Impact of Network Topology on Anonymity and Overhead in Low-Latency Anonymity Networks

Claudia Diaz[1], Steven J. Murdoch[2], Carmela Troncoso[1]

[1] K.U.Leuven, ESAT/COSIC
[2] University of Cambridge / The Tor Project

The problem

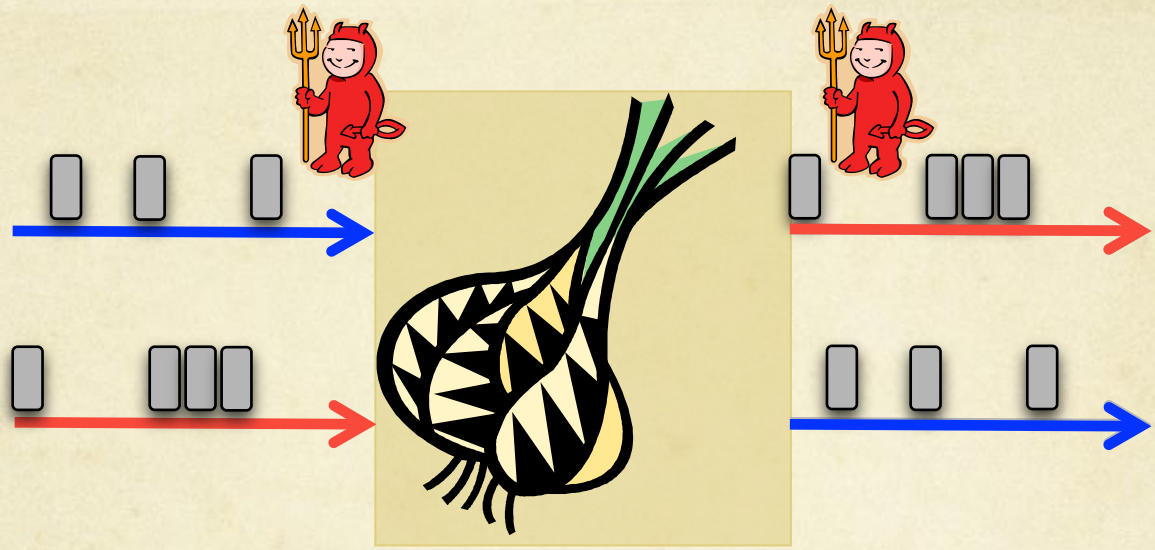

- Packet counting
- Inter-packet delays
- Start and end of streams
- Traffic watermarking (active attack)

Padding to resist traffic analysis

- Independent Link Padding (ILP)
  - Constant rate
  - Poisson
  - Any distribution as long as output rate is independent of input rate
  - If traffic rate highly variable this is very inefficient
    - Lots of padding: wastes bandwidth
    - Little padding: drop/delay real packets (bad QoS)
- Dependent Link Padding (DLP)
  - output rate dependent on inputs [VT08,WMS08]
- Synchronous start and end of communications !

Dependent Link Padding (DLP)

