

MANTIS


A server-anonymity preserving,
searchable, P2P network.

Christopher Soghoian
Johns Hopkins University



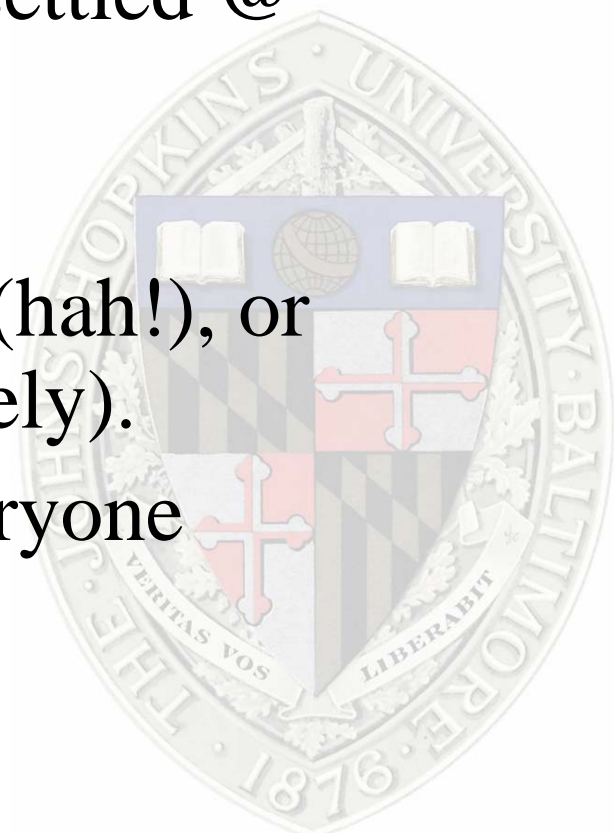
Privacy Online – it never existed

- You never had privacy online.
- Filesharers were always exposed....
- But no one really cared to stop them.
- People had a false sense of security, and thought they were safe from discovery.



The RIAA Attacks!

- The RIAA has been filing lawsuits left and right – against sharers of 1000+ files.
- Nearly 3000 US lawsuits, 486 settled @ ~\$3k.
- This is having a chilling effect.
- People either stop trading files (hah!), or switch to leech mode (more likely).
- Networks cannot survive if everyone leeches.




A Panopticon Society?

- Idea first introduced by British philosopher Jeremy Bentham in 1748.
- A new design for a prison, where prisoners are always under the threat of surveillance, but can never know when they are actually being watched.
- The goal is for for the inmate to internalize the mechanism of surveillance which the building establishes.
- Is the threat of lawsuits creating a Panopticon Internet?




You're helping illegal filesharers?

- There are several non-infringing uses of anonymous p2p networks:
- Those stuck behind the great firewall of china.
- Dissident radio networks.
- Diebold source-code sharers.
- Country Music Fans (explain).




P2P legalities

- Sony Betamax case.
Contributory/Vicarious infringement vs.
legitimate use.
- Napster: file indexes stored on napster
servers. Because of this, they are aiding
the piracy, and can stop it at any time.
- Benefited financially from piracy



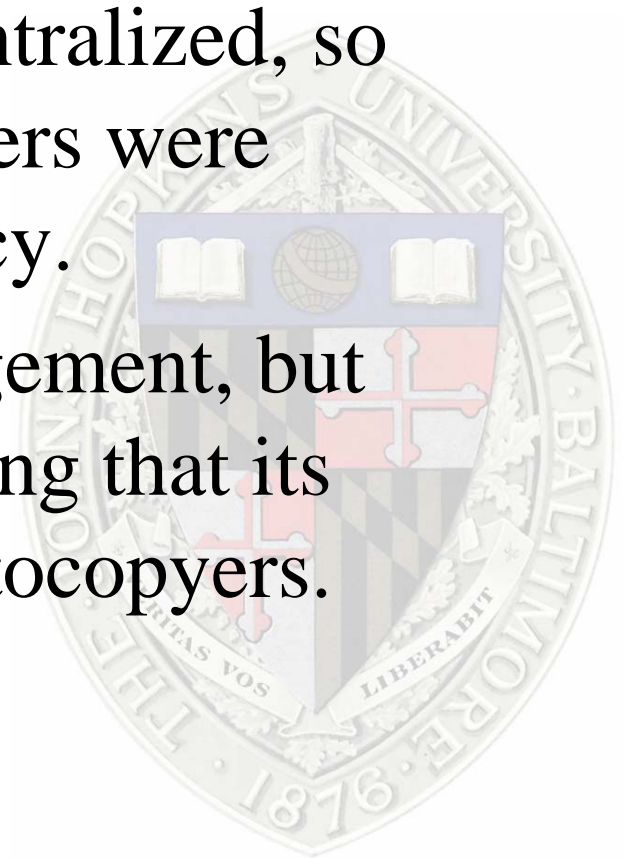
P2P legalities ctd..

