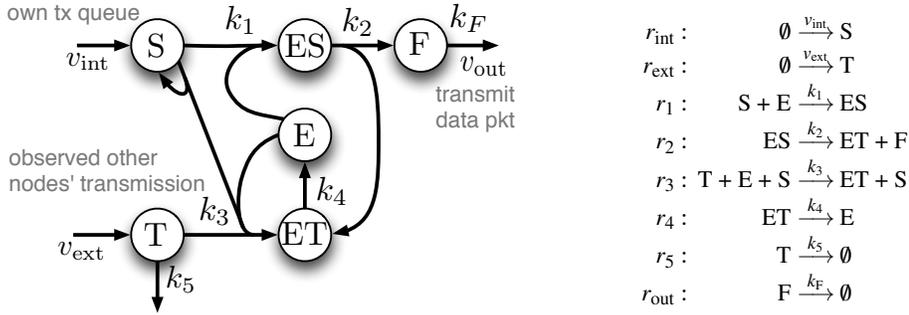


Figure 3: Adaptive chemical rate controller.

The combination of LoMA and the two conservation laws above is sufficient to provide us with a direct description of the system in Fig. 2 through the following equation, which relates the output rate to the control parameters and to the fill level of the input queue (c_S concentration):

$$v_{\text{out}} = k_2 e_0 \frac{c_S}{c_S + k_1/k_2} \quad (1)$$

where $c_E + c_{ES} = \text{const.} = e_0$ according to the mass conservation law (**L3**). The product $k_2 e_0$ defines the upper rate limit v_{max} that can be achieved, whereas the ratio of control constants (k_2/k_1) specifies how aggressive the rate regulation is.

3 Adaptive Chemical Rate Controller Design

We are now ready to describe the reaction network that provides a core design of a distributed adaptive rate controller and provide an analysis of its behavior. By parametrizing this design or extending its basic set of reactions, a family of rate controllers can be derived.

Figure 3 presents the complete reaction network implementing the functionality of the adaptive rate controller. Despite the perceived complexity, the underlying structure is a simple combination of a number of chemical design patterns that extend the fixed-rate chemical token bucket of Fig. 2. Exploring part-by-part this reaction network, we first identify two inter-weaved reaction loops, one of which is encapsulated in (part of) the other. This pattern is highlighted in Fig. 4(a). We will denote the sum of concentrations along the larger loop as $E_0 = c_{ES} + c_E + c_{ET} = \text{const.}$ according to **L3**, and the sum of concentrations along the smaller loop as $e_0 = c_E + c_{ET}$. Note that according to **L3**, e_0 is not constant because E is consumed by two reactions (r_1 and r_3) but only produced by one (r_4); also ET is consumed by only one reaction (r_4) but produced by two (r_2 and r_3). The difference of the two concentrations can be expressed as a function of the ratio of rates of the two reactions that consume E-molecules.

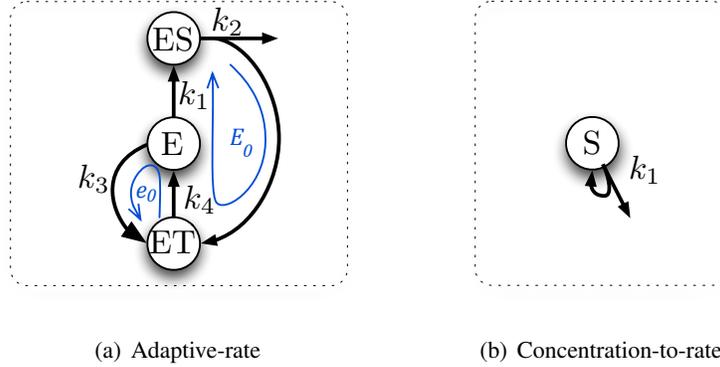


Figure 4: Chemical motifs for the distributed, adaptive rate controller.

When the system is in equilibrium all rates are stationary. The rate-conservation law **(L2)** applied to species E and ES respectively gives us

$$k_4 c_{ET} = k_1 c_E + k_3 c_E \quad (2a)$$

$$k_2 c_{ES} = k_1 c_E \quad (2b)$$

Furthermore, the concentrations along the two embedded loops are related with the equality $E_0 - e_0 = c_{ES}$, which combined with system (2) (and eliminating c_{ES} and c_{ET}) gives us

$$E_0 - e_0 = \frac{k_1 + k_3}{k_2} \cdot \frac{c_E}{v_3/v_1 + 1}$$

where $v_3 = k_3 c_E$ and $v_1 = k_1 c_E$. Given set coefficients k_i and on-average constant concentration c_E (the system being in equilibrium), this equation shows the relationship between the ratio of rates v_3/v_1 (that E molecules get consumed) and the difference of mass $E_0 - e_0$ in the two loops. This difference determines the number of tokens available for dispatching packets from the input queue S.

Thus, by replacing the single conservation loop of Fig. 2 with the extended two-loop pattern of Fig. 4(a) (and replacing e_0 with $E_0 - e_0$ in (1)), the output rate v_{out} becomes dependent on the ratio of outflow rates from E-species. By adjusting this ratio we can adapt the maximum rate limit $v'_{\text{max}} \leq v_{\text{max}}$, where $v_{\text{max}} = k_2 E_0$ and $v'_{\text{max}} = k_2 (E_0 - e_0)$.

In order to regulate rate v_3 that controls the “speed” of the small reaction loop, reaction r_3 depends on two additional inputs. One of them is the backlog of S-molecules, which builds up when reaction r_1 cannot process them fast enough. This represents the amount of locally queued packets. The other is the concentration c_T , which increases at a rate that reflects the “sensed” aggregate utilization of the network of other transmitters. (How the sensing function is implemented is beyond the scope of this design. Depending on the application environment different solutions may be adopted from the literature (e.g., carrier sensing, multicast, gossip, explicit out-of-band signaling, and others).)

