The Math Citadel

Chemical Kinetics in Computer Networks

Rachel Traylor, Ph.D.
Research Assistant Professor, Marquette University
Co-founder/Chief Scientist
www.themathcitadel.com

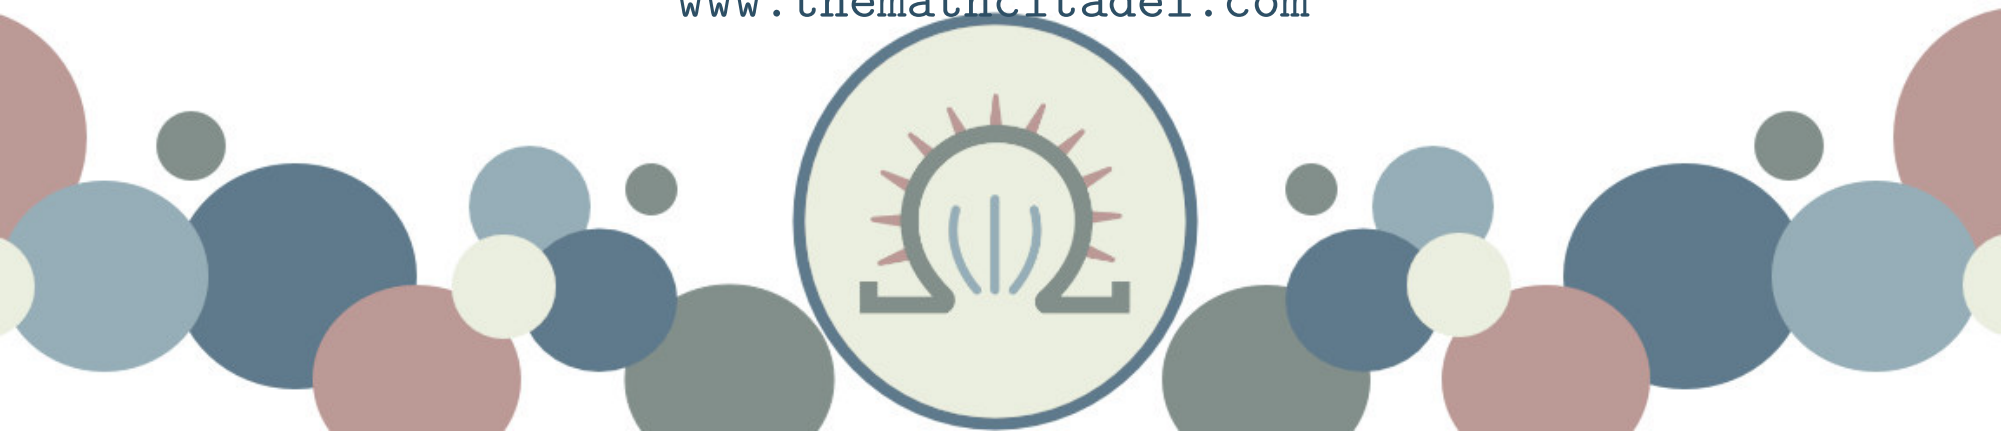

Content

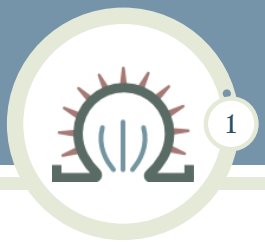

Computer Networks and ... Chemistry?

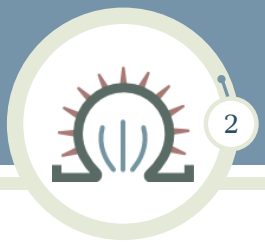

- Don't re-invent the wheel
- Work smarter, not harder
- ...and other such cliches

Let nature do some of the work

- Packet flow in a network has a beautifully natural analogue living inside chemical kinetics and reaction mechanisms
- Like particles or molecules, we have microscopic (or even quantum) behavior and macroscopic dynamics
- Currently, we notice issues in trying to scale microscopic (and generally more accurate) analysis. It quickly becomes intractable.
  - Smooth it away! (Comes at a steep cost.)

Overview of the talk

- Discuss the use of the Law of Mass Action in chemical kinetics in creating an artificial packet chemistry
- Note current applications of this model
- Briefly discuss limitations and future research

Background

LoMA in Chemistry

If we have a chemical reaction with reactants $A_1, \ldots A_m$, and products $B_1, \ldots B_n$, we can express the reaction (simply) as

$$A_1 + A_2 + \ldots + A_m \longrightarrow B_1 + B_2 + \ldots B_n$$

Example

$$HCl + NaOH \longrightarrow NaCl + H_2O$$

(plus lots of heat...)