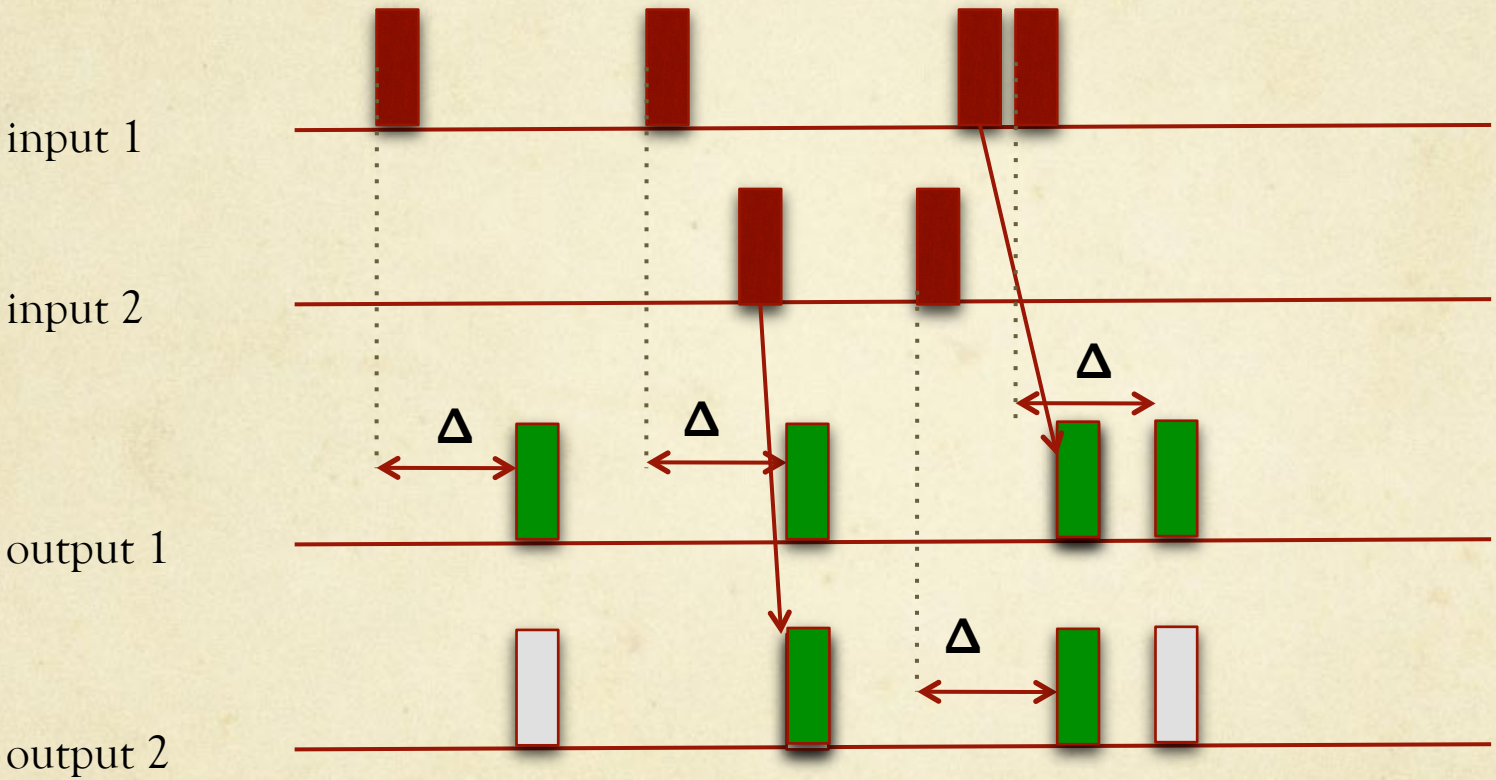

Our contribution

- If we implement DLP in a network, are some network topologies better than others?
  - Overhead
  - Anonymity (how to compute it? [TD09])
- Low-latency (circuit-based) anonymity networks multiplex the circuits between two routers over the same link
  - Can this help to further reduce overhead?
- Can Tor support DLP? Which modifications would be needed ?

Network topologies

- Evaluation through simulations
  - Same (average) traffic load per node
  - Same traffic load for the network as a whole
- Input: real Tor traces
  - Packet timestamp per circuit (bi-directional)

