A distributed approach to bandwidth allocation in logical networks

Maya Ramakrishnan

Supervisor: Prof Peter Taylor

Co-supervisor: Dr Andre Costa

The University of Melbourne

AUSTRALIAN RESEARCH COUNCIL
Centre of Excellence for Mathematics
and Statistics of Complex Systems

Outline

1. Motivation
2. Problem framework
3. Buy/Sell heuristic
4. Continuous re-allocation scheme
5. Discretised re-allocation scheme
6. Work in progress

Motivation

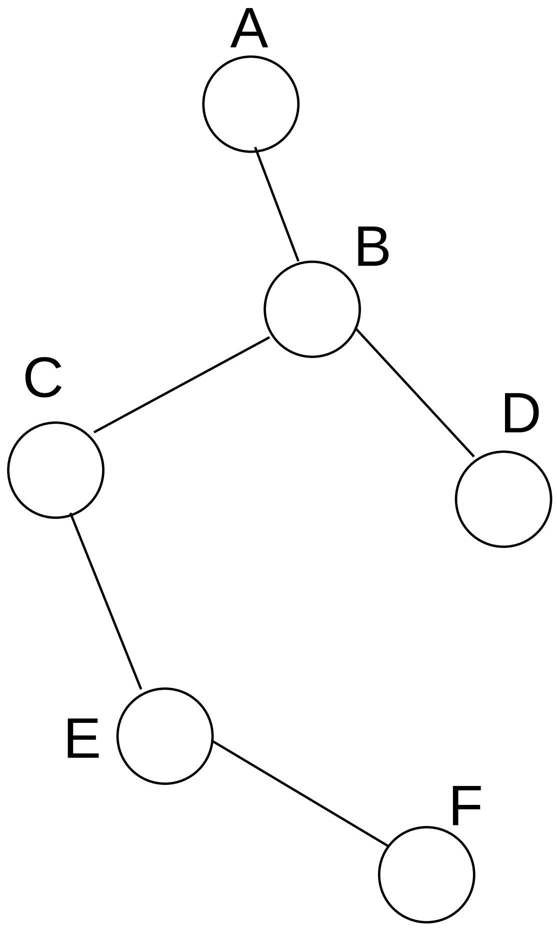

- Nodes $A \rightarrow F$

Motivation

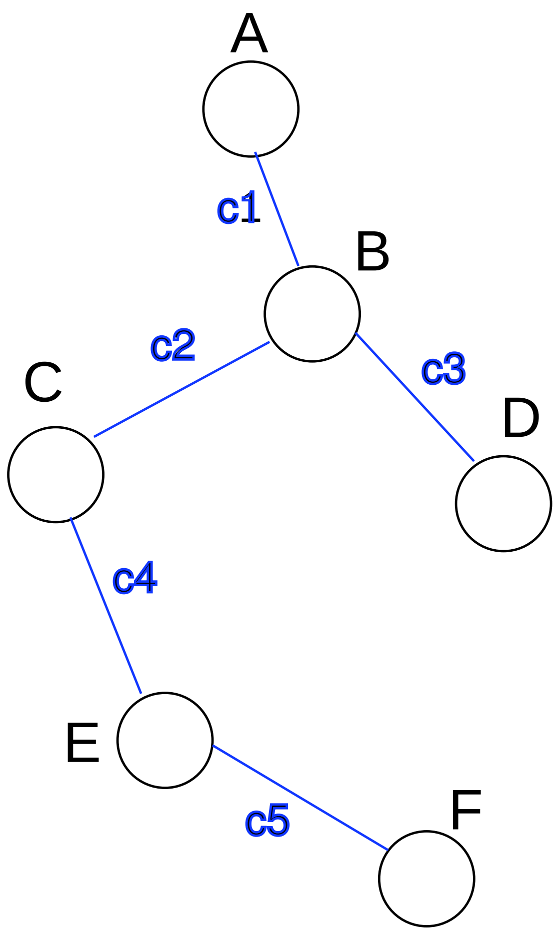

- Nodes $A \rightarrow F$
- Connected by a set L of physical links, each with capacity c_l

Motivation

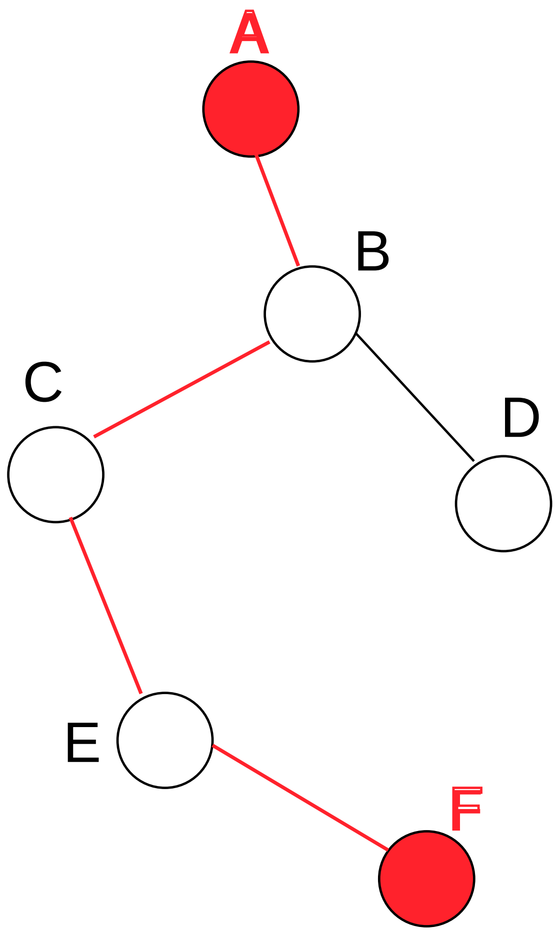

- Nodes $A \longrightarrow F$
- Connected by a set L of physical links, each with capacity c_l
- A route r is a non-empty subset of the physical links, with Poisson arrival rate λ_r and mean connection time μ_r^{-1}

