TELE302 Lecture 17

Flow Analysis I: Concepts

Jeremiah Deng

TELE / InfoSci, University of Otago

7 September 2015

Lecture Outline

1. Basics
2. Flow Models
3. Flow Boundaries
4. Flow Distributions
5. Flowspec Algorithms

Flows

- Flow: a set of application or protocol information transmitted during a single session.
- Flows are end-to-end, between source and destination applications/hosts.
- Can be examined link-by-link or network-by-network.

Flow Analysis

- Flow analysis provides an *end-to-end perspective on requirements*, and shows where requirements combine and interact.
- An integral part of the analysis process.
- Gives insight into degree of *hierarchy* and *redundancy* needed in network design.
- Provides information on interconnection strategy (switching/routing/hybrid).

Network Flows

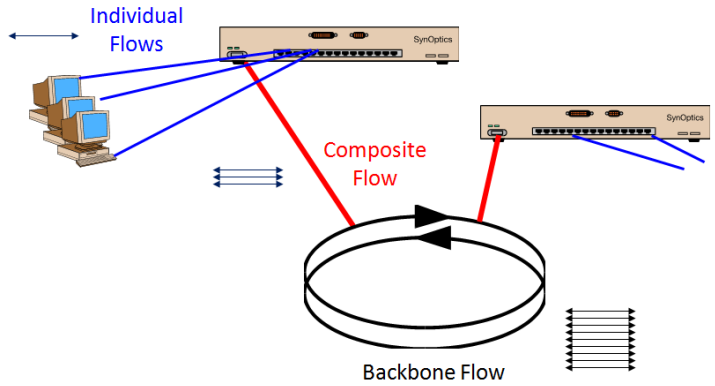

Different Types of Flows

- An **individual flow** is the flow for a single session of an application.
  - Best-effort delivery / specified characteristics.
  - Derived or estimated from requirement analysis.
- **Composite flow**: combination of individual (usu. best-effort delivery) flows sharing the same path, link, or network.
  - Used in capacity planning.
- **Backbone flow**: formed by composite flows when the network reaches a certain degree of hierarchy.
  - Indicates locations for special consideration in network capacity planning.
  - Useful for developing routing and addressing plan.

Data Sources and Sinks

- Data sources *primarily* produce network data.
- Data sinks *primarily* accept data from the network.

Flow Models - Peer-to-Peer

- No directionality or hierarchy
- Best-effort likely
- Examples:
  - LANs sharing file systems and printers
  - Early Internet
  - A set of Linux hosts supporting FTP, HTTP and NFS etc.
  - Gnutella/Bittorrent
  - Teleconferencing (involves Server)

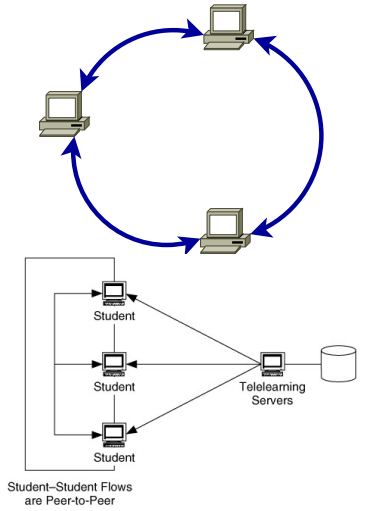

Flow Models - Client-Server

- Most generally applicable
- Asymmetric
- Internet becoming more client-server-like
- Add hierarchy → Cooperative Computing

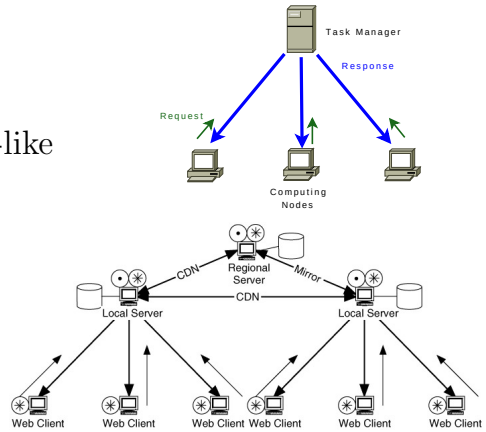

Parallel Processing

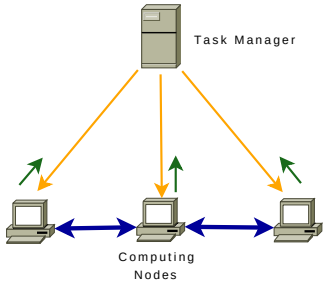

- Active connections between computing nodes
- Nodes likely to be both Source and Sink