More explicitly, concentration c_S participates in reaction r_3 with the so-called *concentration-to-rate* pattern [19] (Fig. 4(b)). The looping arrow signifies that for every molecule consumed from concentration c_S in a reaction, a new S-molecule is produced. Thus c_S remains constant, and because of the LoMA, the reaction rate is directly proportional to c_S .

The effect of this pattern is that when the arrival rate v_{int} at S is in flow-balance with rate v_1 of reaction r_1 , the input queue is drained fast enough and there is no increasing backlog of S-molecules to “fuel” reaction r_3 . The small loop is blocked and the system behaves like a fixed-rate token-bucket.

As soon as the enqueueing rate v_{int} increases too much and there are other flows claiming part of the network capacity, there will be an increasing backlog of S-molecules created that activate reaction r_3 , with a rate proportional to the backlog. As there are contributor species to reaction r_3 , the rate increase will be further regulated according to the utilization of the network by other users through the concentration c_T . The result will be an adaptive drop in the maximum allowed rate v'_{max} .

The last component of the design is the F-species, which implements (in a simple queue) a linear low-pass output filter. When the system is in equilibrium, the instantaneous variation of c_F will follow the variations of the influx rate v_2 , and because of the LoMA the transmission rate will follow: $v_{\text{out}} = k_F c_F$. This means that k_F determines the cut-off frequency for smoothing (dampening) the oscillations of the transmission rate acceleration. This is important for several reasons: (a) it reduces the burstiness of transmissions and their effects on the sensing function of other rate controllers, (b) in case of long delay networks, it “smoothens” the traffic variations over time such that rate controllers respond to them as they develop and not after they have occurred.

4 Analysis of the Chemical Rate Controller

We now analyze the final design of our rate controller in order to develop a clear view of its dynamics and understand which are the key parameters that dictate its behavior. The following analysis is divided in steady-state and transient-state.

4.1 Steady-state analysis

Reaction r_5 does not play a role in the principal operation of the rate controller but rather aims to ensure that the system does not end up with stale state (more in Sect. 4.3). For this reason, and in order to keep the analysis uncluttered we choose to ignore it for now.

The following system of ordinary differential equations (ODEs) describes, based on LoMA, the dynamic behavior of the chemical reaction system.

$$\dot{c}_S = v_{\text{int}} - k_1 c_S c_E \quad (3a)$$

$$\dot{c}_{ES} = k_1 c_S c_E - k_2 c_{ES} \quad (3b)$$

$$\dot{c}_E = k_4 c_{ET} - k_3 c_S c_{ECT} - k_1 c_S c_E \quad (3c)$$

$$\dot{c}_T = v_{\text{ext}} - k_3 c_S c_{ECT} \quad (3d)$$

$$\dot{c}_{ET} = k_3 c_S c_{ECT} + k_2 c_{ES} - k_4 c_{ET} \quad (3e)$$

$$\dot{c}_F = k_2 c_{ES} - \underbrace{k_F c_F}_{v_{\text{out}}} \quad (3f)$$

where $E_0 = c_{ES} + c_E + c_{ET} = \text{const.}$ according to **L3**.

When the system reaches equilibrium (steady) state the concentrations remain on average constant and so by solving the system (3) for zero derivatives, we obtain the concentrations in terms of the k_i parameters, and

the system inputs

$$c_S^* = \frac{v_{\text{int}}k_2k_4}{e_0k_1k_2k_4 - v_{\text{int}}(k_2 + k_4)k_1 - v_{\text{ext}}k_1k_2} \quad (4a)$$

$$c_{\text{ES}}^* = \frac{v_{\text{int}}}{k_2} \quad (4b)$$

$$c_E^* = \frac{e_0k_2k_4 - (v_{\text{ext}} + v_{\text{int}})k_2 - v_{\text{int}}k_4}{k_2k_4} \quad (4c)$$

$$c_T^* = \frac{v_{\text{ext}}k_1}{v_{\text{int}}k_3} \quad (4d)$$

$$c_{\text{ET}}^* = \frac{v_{\text{ext}} + v_{\text{int}}}{k_4} \quad (4e)$$

$$c_F^* = \frac{v_{\text{int}}}{k_F} \quad (4f)$$

System (4) provides some useful insights about the system operation: As concentrations cannot assume negative values, from (4c) we get the following constraint in the steady-state:

$$v_{\text{int}} \leq \frac{e_0k_4 - v_{\text{ext}}}{1 + k_4/k_2} \quad (5)$$

This means that as long as the locally generated traffic remains below this limit value there are E molecules available and there is no rate-limitation imposed. The output rate of the system $v_{\text{out}} = c_F k_F$ follows on average the input rate v_{int} according to (4f). As soon as the equality in (5) is met (i.e. v_{int} becomes higher than the rate of reaction r_1), the backlog of S-molecules increases, which activates reaction r_3 . The concentration of available E-molecules converted to ES-molecules drops, forcing reaction r_1 to slow down further. Unless the sender backs-off, the increasing reaction speed of r_3 will absorb more molecules in e_0 and tend to block transmissions.

On the other hand, a complete block is unlikely (subject to suitable selection of k_i parameters) because as the reaction speed of r_3 accelerates, the concentration of T molecules starts reducing faster, counter-balancing the acceleration rate of r_3 . This gradual reduction of the acceleration will continue as c_T^* approaches, according to (4d) the limit value

$$c_T^* = \frac{k_1(k_2 + k_4)}{k_3k_2(k_4e_0/v_{\text{ext}} - 1)} \quad (6)$$

Thereafter, r_3 will start decelerating as $c_T^* \rightarrow 0$, until eventually r_1 gets re-activated, and a smoothly oscillating pattern emerges as long as the system remains under ‘‘pressure’’ (above the set-point (5)) but never exceeds the maximum rate reinforced by the E_0 loop.