Motivation

- Routes require capacity to carry traffic
- There are situations where one physical link is used in many paths through the network
- However, the physical links have capacity constraints

Motivation

- Routes require capacity to carry traffic
- There are situations where one physical link is used in many paths through the network
- However, the physical links have capacity constraints

The natural question that arises is

> How should capacity be allocated in a network between competing streams of traffic?

Problem framework

- A physical network overlaid by a logically *sufficiently* connected network
- When traffic arrives, we do not "check ahead" for spare capacity
- Reserve capacity between each origin and destination node ($A \to F$)
- This allows us to decouple the network and treat each route r as an Erlang loss system with capacity x_r
  - A continuous time Markov chain, where arrivals that find the system at capacity are lost; they are not queued

Problem framework

- Allocating capacity across routes in the network requires respect of capacity contraints
- Define a matrix A with elements

$$A_{lr} = \begin{cases} 1 & \text{if } l \text{ is in } r \\ 0 & \text{otherwise} \end{cases}$$

- Assume each route r has an associated utility function $U_r(x_r)$ and we wish to maximise utility over the network; ex. optimise quality of service over all routes

Problem framework

The optimisation formulation for the network *as a whole* is:

$$\max_{x_r} \sum_{r \in \mathcal{R}} U_r\left(x_r\right)$$

subject to

$$\begin{aligned} Ax &= C \\ x &\geq 0. \end{aligned}$$

Ex. $U_r(x_r) = \lambda_r \theta_r T\left(1 - E\left(\rho_r, x_r\right)\right)$ and $E\left(\rho_r, x_r\right)$ gives the blocking probability on route r.

Kelly's approach

Kelly et al (1998) Rate control for communication networks: shadow prices, proportional fairness and stability

- Different physical context, but similar mathematical formulation (key difference: inequality constraint)
- Mathematically tractable problem, but there exist *centralisation* issues
- Decomposition into a USER and NETWORK problem, that are tied together using Lagrangian arguments
- Using these ideas, a rate control algorithm is constructed

Kelly's approach

Kelly et al (1998) Rate control for communication networks: shadow prices, proportional fairness and stability

- Different physical context, but similar mathematical formulation (key difference: inequality constraint)
- Mathematically tractable problem, but there exist *centralisation* issues
- Decomposition into a USER and NETWORK problem, that are tied together using Lagrangian arguments
- Using these ideas, a rate control algorithm is constructed

Can we use a similar approach, but exploit our network structure in the solution method?

1. Motivation
2. Problem framework
3. Buy/Sell heuristic
4. Continuous re-allocation scheme
5. Discretised re-allocation scheme
6. Work in progress

Buy/Sell heuristic

- Using the utility function, we can derive buy and sell prices of a unit of capacity
- Assuming a route has capacity x,

$$\text{BUY}(x) = U(x+1) - U(x)$$

$$\text{SELL}(x) = U(x) - U(x-1).$$

Ex. The buy price using the previous example would take the form $\text{BUY}(x) = \theta\lambda T\left(E\left(\rho, x\right) - E\left(\rho, x+1\right)\right)$

Buy/Sell heuristic

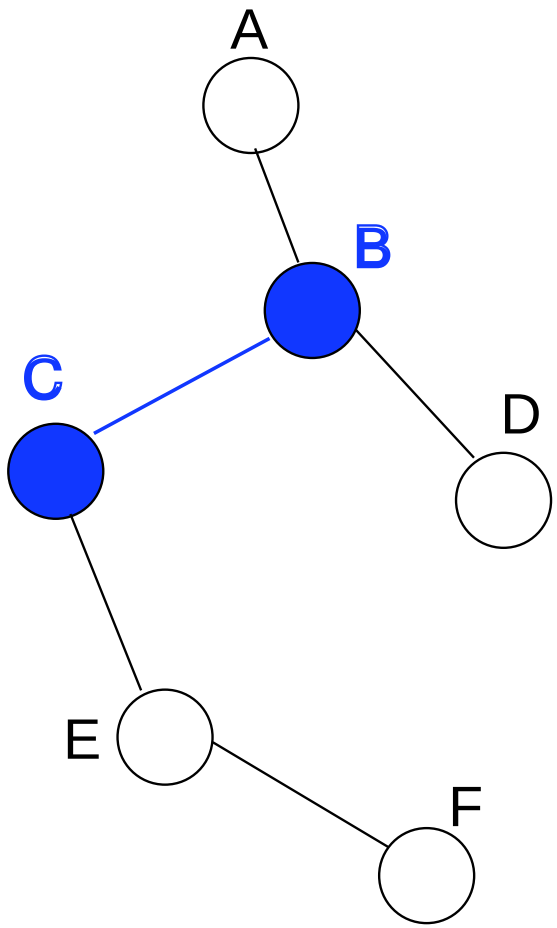

- *Direct routes* are those with origin and destination nodes connected by one physical link only