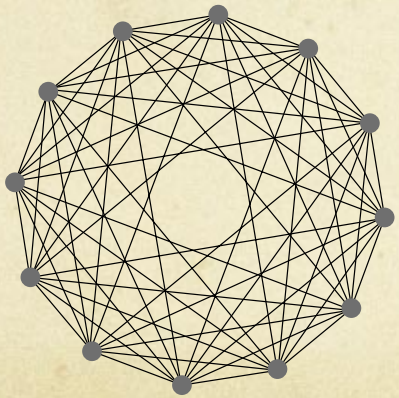

Free Route

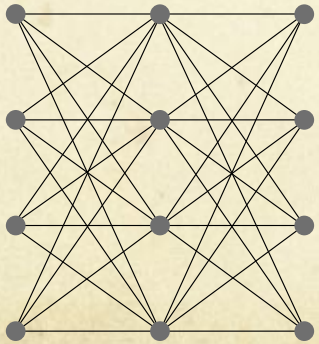

Stratified

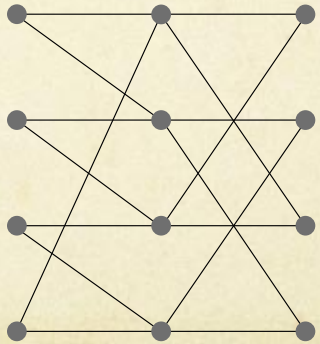

Stratifed Restricted

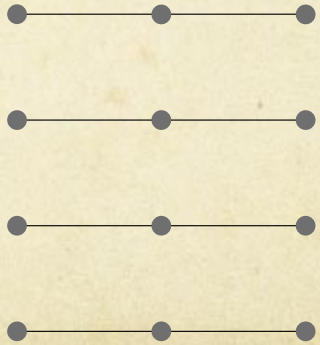

Cascade

Reducing Overhead (RO-DLP)

- Multiple circuits going between two nodes are multiplexed (link encryption)
  - Adversary cannot distinguish which packet belongs to which circuit
- If a link carries more circuits than input packets to be forwarded at time t, then we do not need to send packets on all the circuits
- Send min(C_i, packets)

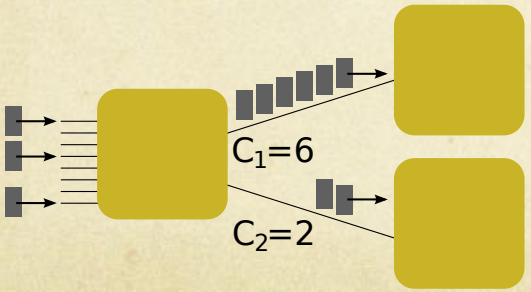

DLP

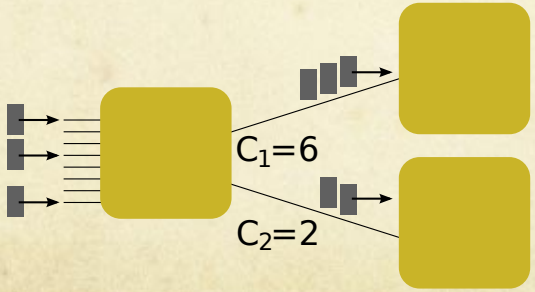

RO-DLP

Overhead (intra-network links)