Computing Cluster

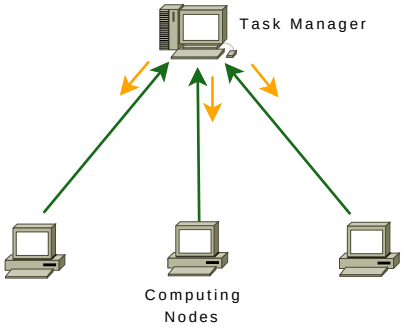

- Less traffic between computing nodes
- Computing nodes as Source; manager as Sink

The Making of *Avatar*

- The rendering done on Weta Digital's supercomputers ranked among Top 500.
- Over 4,000 compute blades running on Ubuntu Linux
- For throughput a 10 Gb network connects 40,000 processor to a large amount of high speed storage.
- A 12 node BlueArc Titan cluster served data at a rate of 8 Gb/sec to the compute grid.
- Final edited version of Avatar is over 2.8 TB of content.
- The multiple raw scenes are in the 100s of Terabytes.

Making *The Tale of Despereaux*

- Used a 200TB clustered file system closely coupled to a 6000 core render farm.
  - Lustre - an open source Linux cluster file system
  - 400TB mirror storage
- The processor cores are split across 1000 Dell servers each core having 2 GB of RAM.
- At the peak of production the cluster was using 4700 of its 6000 cores for work
  - 413138 hairs on the head of Despereaux
- Production data generated 5TB per night.

Flow Boundaries

- Geographic separation
  - LAN/WAN
  - LAN/MAN
  - MAN/WAN
  - Campus/Campus, Bld/Bld, Floor/Floor
- Logical separation
  - Backbones
  - Concentration points
  - WANs (likely to use service providers)
  - Specialized areas with specific service requirements

The 80/20 Rule

- 80% of traffic flow stays within a LAN and 20% is across the WAN.

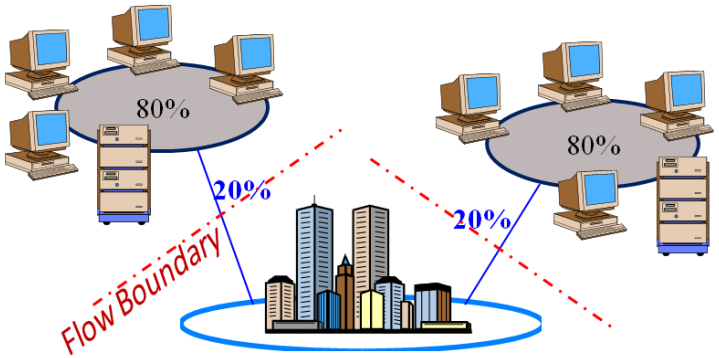

Server Farms: No more 80/20!

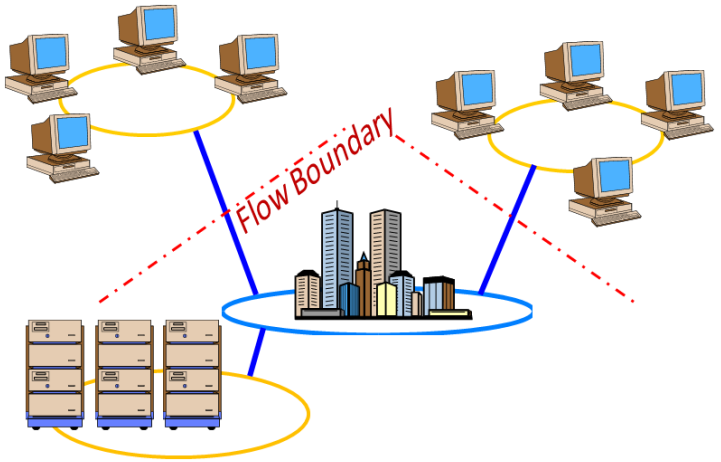

80/20 or not?

- Based on traditional usage patterns across LANs in WAN environment.
- Still a rule of thumb in sizing WANs in relation to the sizes of LANs.
- Can be extended to application requirements.
- Driving factors against 80/20 rule:
  - Enterprise trends in centralizing servers on corporate backbone
  - Increased distance-independent computing
    - Decoupling of physical and logical networks
    - Virtual networks
    - Flows are more likely to be in MAN/WAN

Applying Flow Distributions

- Client-server model?
  - Proportion of the number of clients on remote/local sides of the WAN?
- Distributed computing?
  - Are computing nodes tightly coupled or loosely coupled?
  - Are tightly coupled hosts connected via WAN?
- Peer-to-peer?
  - Traffic flows within LAN or across WAN?