Buy/Sell heuristic

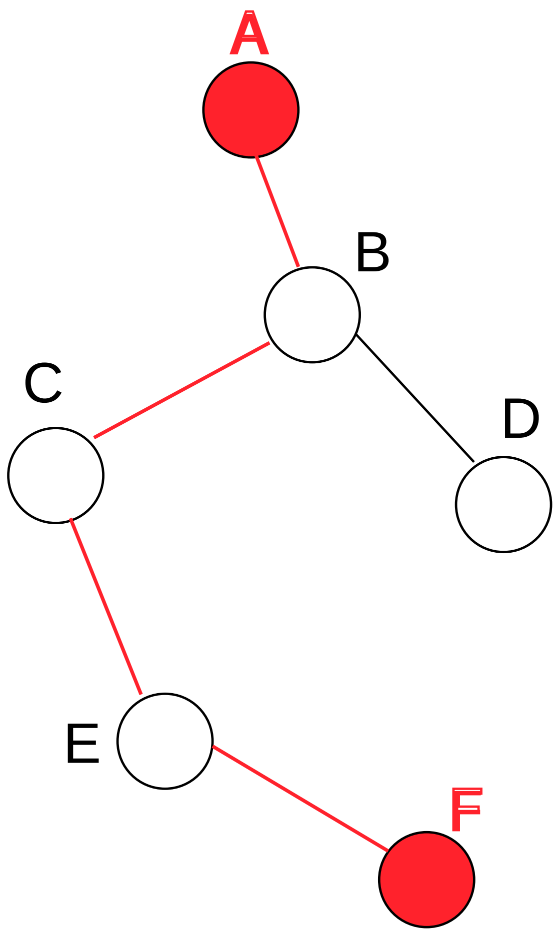

- *Direct routes* are those with origin and destination nodes connected by one physical link only
- *Transit routes* are those which utilise more than one physical link

Buy/Sell heuristic

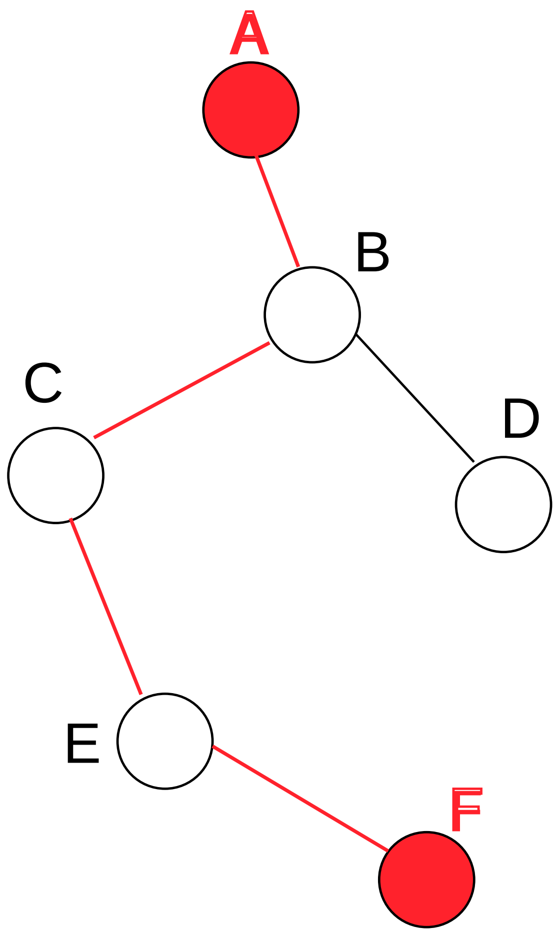

- *Direct routes* are those with origin and destination nodes connected by one physical link only
- *Transit routes* are those which utilise more than one physical link
- The proposed capacity trading scheme operates *locally*, that is transit routes can only trade with their constituent direct routes

Buy/Sell heuristic

- Chiera and Taylor (2002) derived a capacity value function to be used for this type of trading scheme
  - Modelled each route as an $M/M/C/C$ queue
- Chiera et al (2003) showed that this type of local interaction (between transit and direct routes) lowers blocking probabilities of a network
- At this stage, the method is a heuristic – it is not clear whether the global optimum can be reached
- Can we provide some theoretical support to this heuristic?

1. Motivation
2. Problem framework
3. Buy/Sell heuristic
4. Continuous re-allocation scheme
5. Discretised re-allocation scheme
6. Work in progress

Global problem (revisited)

We can split the set of routes $\mathcal{R}$ into $\mathcal{T}$, the set of transit routes and $\mathcal{D}$, the set of direct routes.

$$\max_x \left\{ \sum_{r \in \mathcal{T}} U_r(x_r) + \sum_{l \in \mathcal{D}} U_l(x_l) \right\}$$

subject to

$$\begin{aligned} \sum_{r \in \mathcal{T} \cup \mathcal{D} : l \in r} x_r &= C_l \quad \forall l \in L \\ x &\geq 0. \end{aligned}$$

Reformulation

We can rewrite the global problem solely in terms of capacity on the transit routes.

$$\max_x \left\{ \sum_{r \in \mathcal{T}} U_r(x_r) + \sum_{l \in \mathcal{D}} U_l \left(C_l - \sum_{r:l \in r} x_r \right) \right\}$$

subject to

$$\begin{aligned} \sum_{r \in \mathcal{T}:l \in r} x_r &\leq C_l \quad \forall l \in L \\ x &\geq 0. \end{aligned}$$

