



# Explicit Application- Network Cross-layer Optimisation

**Dimitrios Pezaros**


Computing Department  
Lancaster University  
[dp@comp.lancs.ac.uk]

12 Dec 2006 Progn06



## Overlay Networks and the Underlay Internet

- Increasingly popular to deploy customisable services at the application layer
  - File-sharing, object location, application-level multicast, network-embedded storage, etc.
  - Implicitly or explicitly take control over routing
  - Create highly variant & unpredictable traffic dynamics
- Multiple overlays over the underlay Internet
  - Can destabilise ISPs' TE and load-balancing policies at the underlay
  - Cost-ineffective for overlays to independently probe the underlay to optimise their operation
  - Fast reaction to underlay events ignoring background traffic

12 Dec 2006 LANCASTER UNIVERSITY 

## Cross-layer Issues

- Topology congruence not preserved
  - Black-box view of the underlay
  - Path selection based on an *overlay* cost minimisation function
- Underlay path measurement load
  - E.g. ping storms: 1 GB/day of ping traffic on planetlab
- ISP Traffic Matrices (TM)s
  - Entry duplication: overlay routing through intermediate access PoPs
  - Entry shift: alteration of ISP exit point for inter-domain traffic
  - ISPs assume stable traffic demands over short timescales
- Destabilisation of routing and Traffic Engineering (TE) policies
  - Enforce routing through high-(IP-)cost links
- Traffic oscillation
  - Two closed-loop systems react independently to the same set of events
  - Multiple overlays synchronising reaction to underlay events, e.g. link failures
- Inter-domain decoupling violation
  - Network state in one domain can influence behaviour in another

12 Dec 2006



## Proposed solutions: Underlay evolution vs. Explicit Cross-layer interaction

- Enrich underlay service to alleviate the need for overlays
  - Purposes and services offered by overlays are diverse & global
- Information exchange between overlay and underlay
  - Underlay exposes topology and performance information. Overlay exposes structure. However...
  - Global topology information exposure may not scale
  - Makes ISPs' policies even more vulnerable to opportunistic overlays
  - Conflict of business interests between stakeholders
    - Cross-layer conflict: between ISP and overlay applications
    - Intra-layer conflict: among ISPs

12 Dec 2006



## ISP Hints: Explicit Cross-Layer Interaction

- Explicit cross-layer interaction without exchange of structural information
  - Deploy services at the ISP level to take explicit input from the application and return hints about this input
  - Hints aim at improving network resource utilisation
    - E.g. minimise intensity and length of bursty requests; minimise duplicate inter-domain load
  - Applications should also have incentives to use ISP hints
    - Improved overlay performance and user satisfaction
- It is up to the respective layers
  - What inputs mean, how to use them, how to form hints
  - Each layer retains its own control to:
    - Set policies, optimise given performance metrics

12 Dec 2006



## ISP Hints: Examples

- Distance hint
  - Input: a set of destination IP addresses;
  - Output hint: a distance metric hint for each
    - ISP chooses what the metric is (hop count, AS path length, RTTs, etc.) and its resolution
- Similarity hint
  - Input: a set of destination IP addresses;
  - Output hint: clustered addresses based on a "similarity" metric
    - E.g. cluster all destination seen through the same egress link
- Hints can be composed by:
  - A set of Input/output formats; some semantics description
  - An ISP domain-local implementation
- Much of the information needed is already readily available to the ISPs

