Reducing tandem queues (sketch)

James Hodgson

Under Andrew Moore, Cambridge Computer Labs

2017

Data centres and tandem queues

[Introduction to data centres and the important questions surrounding their functionality]

Data centres and tandem queues

A data centre might be modelled as a (tiered) network of queues. The customers are packets of data, which arrive from a host, travel through the aggregating layers, and are then sent to another host.

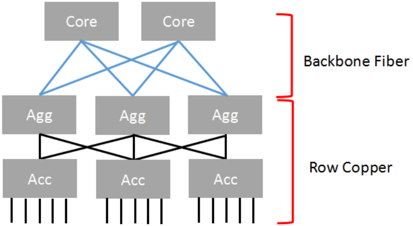

Data centres and tandem queues

Let's focus on an **single end-to-end route** through the network - this can be modelled as a sequence of queues in tandem, with incoming and outgoing cross-traffic.

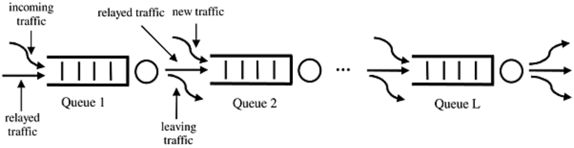

The inner workings of this tandem system are complicated - we might wish to reproduce its behaviour (ie the sojourn or holding times of the customers) more simply).

Data centres and tandem queues

One way to do this is to see if we can replicate the behaviour using a shorter tandem system, the simplest case being a single queue:

This is still quite a general set-up, with totally arbitrary cross-traffic and service distributions, so task does not seem unreasonable.

Analytic results

- One way to approach the problem is to assume a fully specified probabilistic model - arrival, cross-traffic, and service distributions, scheduling disciplines, buffer sizes and so forth
- There is some nice existing work on reducing networks of queues to smaller systems

Reduction Methods for Tandem Queuing Systems (Friedman 1963)

- Tandem queue $(Q_1, \ldots, Q_n)$
- General arrival process,
- m_i servers in queue i, each serving at constant rate s_i
- No cross-traffic, infinite buffers
- First note that the queue order in this system doesn't affect the exit distribution of the customers
- Say Q_i **dominates** Q_j if in the (different!) system $\rightarrow Q_i \rightarrow Q_j$, no customer waits at Q_j for any input process to Q_i

Reduction Methods for Tandem Queuing Systems (Friedman)

- There is a very simple criterion for domination in our constant service queues, namely that $s_j \leq \lceil m_j/m_i \rceil s_i$
- The property of dominance is **transient** and **persistent** (ie it operates across interposed queues)
- Then, for example, in a tandem system in which one queue dominates all the others, rearrange to bring that queue to the front
- Now no waiting occurs downstream, so all other queues can be replaced by a constant delay

Reduction Methods for Tandem Queuing Systems (Friedman)

- In general, reduction proceeds as follows:
- Reorder the Q_i in ascending order of m_i, and sub-order by s_i descending
- For each m_i, the queue with the highest s_i dominates all the others - move these to back of system
- For remaining queues, apply dominance test and move all dominated queues to back of system
- Finally, replace the dominated subsystem by a constant delay

Reduction Methods for Tandem Queuing Systems (Friedman)

Using the notation (s_i, m_i) for Q_i, consider as an example:

$$\to (5;2) \to (4;2) \to (2;1) \to (3;1) \to (7;3) \to (3;1)$$

$$\to (3;1) \to (3;1) \to (2;1) \to (5;2) \to (4;2) \to (7;3)$$

$$\to (3;1) \to (5;2) \to (7;3) \to (3;1) \to (2;1) \to (4;2)$$

$$\to (3;1) \to (5;2) \to (7;3) \to [9]$$

$$\to (3;1) \to (5;2) \to (7;3) \to [9]$$

$$\to (3;1) \to [21]$$

- It is well-known that $\cdot/M/1$ queues in tandem are also interchangeable
- It may be tempting to weaken the notion of dominance to a probabilistic one, and apply similar ideas - but persistence is not a all clear in this case
- In a data centre, service times are very fast, and queueing behaviour is generated by colliding cross-traffic'

Further analytic results

Further discussion of other queue reduction results, eg:

- A Tandem Network of Queues with Deterministic Service and Intermediate Arrivals (Shalman and Kaplan 1984)
- Delay analysis for tandem networks (Konheim and Reiser 1977)
- Reduction of a polling network to a single node (Beekhuizen, Denteneer, Resing 2008)

Limitations

- Fairly strong assumptions are needed on the nature of the arrival processes, service processes and other attributes of the system to make analysis tractable
- Tend to be exact, but we could be satisfied with approximations
- Generally we don't have complete information about the inner workings of a data centre
- Often limited to preserving only first moments - more accuracy is needed to produce good simulations

Turning the problem around

- Rather than assuming a fully specified probabilistic model, suppose instead we gather data describing the packet sojourn times through the tandem system
- Can we use this to construct a single queue with the desired properties?
- In other words, assuming the results came from a single queue, what can we infer about that queue?

Inverse problems in queueing

- Inference by probing
- The M/G/1 queue
- Lindsley's equation

The G/G/1 queue without cross-traffic

- The G/G/1 queue is the most natural simple model with which to begin
- Suppose we have a access to a set of arrival and departure times for the customers
- Then the service times are fully specified
- How good/poor is the approximation that these are independent samples from some distribution?

The G/G/1 queue with cross-traffic

- Discussion of several methods of determining cross-traffic and service distributions, with numerical examples, will make up the rest of the talk