Furthermore, we can express the steady-state concentrations of all other molecule species in terms of the input concentrations c_S and c_T by solving system (4) so as to eliminate v_{ext} and v_{int} . Of particular interest is the concentration c_F^* that, because of the LoMA, is directly proportional to the output transmission rate:

$$c_F^* = \frac{k_1k_2k_4e_0c_S}{(k_2k_4 + c_S(k_1(k_2 + k_4) + c_T^*k_2k_3))k_F} \quad (7)$$

At the boundary condition, when the input rate is $v_{\text{int}} = 0$, there is no backlog ($c_S = 0$) and consequently the output rate is $v_{\text{out}} = k_F c_F^* \stackrel{(7)}{=} 0$, too. When on the other hand, the local data generation rate v_{int} is extremely high and rate control takes place, c_S would theoretically increase without bounds. However, the output rate

cannot exceed an asymptotic upper bound, which depends on the current average utilization of the resource by all other transmitters

$$c_S \longrightarrow \infty : \quad v_{\text{out}} = k_F c_F^* \stackrel{(7),(6)}{=} \frac{k_4 e_0 - v_{\text{ext}}}{1 + k_4/k_2} \quad (8)$$

If we further assume that the local transmitter uses the resource alone, i.e., $v_{\text{ext}} = 0$, then from (8) we can infer a set of parameters that allow exploiting the (asymptotic) maximum available capacity of the resource: $k_2 \gg k_4$ and $k_4 e_0 = v_{\text{max}}$. The blocking case of $k_4 e_0 = v_{\text{ext}}$ is theoretically possible only if there are non rate-controlled transmitters, or there is an infinite number of (chemically) rate-controlled transmitters.

4.2 Analysis of transient behavior

By solving numerically the system (3), we could predict the transient behavior of the protocol in the time-domain for a specific variation of the input-rates. However, this does not allow us to evaluate how the rate-controller responds to arbitrary variations (perturbations) of the input rates. To develop a general model of the transient behavior, we therefore start from a different perspective – the frequency domain.

In general our system is non-linear as can be seen by the ODEs and also by the fact that its output is bounded. However, as it cannot move arbitrarily away from the output region $[0, v_{\text{max}}]$, true instability caused by the system operating far away from its steady-state is not a concern. We can therefore focus our analysis around the steady state where we can see it as a linear system. The system is also causal and time homogeneous because its k_i parameters remain the same and when its inputs go to zero the concentrations tend to reset towards its vanilla state (more on this in Sect. 4.3).

These conditions enable us to follow a systematic control-theoretical analysis that consists of the following steps: (1) express the system of ODEs in a linear form (linearization), (2) treat the system as a linear time invariant (LTI) one, and express the observable parameters as a function of its state, and (3) produce its transfer function (TF) in state space using the Laplace transform. For an LTI system the TF completely describes the system and allows us to make general observations.

4.2.1 Linearization around the steady-state

We can represent the system of ODEs (3) compactly in matrix form as follows

$$\dot{\mathbf{c}}(t) = \Psi \cdot \mathbf{v}(\mathbf{c}(t), \mathbf{p}(t)) \quad (9)$$

where $\mathbf{v}(t)$ is a column vector containing all reaction rates as functions of molecular species concentrations, $\mathbf{c}(t)$, and the input rates $\mathbf{p}(t) = [v_{\text{int}} \ v_{\text{ext}}]^T$. Ψ is the so-called *stoichiometric matrix* that describes the reaction network topology, i.e., which reaction consumes/produces which molecules and how many of them; its rows correspond to molecular species, and its columns to the reactions, including the two system inputs.

For small fluctuations around the steady-state, we define $\mathbf{x}(t) = \mathbf{c}(t) - \mathbf{c}^*$, and $\mathbf{u}(t) = \mathbf{p}(t) - \mathbf{p}^*$. By substitution in (9), the system of ODEs is reduced to a linear form

$$\dot{\mathbf{x}}(t) = \Psi \cdot \left. \frac{\partial \mathbf{v}}{\partial \mathbf{c}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)} \cdot \mathbf{x}(t) + \Psi \cdot \left. \frac{\partial \mathbf{v}}{\partial \mathbf{p}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)} \cdot \mathbf{u}(t) \quad (10)$$

4.2.2 Linear Time Invariant Representation

The state space description of an LTI system is given by the linear system

$$\dot{\mathbf{x}}(t) = \mathbf{A} \cdot \mathbf{x}(t) + \mathbf{B} \cdot \mathbf{u}(t) \quad (11a)$$

$$\mathbf{y}(t) = \mathbf{C} \cdot \mathbf{x}(t) + \mathbf{D} \cdot \mathbf{u}(t) \quad (11b)$$

where \mathbf{y} represents the observable quantities of interest (in our case the variation of reaction/output rates). State matrices \mathbf{A} and \mathbf{B} express how changes in the internal state and external inputs respectively affect the future state of the system. From (10) we evaluate \mathbf{A} and \mathbf{B} at the fixed point as

$$\mathbf{A} = \Psi \cdot \left. \frac{\partial \mathbf{v}}{\partial \mathbf{c}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)} \quad \mathbf{B} = \Psi \cdot \left. \frac{\partial \mathbf{v}}{\partial \mathbf{p}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)}$$

The output matrix \mathbf{C} and the feedforward matrix \mathbf{D} express the observed parameters (rates) in terms of the system's state and the inputs respectively. Thus

$$\mathbf{C} = \left. \frac{\partial \mathbf{v}}{\partial \mathbf{c}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)} \quad \mathbf{D} = \left. \frac{\partial \mathbf{v}}{\partial \mathbf{p}} \right|_{(\mathbf{c}^*, \mathbf{p}^*)}$$