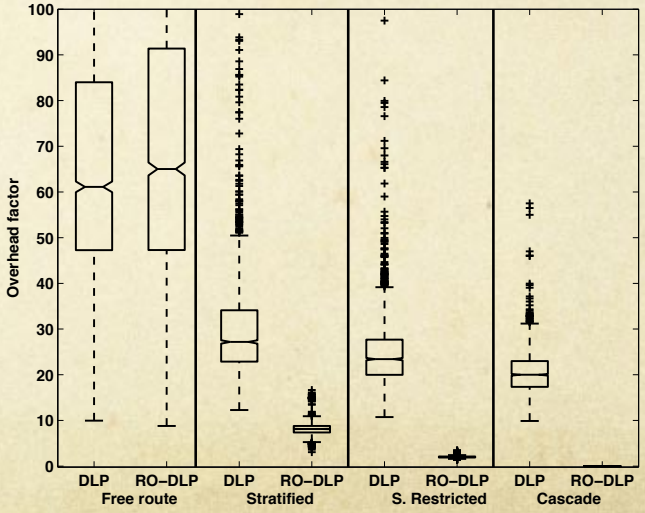

Feedback Effects in Free Routes

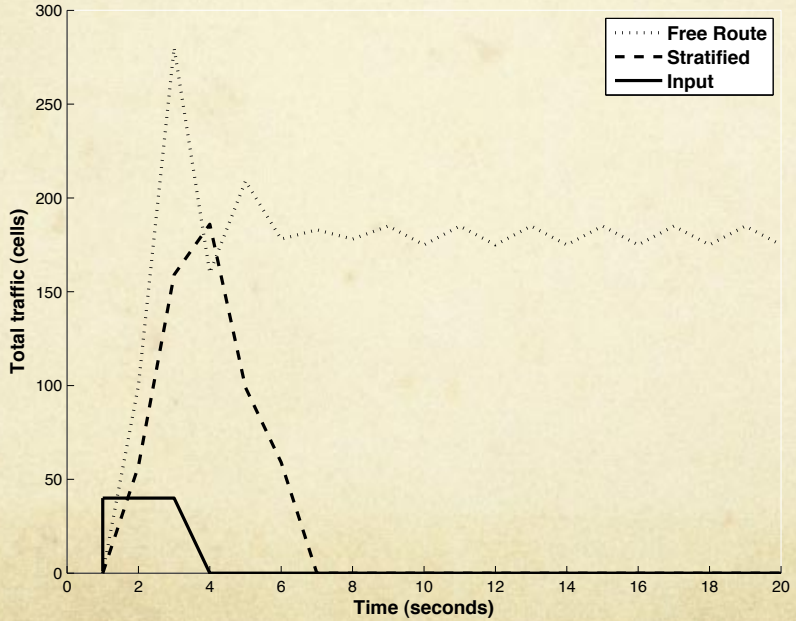

Comparison Topologies

- Anonymity loss: difference with maximum achievable ($\log_2 N$, where N is the total number of circuits in the network
- Overhead factor: number of dummy packets generated per real packet