Karush-Kuhn-Tucker conditions

A Karush-Kuhn-Tucker point in this case will be the global optimum on the feasible region, with capacities x and Lagrange multipliers λ and η satisfying

(1) $U_r'(x_r) - \sum_{l:l\in r} U_l'\left(C_l - \sum_{r:l\in r} x_r\right) - \sum_{l:l\in r} \lambda_l + \eta_r = 0$

(2) $\lambda_l \geq 0$ and $\eta_r \geq 0$

(3) $\lambda_l \left(C_l - \sum_{r:l\in r} x_r\right) = 0$ and $\eta_r x_r = 0$ *(C-S)*

Karush-Kuhn-Tucker conditions

- Assume first that $x_l > 0$ and $x_r > 0$ for all routes in the network
- The KKT conditions specify, the optimal allocation satisfies
  $U_r'(x_r) = \sum_{l:l\in r} U_l'\left(C_l - \sum_{r:l\in r} x_r\right)$
- This is equivalent to a trading scheme where "infinitesimal" chunks of capacity can be traded
- Encouragement that the buy/sell heuristic was on the right track

Dynamics of capacity

Let the dynamics of transit route capacity be described by the system below

$$\frac{dx_r}{dt} = \kappa \left(U_r'(x_r) - \sum_{l \in r} U_l'(C_l - \sum_{s:l \in s} x_s) - \sum_{l \in r} \lambda_l + \eta_r \right)$$

- The fixed point of this system is equivalent to the KKT point

Dynamics of capacity

Let the dynamics of transit route capacity be described by the system below

$$\frac{dx_r}{dt} = \kappa \left(U_r'(x_r) - \sum_{l \in r} U_l'(C_l - \sum_{s:l \in s} x_s) - \sum_{l \in r} \lambda_l + \eta_r \right)$$

- The fixed point of this system is equivalent to the KKT point
- What are the parameters λ_l and η_r ?

Continuous scheme

- Applying the l_2 penalty method to our optimisation problem helps solve this problem

$$\lim_{k\to\infty} k(\sum_{s:l\in s} x_s - C_l)^+ = \lambda_l^*$$

$$\lim_{k\to\infty} k(-x_r)^+ = \eta_r^*$$

- From a practical perspective, we cannot set $k \to \infty$. Instead, we choose a large value K

Continuous scheme

- The fixed point of the system is *exactly* the optimal solution, when the solution is in the strict interior of the feasible region
- When the solution is on the boundary, setting K to be very large, arbitrarily closely approximates the optimal solution

Continuous scheme

- The fixed point of the system is *exactly* the optimal solution, when the solution is in the strict interior of the feasible region
- When the solution is on the boundary, setting K to be very large, arbitrarily closely approximates the optimal solution
- The fixed point is *attracting* – this is established using Lyapunov arguments