4.2.3 Transfer function in the state-space

We are now in the position to produce the system's TF in the frequency domain by means of the Laplace transform. Under zero variance of the input rates, the TF is $\mathbf{H}(s) = \mathbf{y}(s)/\mathbf{u}(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B} + \mathbf{D}$. Here, we do not present the entire TF matrix $\mathbf{H}(s)$ but instead we focus only on the row that corresponds to reaction r_{out} , which gives the system's output rate as a function of the input rates:

$$H_{\text{out-int}}(s) = \frac{v_{\text{int}}}{c_{\text{S}}^*} \cdot \frac{k_2 k_{\text{F}} (s + \frac{k_3}{k_1} v_{\text{int}}) (s + k_4)}{(s + k_{\text{F}}) \cdot \sigma(s)} \quad (12a)$$

$$H_{\text{out-ext}}(s) = \frac{v_{\text{ext}} v_{\text{int}}}{c_{\text{E}}^* c_{\text{T}}^*} \cdot \frac{k_2 k_{\text{F}} s}{(s + k_{\text{F}}) \cdot \sigma(s)} \quad (12b)$$

where,

$$\begin{aligned} \sigma(s) &= s^4 + \left(k_2 + k_4 + \beta + k_3 (v_{\text{int}} + c_{\text{T}}^*) \right) s^3 + \\ &+ \left(\beta (k_4 + k_2 + k_3 v_{\text{int}}) + k_2 k_3 (v_{\text{int}} / k_1 + c_{\text{T}}^*) \right) s^2 + \\ &+ \left(\frac{k_2 k_4}{k_1} (c_{\text{E}}^* + k_3 v_{\text{int}}) + \beta k_2 v_{\text{int}} (k_3 + k_4) \right) s + \\ &+ k_2 k_3 k_4 v_{\text{int}} c_{\text{E}}^* \\ \beta &= \frac{(k_2 + k_4)^2 k_1 c_{\text{S}}^*}{k_2 k_4} + \frac{k_2 k_4}{c_{\text{S}}^*} \end{aligned}$$

By looking at (12) we realize that the controller is an adaptive filter whose parameters are a function of the input rates. Specifically, it behaves as a low-pass filter (LPF) for the variations of v_{int} (with unit gain in the zeroth frequency), and as a band pass filter (BPF) for the variations of v_{ext} (with a very high attenuation for the whole spectrum). The effects of the output stage filtering (F-molecules) is confirmed in the TF equations (pole at $-k_{\text{F}}$ that increases the order of the filters).

When rate regulation takes place, as v_{int} increases beyond the bifurcation point threshold (eq. (5)), by solving (12) for $c_S \rightarrow \infty$ we get

$$H_{\text{out-int}}(s) = 0 \quad (13a)$$

$$H_{\text{out-ext}}(s) = \frac{-k_2 k_F}{(s + k_F)(k_2 + k_4 + s)} \quad (13b)$$

In this case therefore, the system becomes “deaf” to the variations of the offered load v_{int} . The variations of the sensed traffic v_{ext} on the other hand, control the actual transmission rate of the system v_{out} , through a fixed second order low-pass filter (its poles being at $-k_F$ and $-(k_2 + k_4)$).

As the TF has always negative real poles the system exhibits Bounded-Input-Bounded-Output (BIBO) stability: for input bounded variations the output transmission variations are also bounded. Moreover, if the input is not bounded as (13) suggests, the output variation is guaranteed to be bounded at zero variation at the maximum permitted transmission rate, which obsoletes the theoretical need for more complicated stability tests, such as Liapunov’s second method [21].

4.3 Setting the system parameters

As we have mentioned, coefficients k_2 , k_4 , and E_0 are key parameters for setting the upper rate limit imposed by the chemical controller: $k_2 \gg k_4$ and $k_4 E_0 = v_{\text{max}}$.

The coefficient k_F sets the cutoff frequency of the output stage low-pass filter implemented by c_F to smooth bursty behavior and reduce fast oscillations of the transmission rate. The choice of k_F therefore is a tradeoff between reaction speed and stability, and needs to be adjusted according to the maximum path delay of the feedback channel.

Coefficients k_1 and k_3 have no evident effects on the system steady states but they affect its transient state. By increasing either k_1 or k_3 compared to the other, the response time is shorter, enabling prompt state changes; yet this can be a trade off for stability in face of feedback delays.

Initial conditions (i.e., species concentrations at time $t = 0$ s) should not influence the behavior of the system. To guarantee that the system does not remember and respond to stale state after a period of inactivity, we want that it always tries to reset itself towards the vanilla state, where $c_S = c_F = c_T = c_{ES} = c_{ET} = 0$ and $c_E = E_0$. This happens automatically when $v_{\text{int}} = 0$: the c_S concentration is drained by the inertia of r_1 and r_3 , followed by the draining of c_{ES} by r_2 and in turn of c_F by r_{out} , which flushes the remaining data to the network. The unused E molecules are placed back to ET by r_3 and later to E by r_4 , given that there is no sensed traffic. Finally to avoid infinitely accumulating molecules in c_T when there are active transmissions in the network, r_5 is used to drain them. This brings the system to a graceful stand-by of zero memory.

5 Evaluation of the Chemical Rate Controller

In this section we present an evaluation of the chemical rate controller both in simulations as well as in a real deployment. For the simulation-based tests we have used the system settings shown in Tab. 1 and we have based the system operation on packet rates. The total shared rate limit v_{out} was set through the E_0 mass at 1000pkts/s.