12 Dec 2006



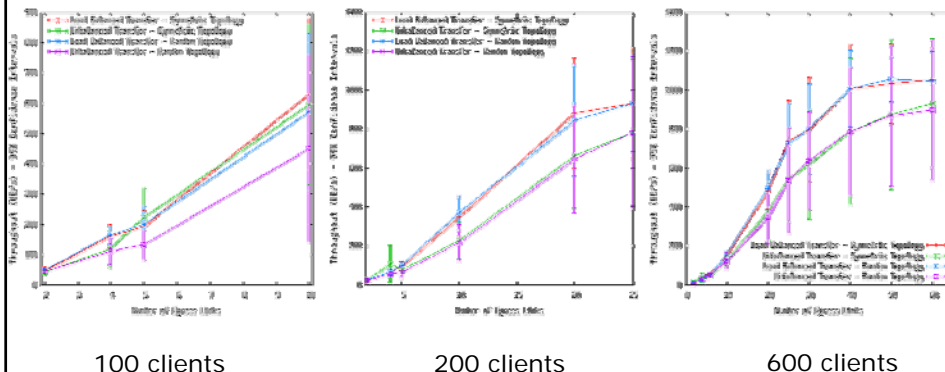
# Applications Using ISP Hints

- Client-server
  - Use similarity hint to cluster requests based on the egress links response traffic is routed through
  - Employ intra-cluster FCFS
    - Favouring inter-similarity cluster concurrency increases load balancing
  - Better handling of flash crowds
  - Increased effective network capacity
  
- Peer-to-peer file-sharing overlay
  - Employ similarity and distance hints simultaneously
    - Cluster requests per-egress link
    - Prefer minimum AS-hop distance providers among alternatives
  - Improved overlay/underlay topology congruence
  - Spread popular content faster to physically diverse segments of the Internet

12 Dec 2006



# Client-Server Performance using ISP Hints



12 Dec 2006




## P2P File Sharing Overlay

- Designed a bittorrent-like overlay to examine performance improvements of employing ISP hints
  - For the underlay
  - And the overlay
- Examined how a chunk (1MB) of data is distributed to the overlay nodes
  - During flash-crowd and steady-state phases
- Three different algorithms compared and contrasted
  - Simple overlay: no cross-layer interaction
  - ISP hints using domain-local knowledge
  - ISP hints using overlay-global knowledge