LoMA in Chemistry

In a very simple reaction $A \longrightarrow P$, the rate this reaction occurs is a nice differential equation in terms of the concentration of the reactant:

$$-\frac{\mathrm{d}c_A}{\mathrm{d}t} = k \cdot c_A$$

where c_A is the concentration of the reactant A.

Put simply

The rate at which the reactant disappears is proportional to the current concentration.

LoMA in Chemistry

For a reaction with multiple reactants, the rate of a forward reaction is proportional to the concentrations of the reactants:

Law of Mass Action

$$r_f = k_f c_{A_1} \cdots c_{A_m}$$

and this concept is what we're going to use.

Queueing Analogue to Chemical Kinetics

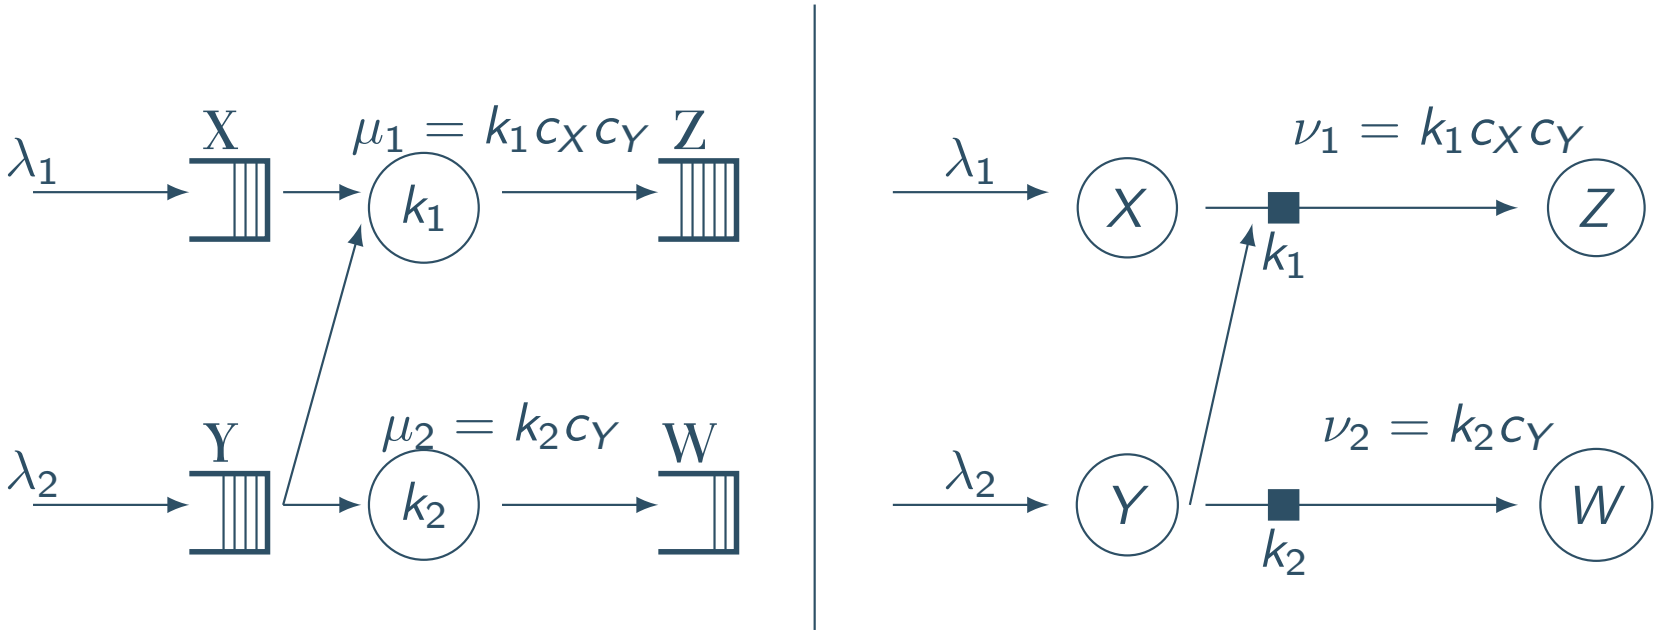

Description

Two packet queues served by two servers, the top extracts a packet from both queues.