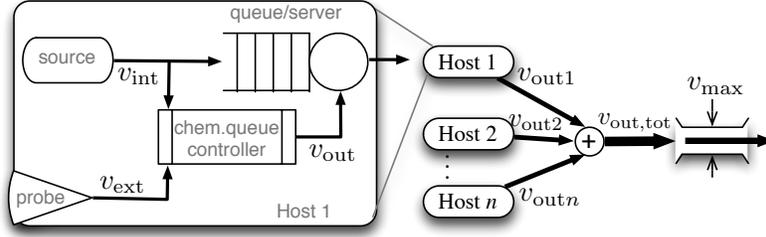


Figure 5: System deployment in an experimental scenario of N hosts sharing a bandwidth-delay product v_{out} .

$c_S(0s) = 0$	$c_E(0s) = E_0$	$k_1 = 1^1/s$	$k_4 = 1^1/s$
$c_T(0s) = 0$	$c_{ET}(0s) = 0$	$k_2 = 100^1/s$	$E_0 = 1000$
$c_{ES}(0s) = 0$	$c_F(0s) = 0$	$k_3 = 1^1/s$	$k_F = 1^1/s$

Table 1: Configuration parameters for simulation tests

For the real deployment tests we have implemented a custom queuing discipline for the Linux kernel, which operates on byte rates to account for varying size packets. Instead of replacing the existing queues in Linux' traffic control system, our software implements the chemical reaction network as an external manager module that controls the service process of the egress link queue or a certain class of traffic (see Fig. 5); concentrations in the chemical manager module are mere byte counters.

5.1 Empirical validation of system's step response

To confirm the validity of the transient analysis we configured a single chemical rate controller with the parameters of Tab. 1. We stimulated its inputs (v_{int} , v_{ext}) and observed the effects at its output v_{out} in two separate experiments. We then plotted the results against those estimated analytically through the TF.

Figure 6(a) shows the step response in the time domain for the first experiment, when initially $v_{int} = 500 \text{ pkts/s}$, $v_{ext} = 300 \text{ pkts/s}$, and then suddenly increased to $v_{int} = 550 \text{ pkts/s}$ at $t = 10 \text{ s}$ and to $v_{ext} = 350 \text{ pkts/s}$ at $t = 15 \text{ s}$. In this case the system operated at the steady state, respecting the condition (5). We observe that the increase of v_{int} at $t = 10 \text{ s}$ had an effect of increasing the output rate smoothly and fully conforming to the analytical prediction of the TF in (12a). The increase of v_{ext} at $t = 15 \text{ s}$ on the other hand had no perceived effect since the cumulative rate was still below the total rate limit (as announced by the high attenuation exhibited by the BPF described in (12b)).

In Fig. 6(b) we plot the step response in the time domain for the second experiment, when initially $v_{int} = 1400 \text{ pkts/s}$ and $v_{ext} = 600 \text{ pkts/s}$, and then suddenly increased to $v_{int} = 1450 \text{ pkts/s}$ at $t = 10 \text{ s}$ and to $v_{ext} = 650 \text{ pkts/s}$ at $t = 15 \text{ s}$. The system operated above its maximum allowed rate limit, and condition (5) was not respected. The variation of the output rate was zero (confirming (13a)) and insensitive to the change of v_{int} at $t = 10 \text{ s}$. This means the system was transmitting at a constant rate that did not exceed its allocated share. However, the variation of v_{ext} at $t = 15 \text{ s}$ had an attenuating effect (confirming (13b)) that lead to an adjustment of the maximum output rate share.

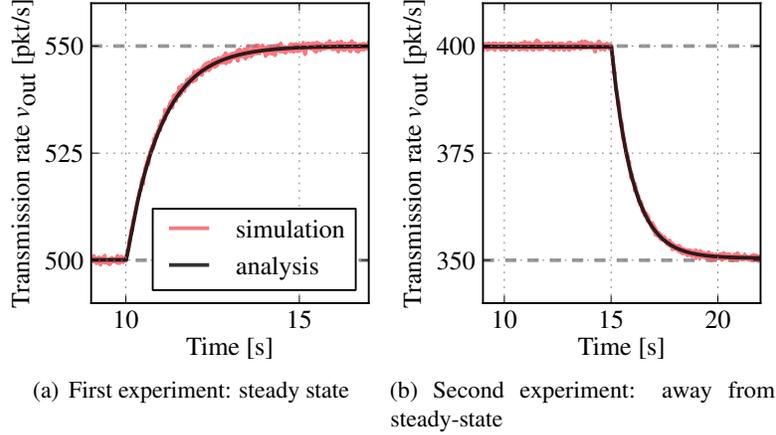


Figure 6: Empirical validation of the step response of a single chemical rate controller for maximum capacity of $v_{\max} = 1000 \text{ pkt/s}$.

- (a) Steady-state: $t = 0 \text{ s}$: $v_{\text{int}} = 500 \text{ pkt/s}$, $v_{\text{ext}} = 300 \text{ pkt/s}$; $t = 10 \text{ s}$: $v_{\text{int}} \rightarrow 550 \text{ pkt/s}$; $t = 15 \text{ s}$: $v_{\text{ext}} \rightarrow 350 \text{ pkt/s}$.
(b) Away from steady-state: $t = 0 \text{ s}$: $v_{\text{int}} = 1400 \text{ pkt/s}$, $v_{\text{ext}} = 600 \text{ pkt/s}$; $t = 10 \text{ s}$: $v_{\text{int}} \rightarrow 1450 \text{ pkt/s}$; $t = 15 \text{ s}$: $v_{\text{ext}} \rightarrow 650 \text{ pkt/s}$.

5.2 Multiple transmitters, multiple delay feedback

Next, we put into our simulation environment 11 hosts that started generating data packets with different rates at different times. The nodes shared a maximum capacity of $v_{\max} = 1000 \text{ pkts/s}$. Packets were generated in each host ($j \in \{0, \dots, 10\}$) using a Poisson model. Each host's offered load $v_{\text{int},j}$ was regulated by a chemical rate controller configured based on Tab. 1. Finally, each host sensed the utilization of the shared capacity $v_{\text{ext},j}$ with a variable delay from a uniform distribution of $(0, 60)$ ms.