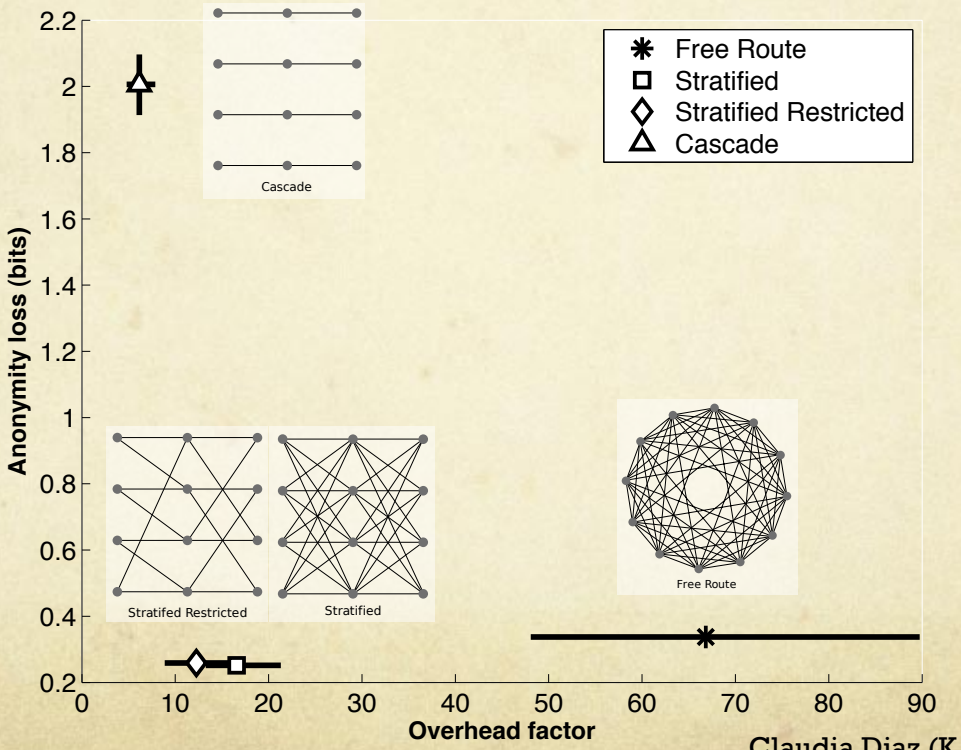

Why Free Routes provide worse anonymity than Stratified

- In Stratified topologies, a node is always in the same position for all the circuits it routes
  - Result: circuits always "mix" in all routers
- In Free Routes, two circuits may pass by the same router and not be "mixed"

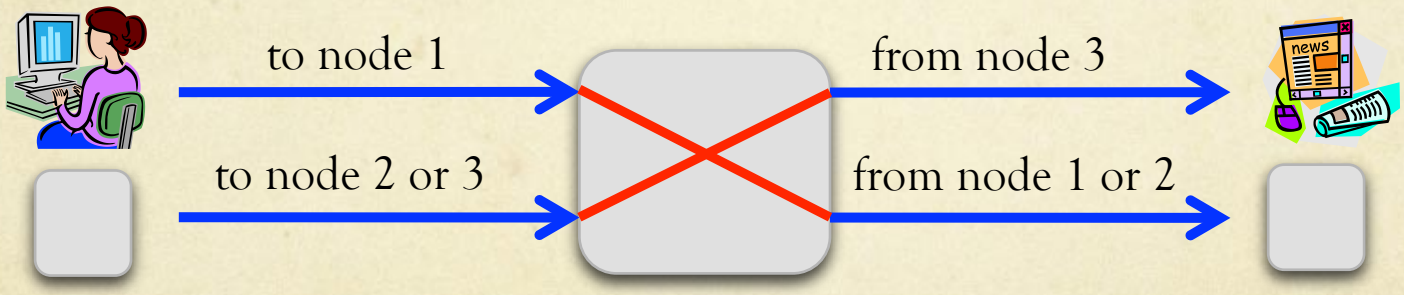

Applying DLP to Tor

- Topology: Tor was originally designed as a Free Route network, but:
  - Only a subset acts as entry
  - Only a subset acts as exit
  - In practice, the topology is close to Stratified
- Padding modes
  - Supports link and circuit padding, but not used in practice
  - Neither padding scheme could be used to support DLP
    - Intermediate nodes must be able to inject padding in circuits that is only detected as padding at the destination
    - AES CTR mode: counter desynchronized if cells added
    - Change to per-cell IV instead of CTR mode (per-stream IV)

Conclusions

- Possibility of reducing overhead by taking advantage of multiplexing
- Impact of network topology for implementing DLP
  - Partitioning of anonymity sets in Cascades
  - Feedback effects in Free Routes (huge overhead, worse anonymity than Stratified)
  - Stratified: best anonymity/overhead tradeoff
- Network scalability
  - Good news: very good anonymity as network grows (except for Cascades)
  - Bad news: high increase in overhead (except for Cascades)
- Applicability to Tor: possible with small modifications
- Open question: resistance of RO-DLP to corrupted nodes
  - Strategies for assigning padding to circuits in a smart way?