Queueing Analogue to Chemical Kinetics

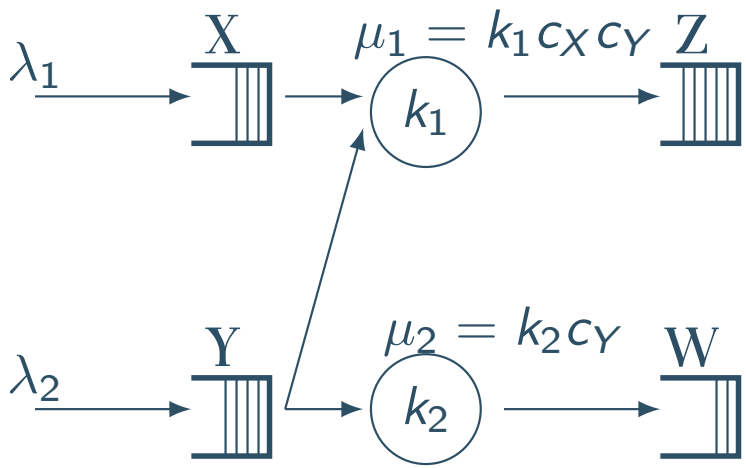

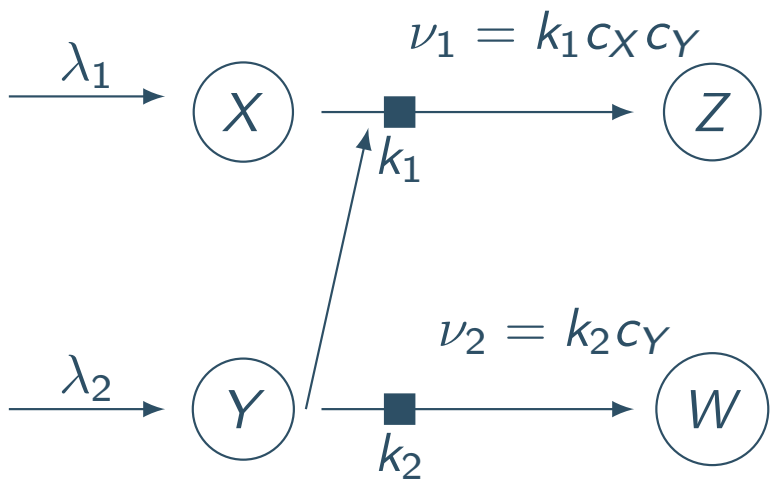

Notation

λ_i – arrival rate/inflow $\quad$ μ_i – service rate

ν_i – reaction rate $\quad$ c_X, c_Y – concentrations/fill level

Queueing Analogue to Chemical Kinetics

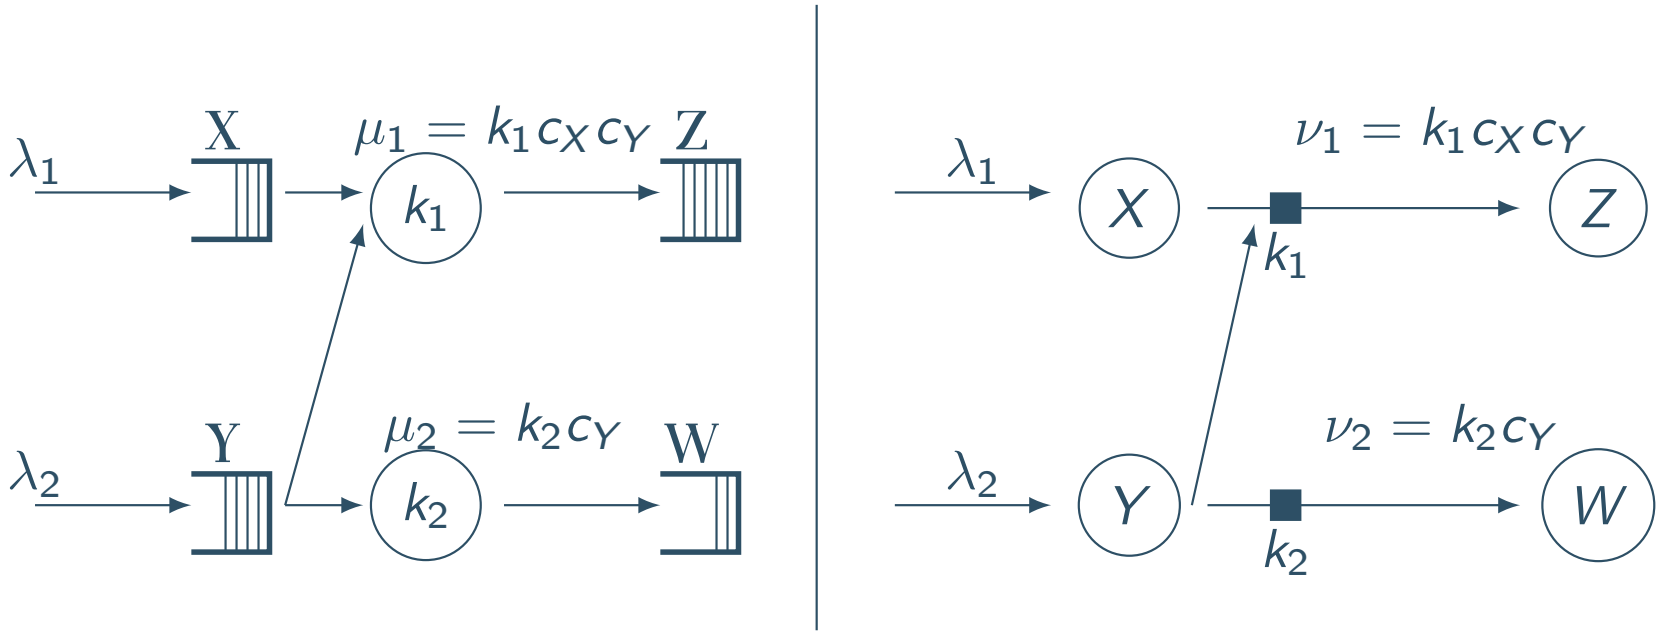

Production rate

The reactions are coupled, which can make analysis and scheduling complicated due to the coordination required here.

Assumptions

- Non-work-conservation: the server may stay idle even if the queue contains packets
- M/M/1 queues (noteworthy for mathematical analysis)
- FIFO queueing discipline
- equilibrium/steady state

What is being accomplished

- natural synchronization and scheduling
- ability to handle competing servers
- indirect communication through fill level of the buffer

Overview of Scheduling (Disperser example)