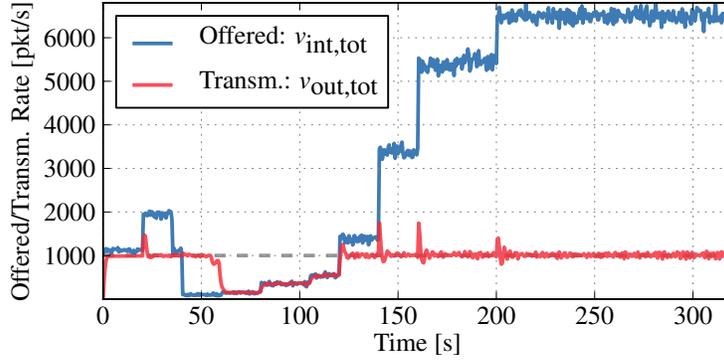
Figure 7 shows that the *capacity allocation is efficient*: The figure depicts the cumulative input and resulting output of the system, i.e., the total offered load $v_{\text{int,tot}} = \sum_i v_{\text{int},i}$ and the total utilization of the capacity $v_{\text{out,tot}} = \sum_i v_{\text{out},i}$. We observe that on average the total utilization never exceeds the configured limit: $v_{\text{out,tot}} \leq v_{\max}$.

Figure 8 shows that the *capacity allocation is fair*: The transmission rate of each node $v_{\text{out},i}$ converges to its share in response to the total offered load. The low-pass filter stage at the output of each controller effectively suppressed bursts in the offered load in order to prevent oscillations. Further experimentation with the parameter k_F showed its critical role. When the sensing delay is high, variations of k_F in the range of $(0, 1] \text{ }^1/\text{s}$ can make the system too sensitive leading to high oscillations around the rate limit, or too slow in its adaptation. It is likely that in some cases a small statistical variation of k_F in response to path variability might yield best results.

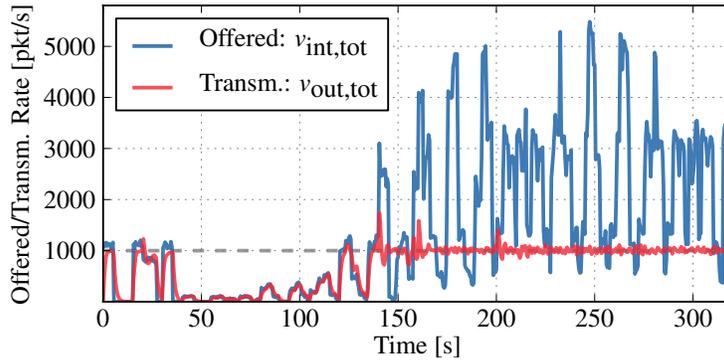
Finally, Fig. 9 shows for a few individual nodes the relation between transmission rate and system load: Host 3 who was not overloaded (as it generated packets at a very low rate) did not experience rate limitation. Host 4 who generated data slightly above its permissible transmission rate, experienced a slight rate regulation. Host 10 who was significantly overloaded, experienced major rate regulation (it even did not match its packet generation rate as it entered the network).

5.3 Real world deployment

The predicted behavior of the rate controller has been validated also by early trials in a real world setting with our chemical queueing discipline in the Linux kernel. The intent of this test was to confirm our simulation



(a) Continuous traffic generation



(b) Bursty traffic generation

Figure 7: Total offered load of 11 nodes and total utilization of the network. The rate limit of $v_{\max} = 1000 \text{ pkt/s}$ is fully respected on average.

results with more realistic capacities/speeds, and also to get a first impression of how the adaptive behavior of the chemical rate controller interacts with TCP’s control loop.

The setup involved four hosts, one on a wired network and the other three on a wireless network. The feedback was provided by an out-of-band UDP channel, and the rate controller was configured to regulate the node’s ssh/scp traffic class (destination port 22). Each node used the ssh service for remote login and transferred a large file at different time intervals and for various time durations. The maximum aggregate rate limit for all rate-controlled traffic was set to $v_{\max} = 2 \text{ Mbps}$ (via E_0), the feedback channel delay was 20 ms, and the low-pass filter coefficient was set to $k_F = 100 \text{ 1/s}$.

Figure 10 reports the results and confirms the correct operation of the system. Every time a new node accesses the service, the total capacity is divided fairly among the nodes on average: one node at 2 Mbps, two nodes at 1 Mbps per node, three nodes at 0.65 Mbps per node, and four nodes at 0.5 Mbps per node. This was also confirmed by the transfer rates reported at each scp session.

One observation is that as the number of nodes accessing the service increases, the stochasticity did so too (although on average it converged to the per-node share). A possible explanation could come from how the TCP exits from the fast-recovery phase when the rate limit is lowered (i.e., whether it succeeds to stay in congestion avoidance or if it drops to slow start and then re-enters congestion avoidance). Further future exploration is needed to establish a conclusive assessment of the situation and refinement of the controller.