Network example

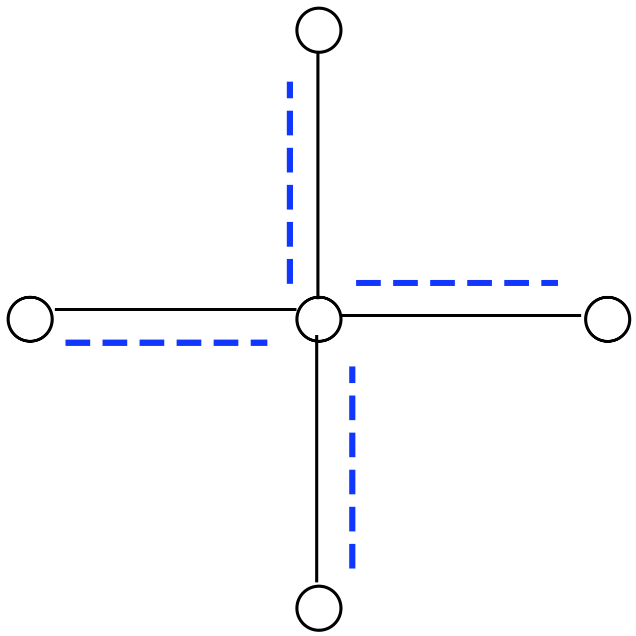

We consider a network with

- 4 direct routes

Network example

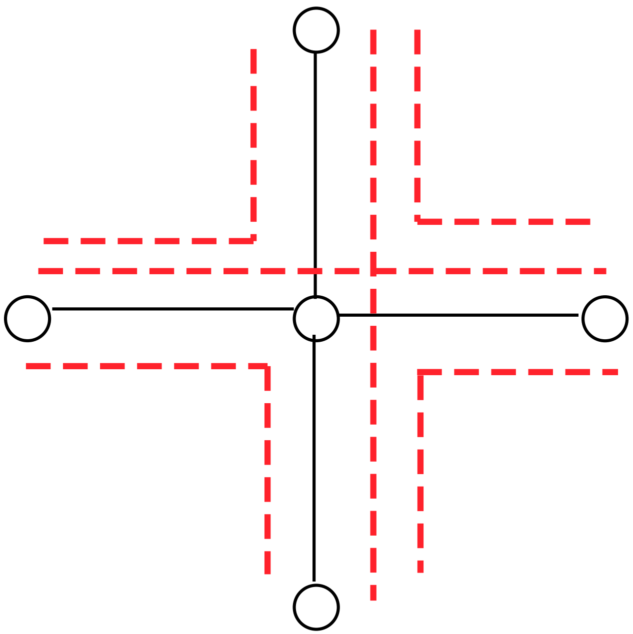

We consider a network with

- 4 direct routes
- 6 transit routes; it is logically fully connected

Network example

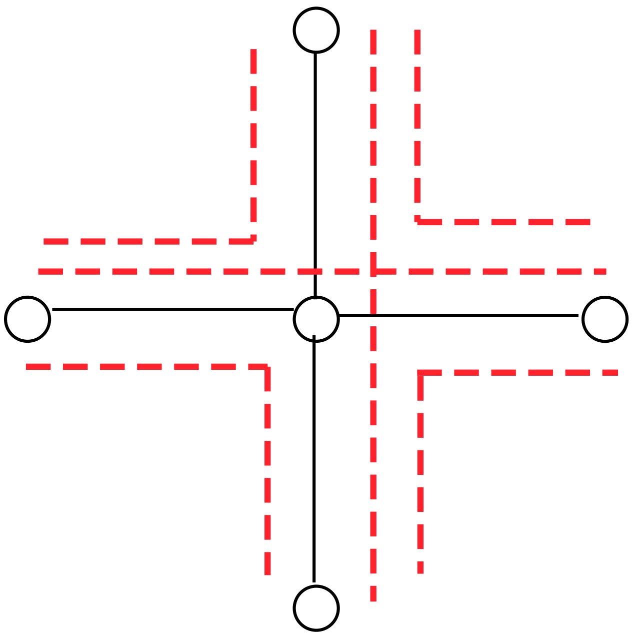

We consider a network with

- 4 direct routes
- 6 transit routes; it is logically fully connected
- Each route has a utility function dependent on parameters λ_r, the arrival rate and θ_r, the revenue generated from each accepted arrival, and is a function of capacity

Network example

The results using the continuous trading scheme are shown below

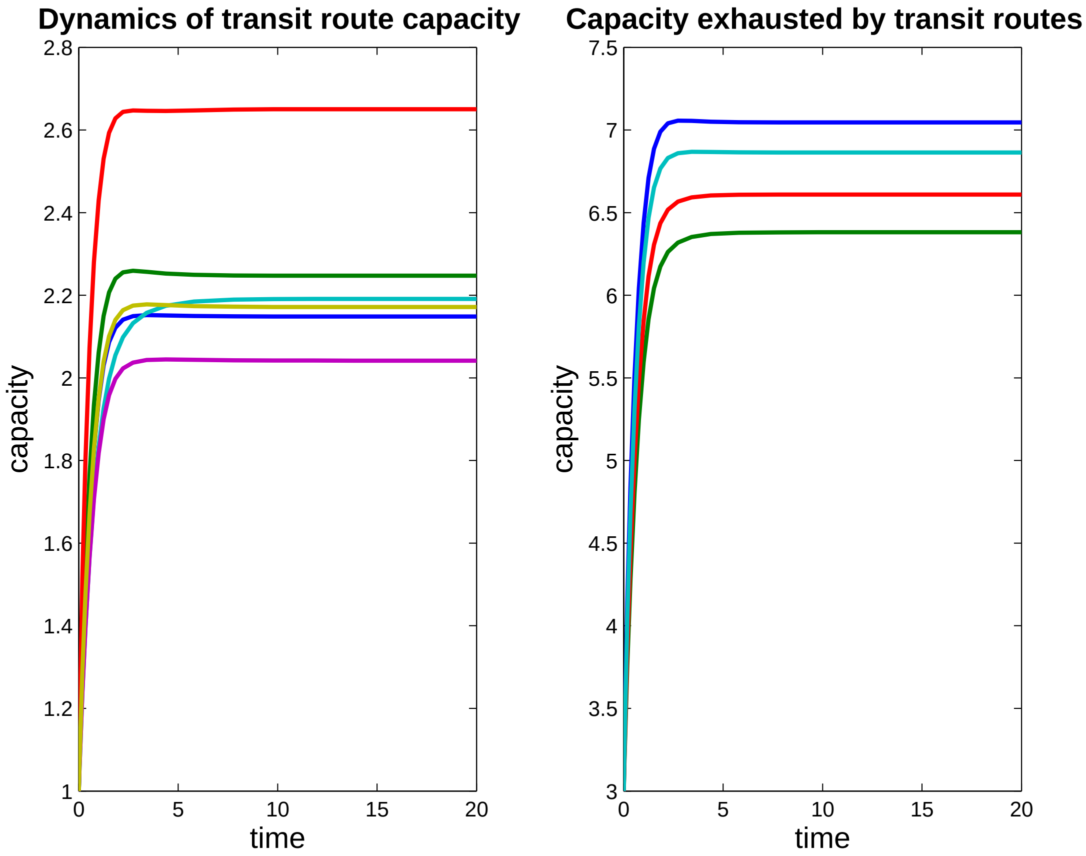

Network example

If the optimal solution is a boundary solution, the results (K = 100) are shown below