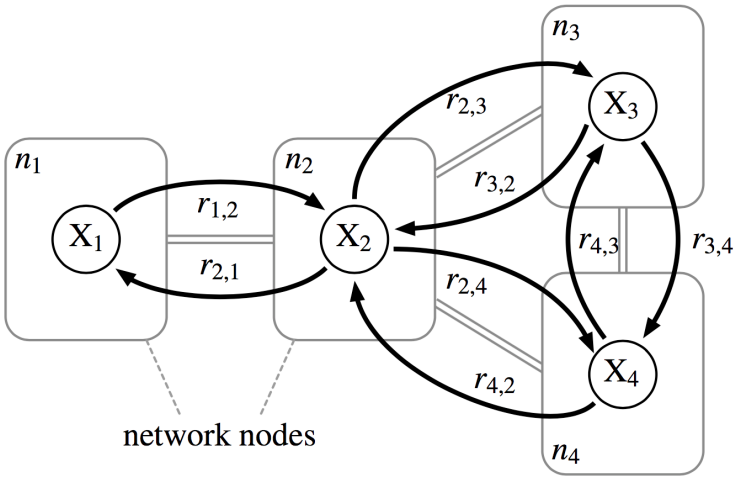

- LoMA competitor to Push-Sum
- Work-conservation would cause bottlenecking at n_2
- Non-work conservation allows for the goal to be achieved: calculating an average
- No symbolic computation: the result emerges in equilibrium from dynamic interaction of packet flow

Artificial Packet Chemistry

Def: Artificial Packet Chemistry

$$\text{PC} = (\mathcal{G}, \mathcal{S}, \mathcal{R}, \mathcal{A})$$

- $\mathcal{G} = (V, E)$ – graph that represents the computer network nodes, where $V = \{n_1, \ldots, n_{|\mathcal{V}|}\}$ are nodes, and $E = \{e_1, \ldots, e_E\}$ are unidirectional network links
- $\mathcal{S} = \cup_{i \in \mathcal{V}} \{S_i\}$ – lists queue instances of node i. Analogous to molecular species
- $\mathcal{R} = \cup_{i \in \mathcal{V}} \{R_i\}$ – gives flow relations among local queues. Analogous to reaction rules.
- $\mathcal{A}$ = reaction algorithm