The Small-Business Example (ref. Lecture 16)

- The flow distribution is ?

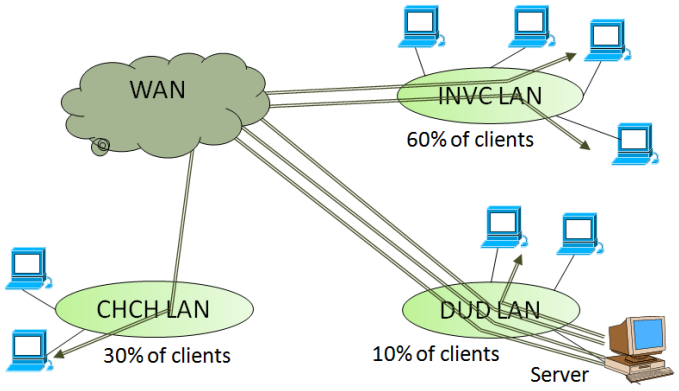

Flow Specification

- After identifying flows and having applications requirements listed, we can determine how to combine requirements for each flow and develop the flow specification (flowspec).

Flow Specification Types

- **Unitary** flowspec contains only best-effort flows.
- **Two-part** flowspec contains both best-effort and predictable flows.
  - Built on the unitary flowspec.
- **Multipart** flowspec describes flows with guaranteed requirements as well as predictable and best-effort flows.
  - Built on the two-part flowspec.

Flowspec Algorithm

- Combines reliability, capacity and delay characteristics for each of the flows to describe the expected overall performance.
- Applies the following conditions:
  - Use only capacity requirements for best-effort calculation. Use all characteristics available to calculate predictable flows. (two-part baseline)
  - Use guaranteed delay and/or reliability requirements individually. (multipart baseline)
  - Include consideration for performance modifiers if necessary.

Unitary Flowspec

- Add all best-effort capacities together
- Result: an overall baseline capacity.

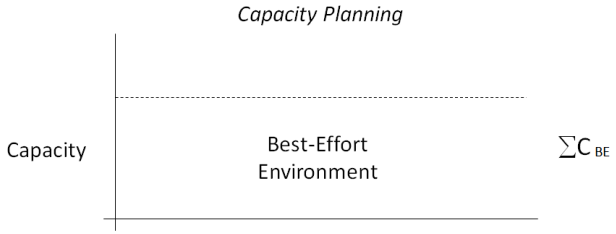

Two-part Flowspec

- *Add* all best-effort capacities together
- *Add* capacities for predictable flows together
- Take *maximum* reliability value and *minimum* delay value
- Result: an overall baseline capacity, an overall minimum delay, an overall maximum reliability

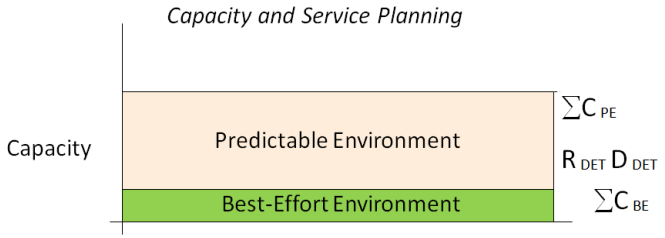

Multipart Flowspec

- Separate guaranteed flows
- Calculate best-effort flows and predictable flows as before
- List each of guaranteed requirements with capacity, reliability and delay.

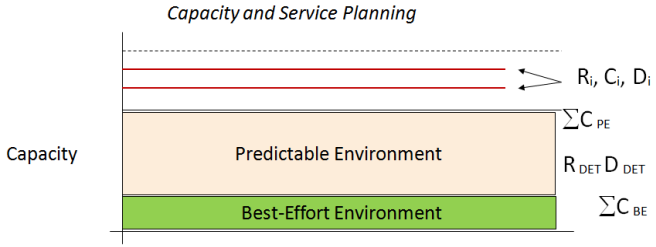

Summary

- Identify data sources and data sinks
- Identify flow models for applications / networks
- Locate flow boundaries
- 80/20 rule and its limit
- Flowspecs and Flowspec Algorithm
- Next: Flow Analysis Practice