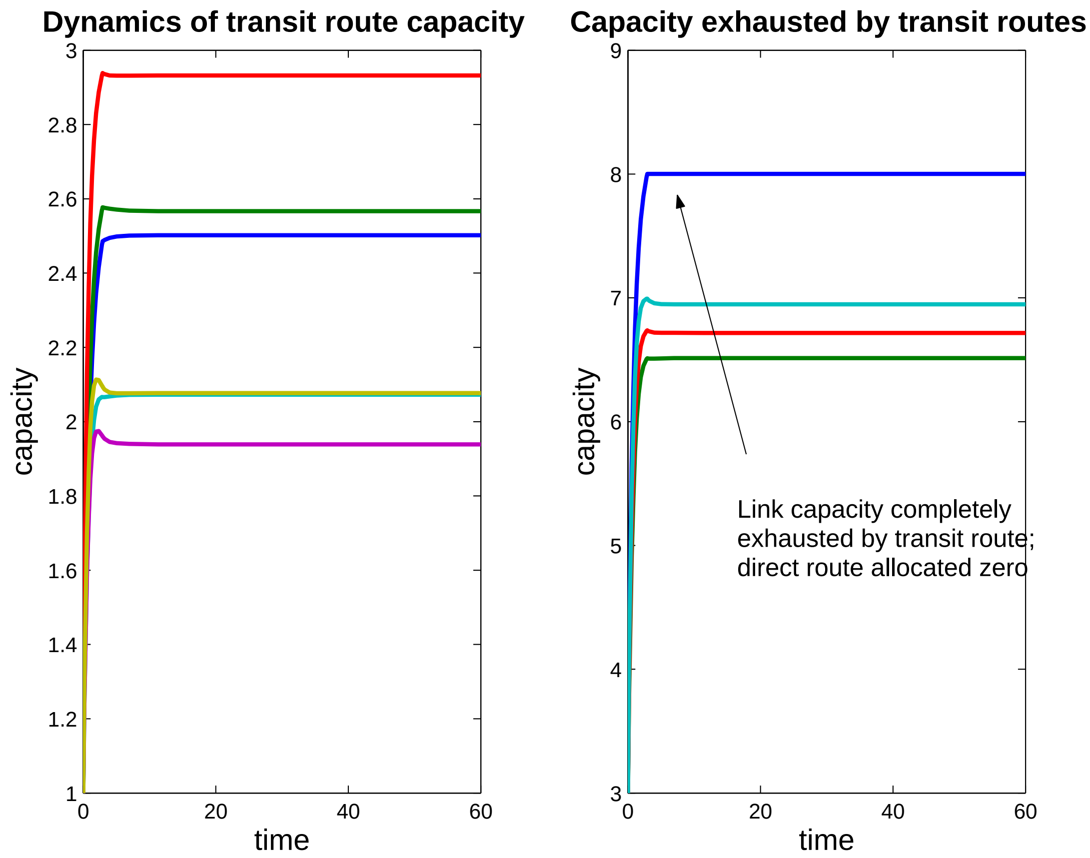

1. Motivation
2. Problem framework
3. Kelly’s capacity allocation method
4. Buy/Sell heuristic
5. Continuous re-allocation scheme
6. Discretised re-allocation scheme
7. Work in progress

Discretised system

- Assume that trades can only occur in amounts of Δ units at a time
- As described in the buy/sell heuristic, a trade occurs if one of the following conditions is satisfied:

Discretised system

- Assume that trades can only occur in amounts of Δ units at a time
- As described in the buy/sell heuristic, a trade occurs if one of the following conditions is satisfied:

$$Buy(x_r) > \sum_{l \in r} Sell(C_l - \sum_{s:l \in s} x_s)$$

$$\Rightarrow x_r := x_r + \Delta$$

Discretised system

- Assume that trades can only occur in amounts of Δ units at a time
- As described in the buy/sell heuristic, a trade occurs if one of the following conditions is satisfied:

$$\sum_{l \in r} Buy(C_l - \sum_{s:l \in s} x_s) > Sell(x_r)$$

$$\Rightarrow x_r := x_r - \Delta$$

Discretised system

- Assume that trades can only occur in amounts of Δ units at a time
- As described in the buy/sell heuristic, a trade occurs if one of the following conditions is satisfied:
- If neither condition is satisfied, no trade occurs.

Discretised system

- If a route has close to zero (or $< \Delta$) allocated capacity and the conditions state the route must relinquish capacity, the trade occurs
- However this yields an allocation that is not feasible

Discretised system

- If a route has close to zero (or $< \Delta$) allocated capacity and the conditions state the route must relinquish capacity, the trade occurs
- However this yields an allocation that is not feasible
- The next time prices are calculated, they will involve a penalty term, taking the same form as in the dynamical system.
- The relinquishing route will be able to acquire capacity at the next trade.