Example

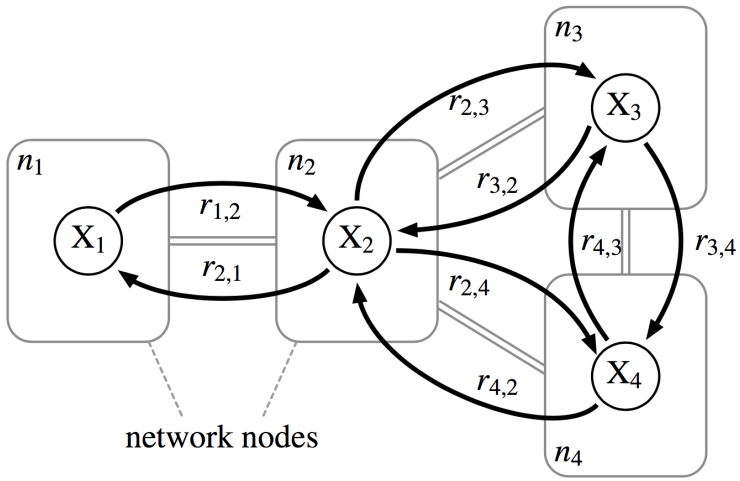

- $\mathcal{G} = (V, E)$, where $V = \{n_1, n_2, n_3, n_4\}$ and $E = \{n_1n_2, n_2n_1, n_2n_3, n_3n_2, n_2n_4, n_4n_2, n_3n_4, n_4n_3\}$
- $\mathcal{S} = \{X_1, X_2, X_3, X_4\}$
- $\mathcal{R} = \{r_{a,b} : X_a \to X_b : ab \in E\}$
- $\mathcal{A} =$ Law of Mass Action

Formal Analysis Techniques

- Microscopic stochastic analysis – Chemical Master Equation (McQuarrie, 1967)
- Deterministic analysis (fluid approximation) can be generated from the topology of the corresponding reaction network

$$\dot{x} = \lambda - kx \text{ with stable fixed point } \hat{x} = \frac{\lambda}{k}$$

- Kirchhoff's current law

Formal Analysis Techniques

- Transient analysis – Metabolic Control Analysis or signal/control theory approaches
- Mesoscopic analysis(estimation of the fill level variance)- Chemical Langevin Equations or linear noise approximation
- General stability analysis – Deficiency Zero Theorem, Chemical Organization Theory

Scheduler

- Each network node has its own scheduler to computer the next occurrence time of each $r \in R_i$ in its local node (processing packets)
- Sorts events into a priority queue, then executes when the first event occurs
- Difficulty: dynamically react and schedule as packets come and go
- Next Reaction Method (Gibson/Bruck): $O(n \log(n))$ time

Scheduler

- The rate constants and fill levels drive the ordering of events
- No timestamp tagging required, because the service rate is proportional to fill level of dependent queues.
- Interleaving and dynamic scheduling are automatic, in the same way chemical reactions (exothermic) happen

Going up a layer: Chemical Control Plane

- Instead of sending packets through a complex queuing network to shape packet flow, create a control plane that handles flow control

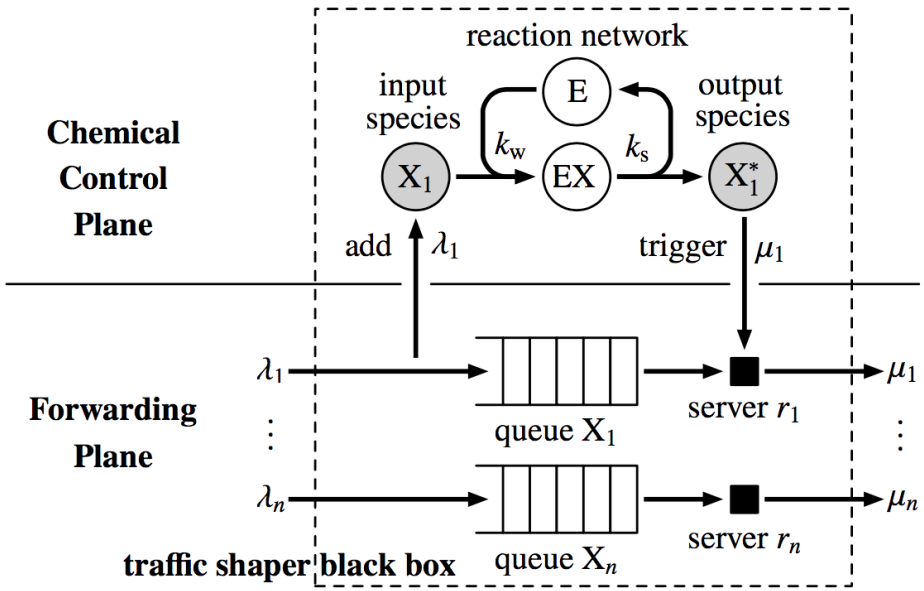

Figure 8: Chemical flow management architecture – The packet flow λ_1 is buffered in a classical FIFO queue while its server r_1 is triggered by a chemical reaction network. The enzymatic reaction implements a rate limitation policy.