Network scalability: anonymity

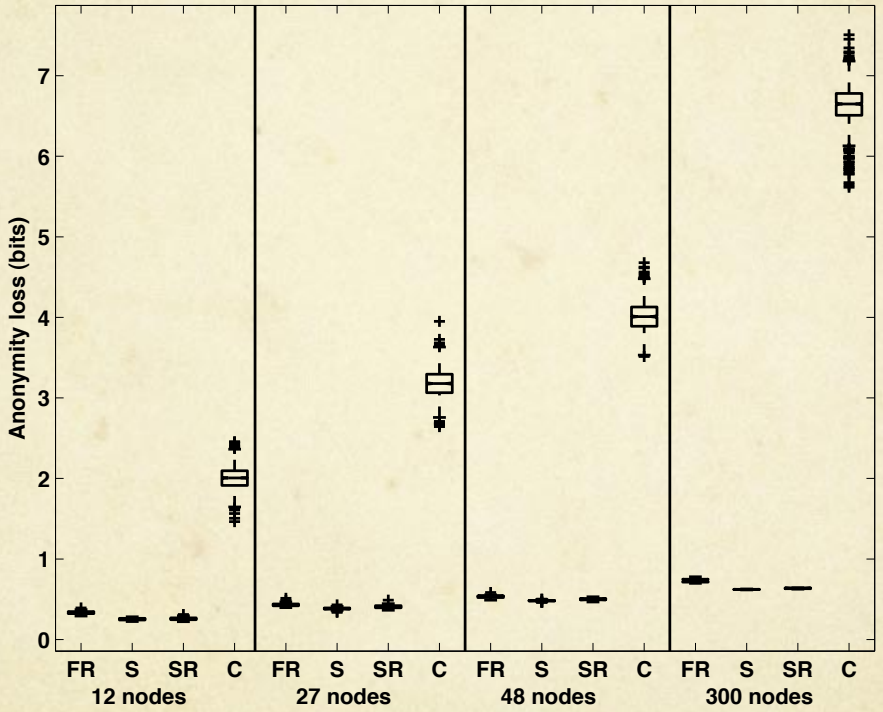

Network scalability: overhead

intra-network links

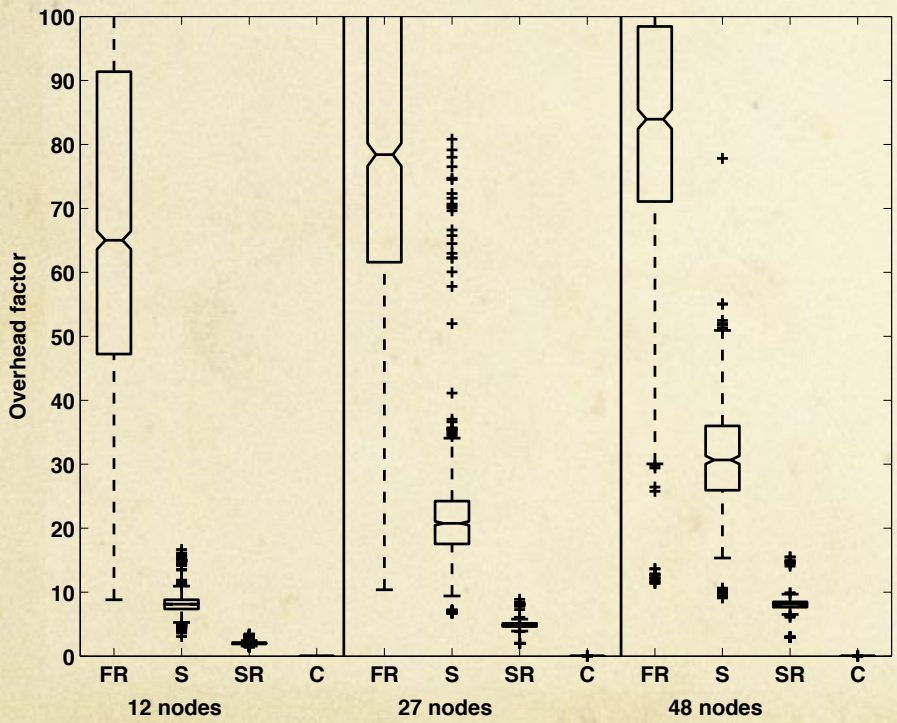

edge links