Network example

Allowing transit routes to instigate trading in a random order,

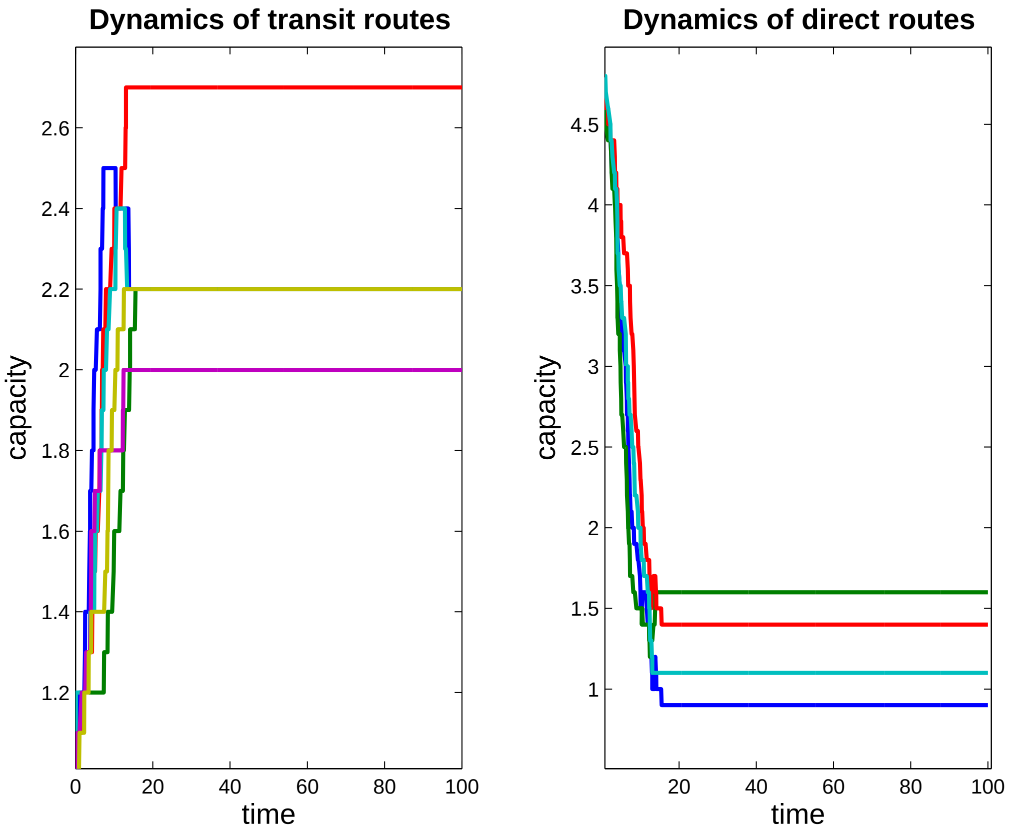

Network example

However, the system can also evolve to

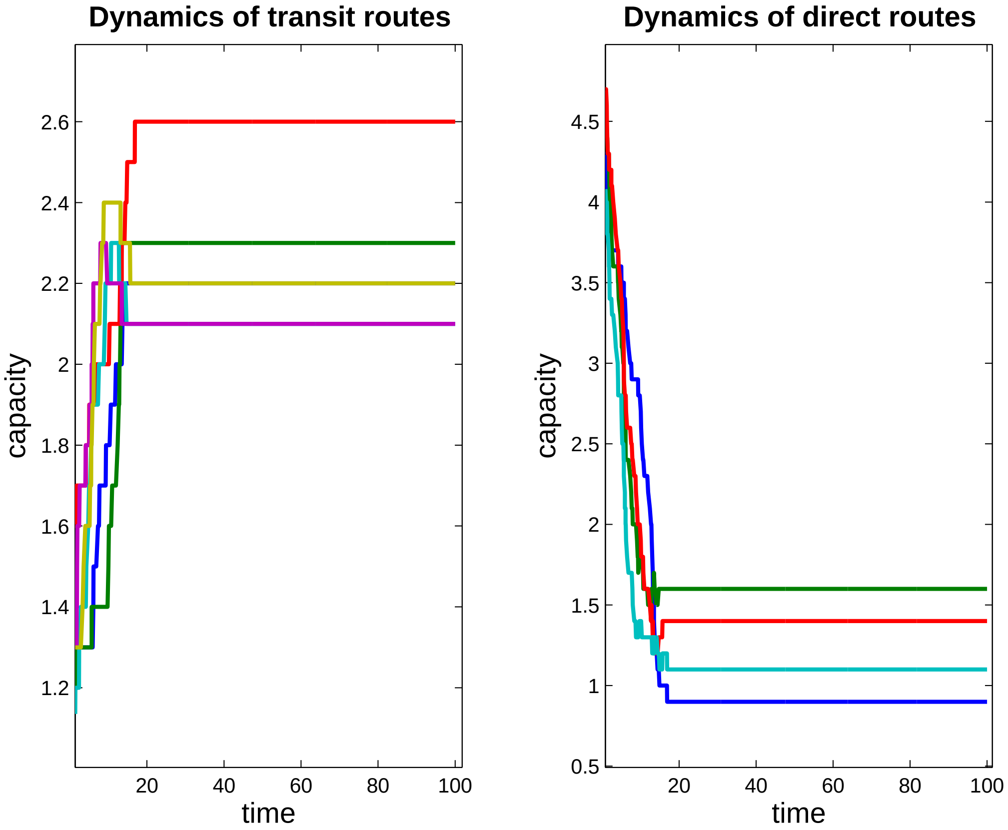

Network example

Oscillatory behaviour arises when a solution is on the boundary of the feasible region