Options

- We can use molecular reactions to naturally create shapes (e.g. low-pass filter with unimolecular reaction)
- Rate limiting via virtual "enzymes", following the Michaelis-Menten equation

$$k_w c_X c_E = k_s c_{EX}$$

where E are enzyme molecules and EX are enzyme-molecule substrates. Then

$$\mu = k_s(c_X + c_E)\frac{c_X}{\frac{k_s}{k_w} + c_X}$$

giving a hyperbolic saturation curve that approaches the M/M/1 queueing behavior in the limit.

Dynamic Self-Sorting via Fraglets

- Fraglets (Meyer/Tschudin, 2003) is a language that allows us to program a reaction network.
- Each packet (molecule) is a string called a fraglet
  - contains instructions and data
  - allows self-sorting into queues
  - tightly coupled with mathematical models

Congestion Control

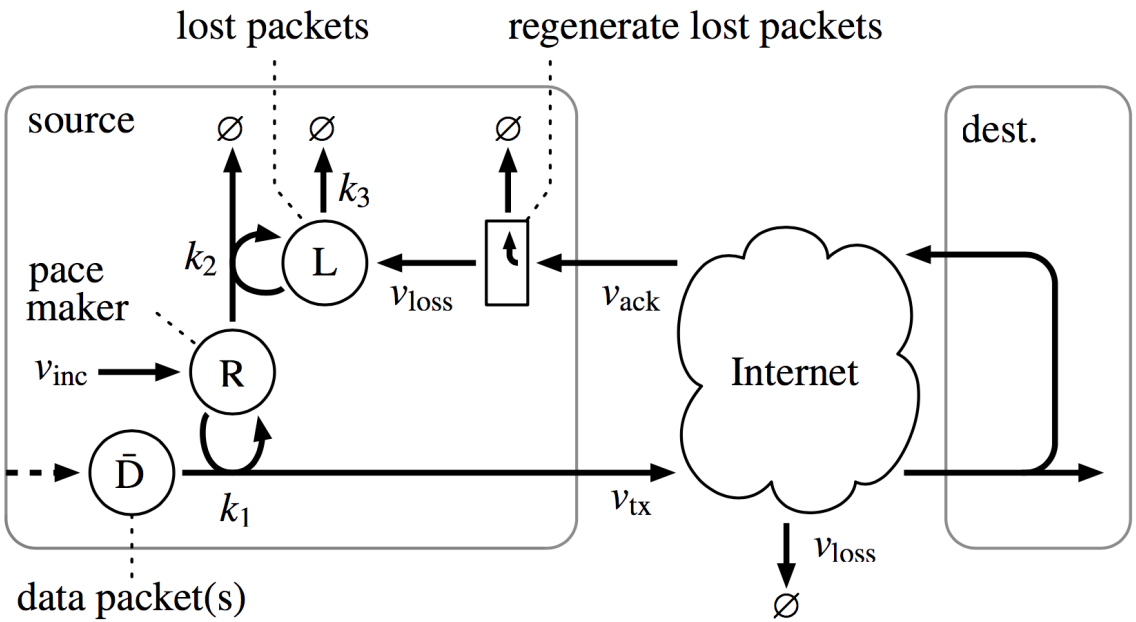

Figure: Meyer/Tschudin (2011)

Chemical Equivalent to TCP Reno

- transmission rate controlled by pace-maker molecules (enzyme) R

Limitations

- infinite capacity assumption
- M/M/1 type-behavior is the "link" (too simplistic?)
- as always, the assumption of independence among the packets themselves

Future Directions

- fuller exploitation of results from chemical physics
- Petri-net models for more complex networks
- finite capacity
- generalized models (M/D/1, G/M/1, heavy traffic perturbation)
- full implementation of reaction designs to facilitate other protocols

Conclusion

- An injection of lateral thinking and creativity, with huge potential
- A testament to the value of interdisciplinary and academic collaboration with the private sector