Resource pooling

Richard Gibbens

Computer Laboratory
University of Cambridge

Resource pooling workshop
Trilogy project
Computer Laboratory
2 July 2010

(Joint work with Frank Kelly and Stephen Turner)

A quick look back 20 years at one paper on resource pooling in loss networks

IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 1, NO. 2, APRIL 1993 261

Dynamic Routing in Multiparented Networks

Richard J. Gibbens, Frank P. Kelly, and Stephen R. E. Turner

Abstract—In this paper, we investigate some of the consequences for dynamic routing schemes of dual- and multiparented networks, in which a call can enter (or leave) the network at two or more points. In particular, we compare bounds on the performance of optimal dynamic routing strategies which respectively ignore and utilize the multiparented structure, and show that simple schemes, easily implemented and analyzed, are able to achieve most of the additional advantages allowed to dynamic routing schemes by multiparenting. Further, we illustrate the robust behavior of these schemes under traffic mismatches as well as multiple link or node failure events.

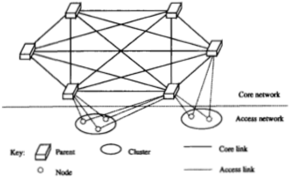

G, Frank P. Kelly and Stephen R.E. Turner
Dynamic routing in mutliparented networks
IEEE Transactions on Networking, 1(2), April 1993, 261–270

Network architecture

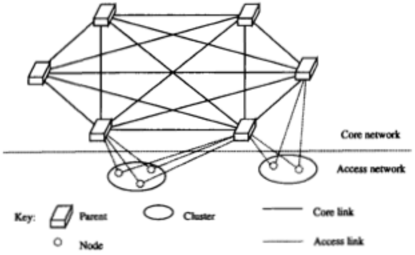

Fig. 1. The network structure.

- Core nodes fully connected act as parents for access nodes
- Access nodes connected to two (or more) parent core nodes
- Architectural drivers:
  - economics of switching and transmission costs ⇒ fully connected core network
  - desire for resilence to link failures ⇒ multiparenting

Routing between access nodes

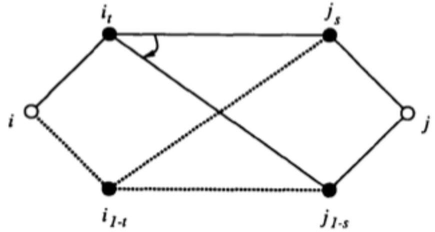

Fig. 2. Routing a call from i to j.

- There are four 3 link paths between access nodes *i* and *j* in distinct non-adjacent clusters.
- Routing strategies:
  - Single-parented routing (SPR)
  - Lastchance & sticky routing principles
  - Least loaded rounting (LLR)

Single-parented routing (SPR)

- Each call chooses its ingress and egress parent nodes without reference to the state of the network.
- Makes no use of multiparenting.
- Call may use any path (of one or more links) between the chosen parents.
- Lower bound on the blocking probability of any SPR strategy (uses a Markov decision problem formulation).
- Simple example strategy:
  - Each link has a trunk reservation parameter depending on its capacity and offered traffic.
  - Route directly if there exists a free circuit on the direct link.
  - Otherwise offer to a randomly ordered sequence of two link paths between the parent nodes any use the first (if any) that meets the trunk reservation condition on both hops.

Lastchance routing principle

- Lastchance: novel use of trunk reservation.
- Use trunk reservation to favour calls which are attempting their last possible route through the network and would therefore be lost if they are blocked.
- Example scheme: call chooses equiprobably a direct path from four available and uses this at low priority else selects the alternatve direct route via the other destination parent as a lastchance call at high priority. It is lost otherwise.

Sticky routing principle

- The sticky routing principle allows for a small amount of state memory.
- Calls with agiven destination have a current preferred ingress parent and the ingress parent has a current preferred egress parent for the given destination.
- If the preferred route fails the ingress parent overflows to the alternative parent. If carried there the ingress parent updates its preferred choice of egress parent.
- If the call can not be carried on the overflow direct route then the preferred choice of ingress parent is updated and that call is lost.
- Can also be combined with the lastchance principle.
- The simplicity of strategy together with a independent blocking assumption gives an accurate analytical approximation for the blocking probability.

Least loaded routing (LLR)

- Consider a family of strategies where a call can in principle select any route.
- Nevertheless, since the call must use at least one link from the core network, a resource pooling bound is to aggregate all offered traffic and all core network capacity into a single virtual link and apply Erlang's formula $E(\sum \lambda_{ij}, \sum C_{ij})$ to give a lower bound of any strategy in this family.
- Example strategy is least loaded routing (LLR), operates as follows: check for each call arrival the four direct paths and choose one, if any, with the most free circuits. If no direct path with free circuits exists then the call is lost.

Examples

Simulations and bounds when core links have capacities $C = 600$

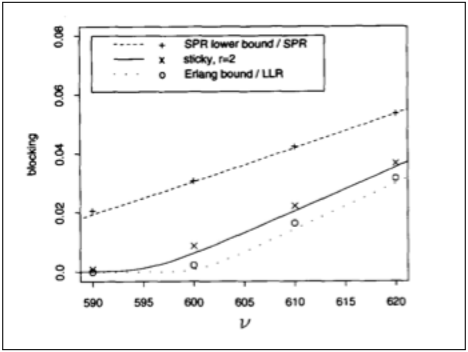

- LLR nearly achieves lower bound so no need to include strategies which consider longer paths.
- LLR can be hard to implement but lastchance/sticky nearly achieves all the benefit of LLR over SPR strategies and is rather easy to implement.