12 Dec 2006

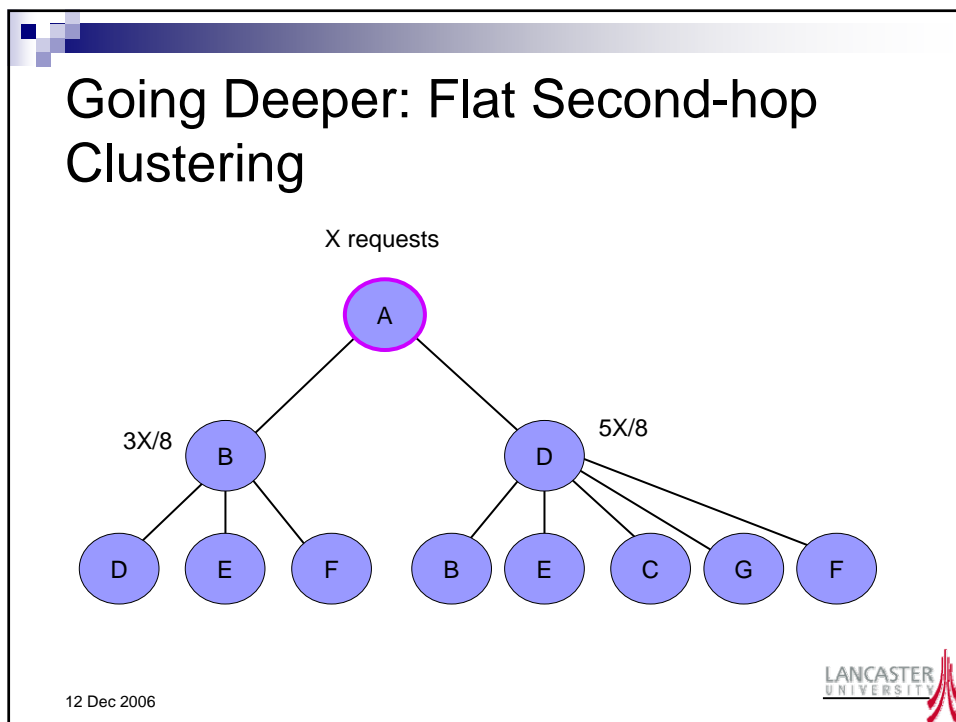
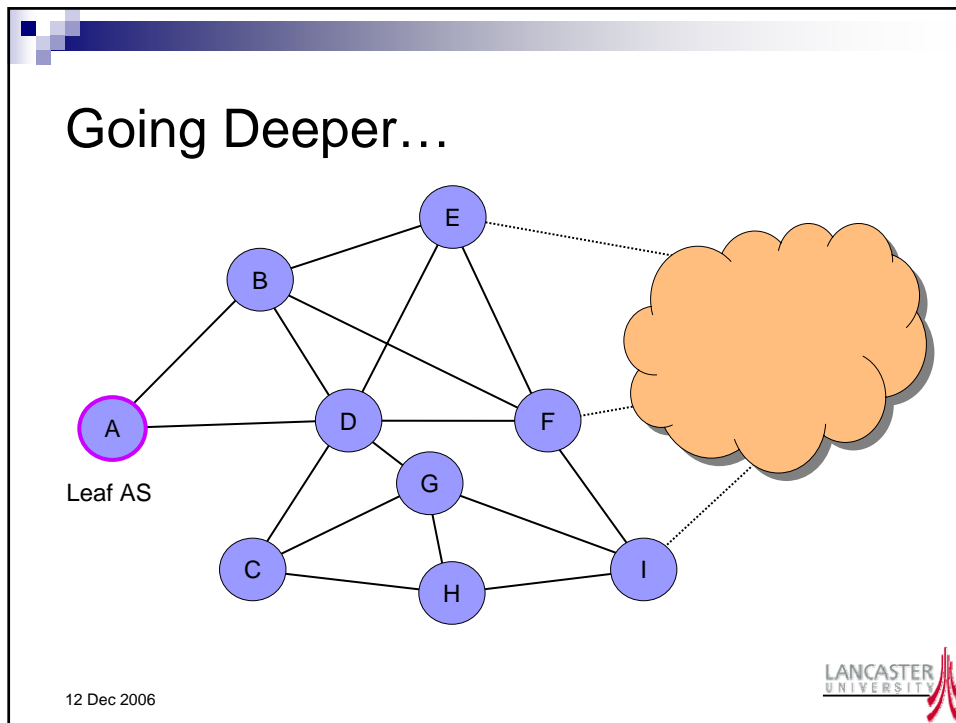


## Simultaneous Internet-wide application throughput and link stress improvement

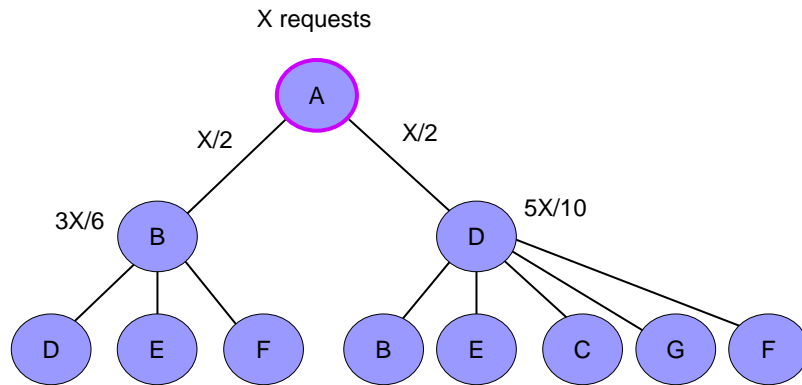


12 Dec 2006






# Going Deeper: Hierarchical Second-hop Clustering



12 Dec 2006



# How much does “deeper” knowledge give back?



12 Dec 2006



## Conclusions

- Synergistic cross-layer operation can improve perceived performance at both layers
- Explicit cross-layer interaction
  - With no explicit layer-specific information exchange
- Even when deployed using readily available information at the ISP level it can make a difference
- Enforcement
  - May not be necessary
  - Can be seen as a value-added service

12 Dec 2006