A second observation is that the allocated aggregate capacity was asymptotically matched but never

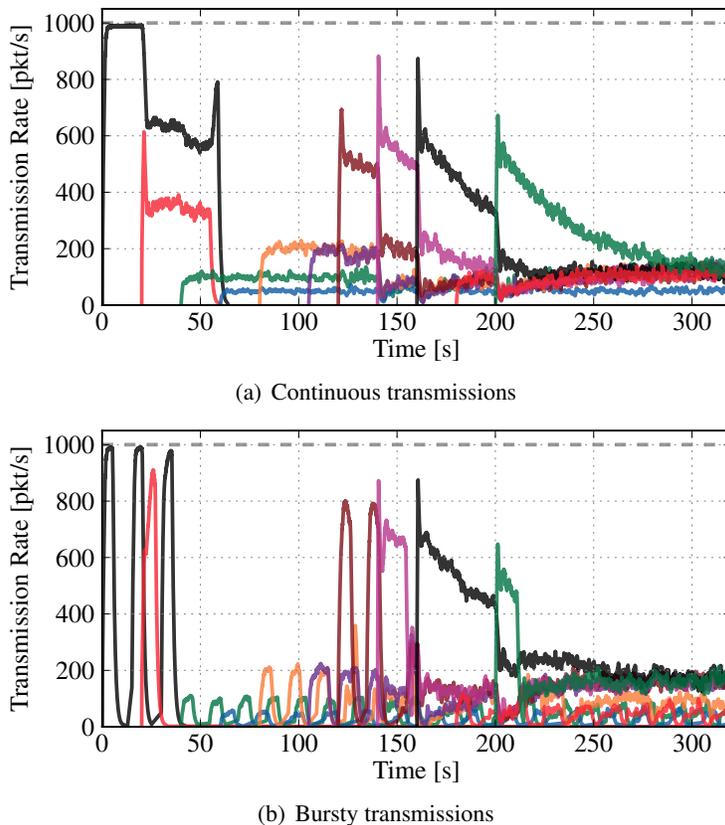


Figure 8: Individual transmission rates $v_{out,j}$ of the 11 nodes responding to each others transmission rates.

100% utilized, even in the single node periods. In fact, this observation also manifested in the simulation scenario. This is probably related to small variance of the feedback signals (expressed through c_T) and c_S that keep reaction r_3 always active even at a very low speed.

6 Discussion

Chemical rate controllers, which are based on the design we proposed, collaborate to optimally meet a globally-predefined rate limit. The reinforcement of this limit is *soft*: the controllers do not synchronize the transmissions (i.e., the instantaneous rates can assume any value). Instead over time, they cooperatively tend to converge to a limit value, which is on average distributed among all, and proportional to the time they have been active. Each controller actively monitors and infers the current network load from the aggregate transmission rate, without requiring knowledge of the number or identity of other transmitters. The inference process does not require any explicit or sophisticated estimation model: the LOMA-based scheduling effectively bridges workloads and rates through a simple predictable relationship that underpins the operation of the reaction network. As concentrations of molecules are a direct linear expression of workload, the granularity at which the system operates can be easily adjusted from the byte level to the packet level (or anything in-between). Sporadic sensing problems of the network utilization have a marginal and instantaneous only effect on the regulation process since the corresponding fractional increase of a molecular

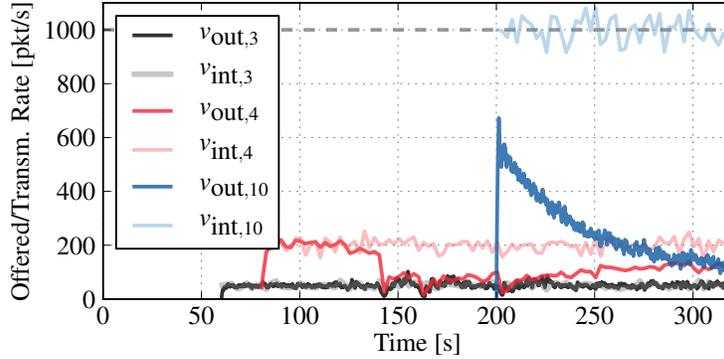


Figure 9: Offered load per node $v_{\text{int},j}$ and effective transmission rate $v_{\text{out},j}$ for the continuous transmission tests. Node 3 (starting at $t = 60\text{s}$) does not experience any rate regulation, node 4 (starting at $t = 80\text{s}$) experiences a small rate regulation, and node 10 (starting at $t = 200\text{s}$) experiences a major rate regulation.

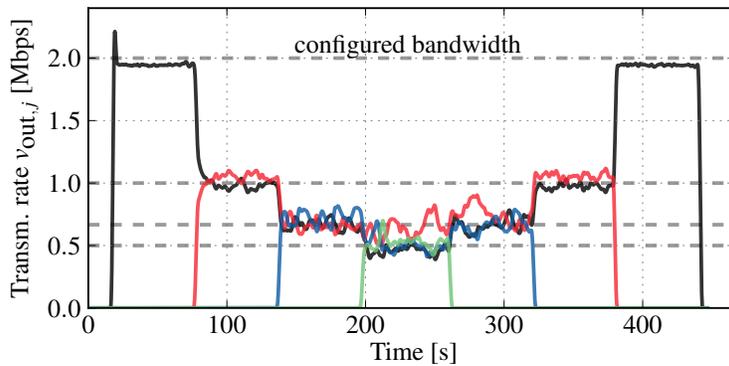


Figure 10: Transmission rates $v_{\text{out},j}$ of four nodes in a real deployment. The chemical rate controller triggers the egress device queue in the Linux traffic control system of the corresponding nodes. Traffic is generated by TCP ssh/scp sessions. The maximum shared capacity is $v_{\text{max}} = 2\text{Mbps}$, the feedback delay is 20ms.

concentration is small with respect to the entire concentration. Even in case of severe problems, the local transmission rate cannot increase beyond the maximum predefined threshold.

The core design we have presented is easily extendable (see [22]) with additional chemical reaction patterns in order to refine the sensing inputs (total utilization, partial utilization), as well as how the system distributes rates (equal, proportionally to traffic class, etc.). This extensibility alongside the parametrizability of the core design can lead to a family of rate controllers suitable for different applications and environments.

At the microscopic level a chemical reaction network as well as a traditional queuing network is described by a stochastic process – a continuous time discrete space Markov jump process. At the same time there is a subtle difference between the two approaches in the design of network systems. Using queuing theory, one starts by engineering the interaction of independent arrival and departure processes of a queue, and observe the resulting effects when queues interact. In chemical design, the departure process links directly to the queue’s fill level and thus, we can directly focus on engineering the interactions (reactions) of queues because the fill-level-proportional scheduling promotes the system’s self-stabilization.