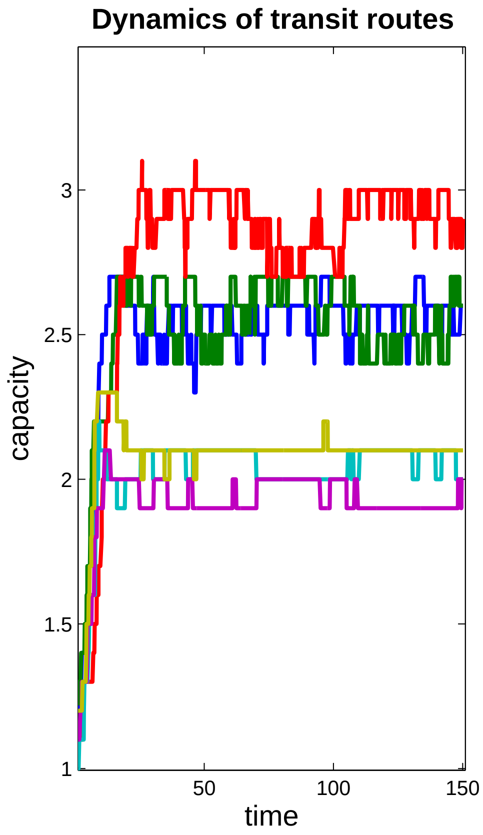

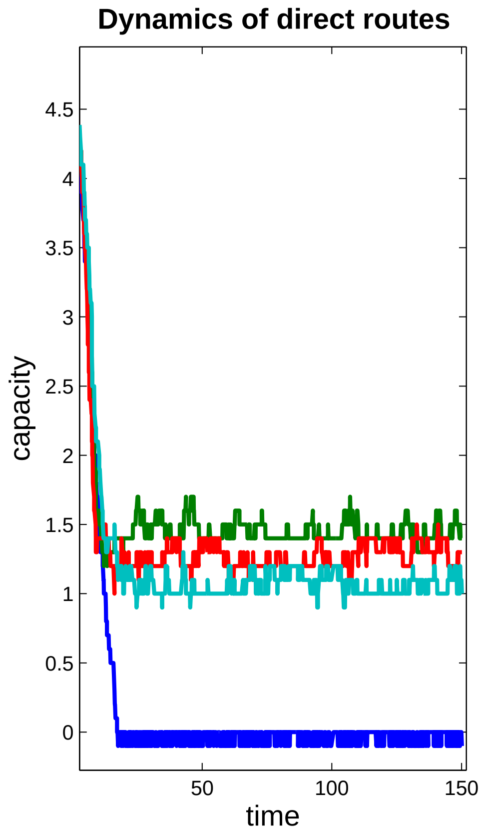

Network example

Introduce a threshold ϵ by which buy prices have to exceed sell prices before a trade occurs

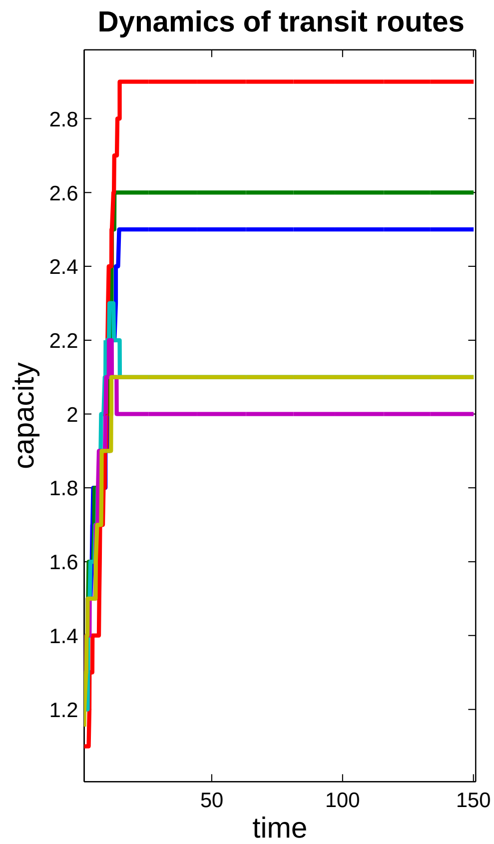

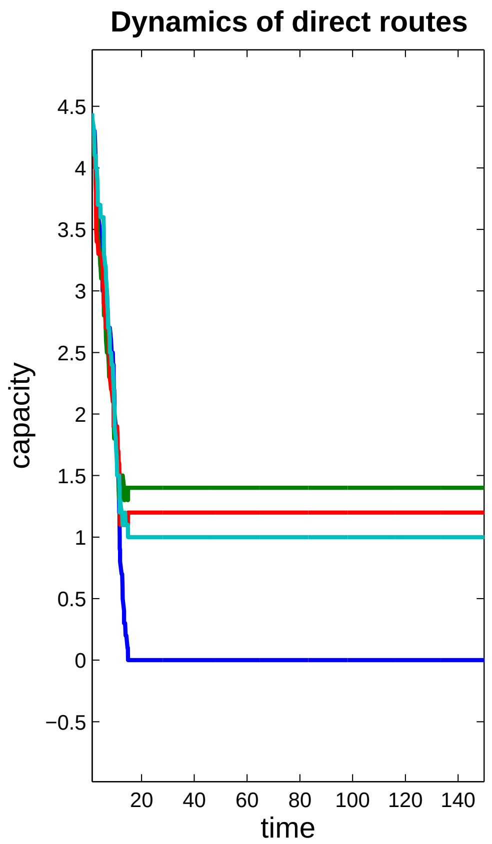

1. Motivation
2. Problem framework
3. Kelly’s capacity allocation method
4. Buy/Sell heuristic
5. Continuous re-allocation scheme
6. Discretised re-allocation scheme
7. Work in progress

Future work

- Can Δ be chosen so that the number of stable equilibria in the discrete system reduces to one? Analyse as the number of absorbing states in a finite-state Markov chain
- How “far” from the optimal can we say an equilibrium solution is, given multiple stable equilibria?
- How can ϵ be chosen using local information, and how much does it affect optimality?

Future work

- Can Δ be chosen so that the number of stable equilibria in the discrete system reduces to one? Analyse as the number of absorbing states in a finite-state Markov chain
- How “far” from the optimal can we say an equilibrium solution is, given multiple stable equilibria?
- How can ϵ be chosen using local information, and how much does it affect optimality?
- Model situation where occupancy on each route r is fluctuating stochastically through modified utility function
- Performance analysis of schemes