Broader picture

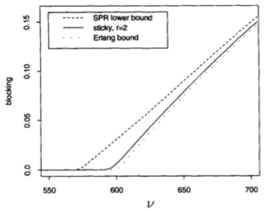

Fig. 5. "Symmetric" network.

- At low offered traffic all strategies achieve negligible blocking.
- At high offered traffic the strategies again perform similarly.
- At intermediate levels of load the LLR strategy can improve significantly better than any SPR strategy — though most of this benefit is achieved by a strategy as simple as the sticky strategy.
- These strategies are essentially pooling resources for all traffic.
- Overload will still occur when this pooled resource is saturated.

Robustness to link failures

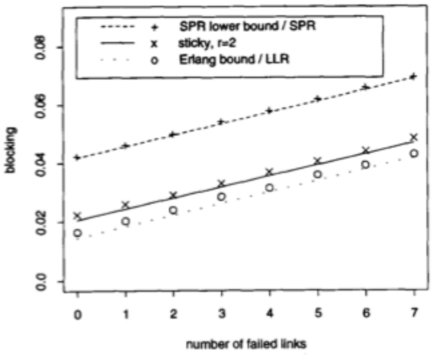

Fig. 6. Link failures.

- Fix the offered traffic at $\nu = 610$ and fail a randomly chosen sequence of core links.
- Blocking tends to increase as the capacity decreases and the relative effect on the strategies is similar.

Robustness to traffic mismatches

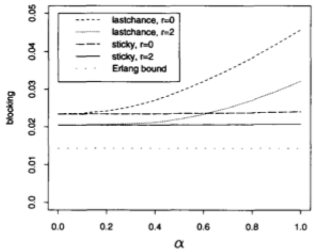

Fig. 7. Traffic mismatches.

- In this example, two pairs of source/destination node swap a random amount of traffic distributed $U[0, \alpha\lambda]$ with $\alpha \in [0, 1]$. Repeat sampling without replacement two pairs of nodes until the entire traffic matrix has been modified.
- Observe that the sticky strategy is able to adjust the routing pattern to the extent that it hardly notices the traffic mismatch (just like the Erlang bound itself).

Multi-path TCP

Modelling multi-path problems

(Invited Paper)

R.J. Gibbens
Computer Laboratory, University of Cambridge
William Gates Building, JJ Thomson Avenue, Cambridge, UK, CB3 0FD

***Abstract*—In this paper we consider the behaviour of both primal and dual multi-path algorithms for a simple network of three resources. We examine the equilibrium behaviour of our models as well as their transient response to the effect of a resource failing. The timescales over which the multi-path algorithms respond to changes in the network conditions are seen to be closely related to the round trip times of the different routes.**

I. Introduction

and

$$\mu_j(t) = p_j\left(\sum_{r:j\in r} x_r(t - T_{rj})\right). \qquad (4)$$

Each resource j has a capacity C_j and a penalty function given by

$$p_j(z_j) = \left(\frac{z_j}{C_j}\right)^{\beta_j} \qquad (5)$$

Modelling multipath problems
42nd Annual Conference on Information Sciences and Systems, CISS 2008, Princeton, NJ, USA, 19-21 March 2008
`http://dx.doi.org/10.1109/CISS.2008.4558492`

Toy network example

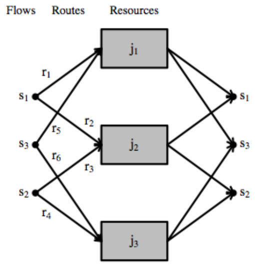

Fig. 1. The example network of three resources and a total of six routes.

- Each resource has capacity $C = 10$.
- Flow level simulations of primal and dual algorithms.
- Consider transient shock at $t = 0$ when link j_2 fails completely. Thus flows $x_2 = x_3 = 0$ for $t > 0$.

Primal algorithm

F. Kelly and T. Voice
Stability of End-to-End Algorithms for Joint Routing and Rate Control
ACM SIGCOMM Computer Communication Review, 35(2), 5–12, Apr 2005.

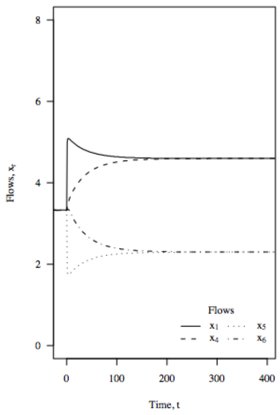

Fig. 2. Flows over time under the operation of the primal algorithm.

Dual algorithm

T. Voice
Stability of Mulit-Path Dual Congestion Control Algorithms
IEEE/ACM Trans. Netw., 15(6), 1231–1239, Dec 2007

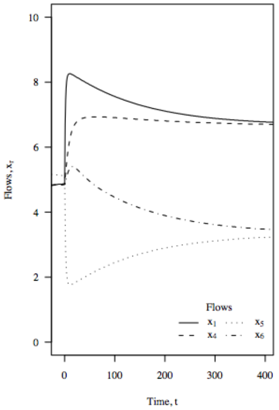

Fig. 4. Flows over time under the operation of the dual algorithm.