The macroscopic control-theoretic analysis approach we followed is rather common-place in chemical engineering, and less systematically used in networking. Probably a reason is that one needs first to develop a suitable fluid model that describes a system accurately enough for the analysis to be credible, which often

turns out to be more of an art rather than science. However, by establishing a close analogy to chemical reactions this difficulty is overcome: with this generic representation, the mathematical fluid model comes for free and control-theory becomes almost a standard methodology.

7 Conclusions and Future Outlook

We have presented the design, analysis and evaluation of an adaptive rate controller that can be used for different purposes and environments where distributed rate control is needed. We have used a unique design methodology that derives from chemical engineering, and we have successfully validated both the performance of the proposed system as well as the credibility of the methodology, through simulations and real world tests.

In the future, more extensive real-world tests are planned to reveal the nature of interactions with existing congestion control schemes and traffic profiles of different services, which may lead to refinements of the core design.

Acknowledgment

The author would like to thank the Swiss National Science Foundation for the partial support of this research through grant #132525.

References

- [1] B. Raghavan, K. V. Vishwanath, S. Ramabhadran, K. Yocum, and A. C. Snoeren, “Cloud-control with distributed rate limiting,” in *Proc. of ACM SIGCOMM Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications*, August 27-31 2007, pp. 337–348.
- [2] H. Abu-Libdeh, Y. Vigfusson, K. Birman, and M. Balakrishnan, “Ajil: Distributed rate-limiting for multicast networks,” Computer Science Dep., Cornell University, Tech. Rep., 2008.
- [3] S. Keshav, “A control-theoretic approach to flow control,” in *ACM SIGCOMM '91*, September 1991, pp. 3–15.
- [4] H. Yin, P. Wang, T. Alpcan, and P. G. Mehta, “Hopf bifurcation and oscillations in a communication network with heterogeneous delays,” *Automatica*, vol. 45, no. 10, pp. 2358–2367, 2009.
- [5] C. Reder, “Metabolic control theory: A structural approach,” *J. Theor. Biol.*, vol. 135, no. 2, pp. 175–201, 1988.
- [6] B. Ingalls, “A frequency domain approach to sensitivity analysis of biochemical networks,” *J. Phys. Chem. B*, vol. 108, no. 3, pp. 1143–1152, 2004.
- [7] K. Jacobsson, L. L. H. Andrew, A. Tang, K. H. Johansson, H. Hjalmarsson, and S. H. Low, “Ack-clocking dynamics: Modelling the interaction between windows and the network,” in *Proc. of IEEE INFOCOM*, 2008, pp. 2146–2152.

- [8] A. Tang, L. L. H. Andrew, K. Jacobsson, K. H. Johansson, H. Hjalmarsson, and S. H. Low, “Queue dynamics with window flow control,” *IEEE/ACM Trans. on Networking*, vol. 18, no. 5, pp. 1422–1435, 2010.
- [9] C. V. Hollot, V. Misra, D. Towsley, and W. Gong, “A control theoretic analysis of RED,” in *Proc. of IEEE INFOCOM*, Apr. 2001, pp. 1510–1519.
- [10] C. V. Hollot, V. Misra, D. Towsley, and W.-B. Gong, “On designing improved controllers for aqm routers supporting tcp flows,” in *Proc. of IEEE INFOCOM*, Apr. 2001, pp. 1726–1734.
- [11] Q. Chen and O. W. W. Yang, “On designing self-tuning controllers for AQM routers supporting TCP flows based on pole placement,” *IEEE Journ. on Selected Areas in Communications*, vol. 22, no. 10, pp. 1965–1974, 2004.
- [12] N. Xiong, L. T. Yang, Y. Zhang, Y. Zhou, and Y. Li, “Design and analysis of a stable queue control scheme for the internet,” in *Proc. of the IEEE Int. Conf. on Embedded and Ubiquitous Computing EUC*, 2008, pp. 378–384.
- [13] B. Briscoe, “Flow rate fairness: dismantling a religion,” *SIGCOMM Computer Comm. Review (CCR)*, vol. 37, no. 2, pp. 63–74, Mar 2007.
- [14] F. P. Kelly, A. Maulloo, and D. Tan, “Rate control in communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [15] A. Charny, D. Clark, and R. Jain, “Congestion control with explicit rate indication,” in *Proc. of the IEEE Int. Conf. on Communications (ICC)*, vol. 3, Seattle, June 1995, pp. 1954–1963.
- [16] S. S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *IEEE/ACM Trans. on Networking*, vol. 1, no. 4, pp. 397–413, Aug 1993.
- [17] S. F. K. K. Ramakrishnan and D. Black, “The addition of explicit congestion notification (ecn) to ip,” RFC 3168, IETF, 2001.
- [18] L. Massoulié, “Stability of distributed congestion control with heterogeneous feedback delays,” *IEEE Trans. on Automatic Control*, vol. 47, no. 6, pp. 895–902, 2002.
- [19] T. Meyer, “On chemical and self-healing networking protocols,” Ph.D. Thesis, University of Basel, Switzerland, 2010.
- [20] T. Meyer and C. F. Tschudin, “A theory of packet flows based on law-of-mass-action scheduling,” in *Proc. of the IEEE Int. Symposium on Reliable Distributed Systems (SRDS)*, Oct. 2012, Irvine, California.
- [21] A. M. Lyapunov, *The General Problem of the Stability of Motion (Orig. 1892)*. London, USA: Taylor & Francis, 1992.
- [22] M. Monti and M. Sifalakis, “Extending the artificial chemistry to design networking algorithms with controllable dynamics,” University of Basel, Tech. Rep. CS-2012-003, July 2012.