- Aimster: Traffic encrypted, so that aimster could not tell what was flowing on their network. Court called this “willful blindness”.
- Tutorials on website showing the technology being used for mp3s.
- Didn't cite any non-infringing uses




P2P legalities

- Grokster/Kazaa: Cited non-infringing uses (e-books, etc).
- Network was completely decentralized, so they didn't know what their users were doing. No ability to block piracy.
- Grokster aware of users infringement, but no different from Xerox knowing that its users copy stuff with their photocopiers.



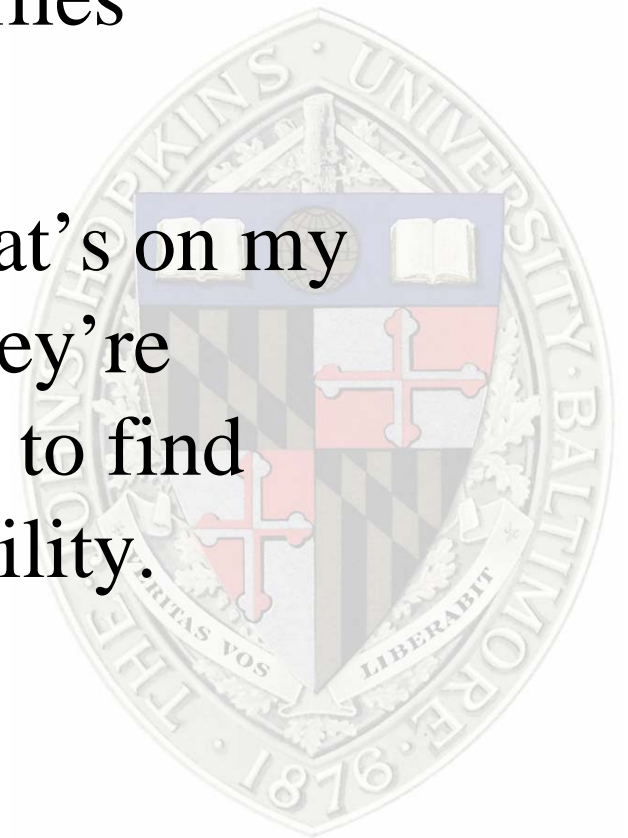
Legal Lessons learned

- You must have non-infringing uses.
- Don't pitch your scheme as a way to share mp3s. At least not publicly.
- Decentralize the network, so you retain no control.
- Don't make any money off it.
- No auto updates (german police)



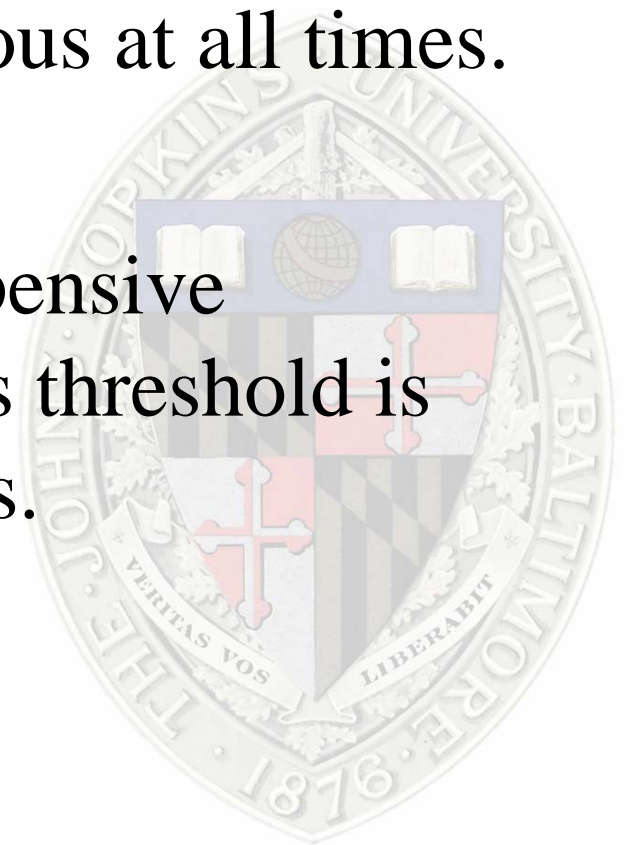
Different Methods for Anonymity

- Freenet: “Sure, I’m on the network, but I don’t know what files are on my computer” – good for popular files
- Mantis & Friends: “I know what’s on my computer, but no one knows they’re talking to me.” – good for hard to find files. Dollar Store style availability.




Prerequisites

- For a P2P network to survive, it has to meet the following requirements:
 - 1. Sharers must remain anonymous at all times.
 - 2. It must be searchable.
 - 3. It must not be excessively expensive (CPU/Bandwidth). Note that this threshold is different for clients/servers/peers.



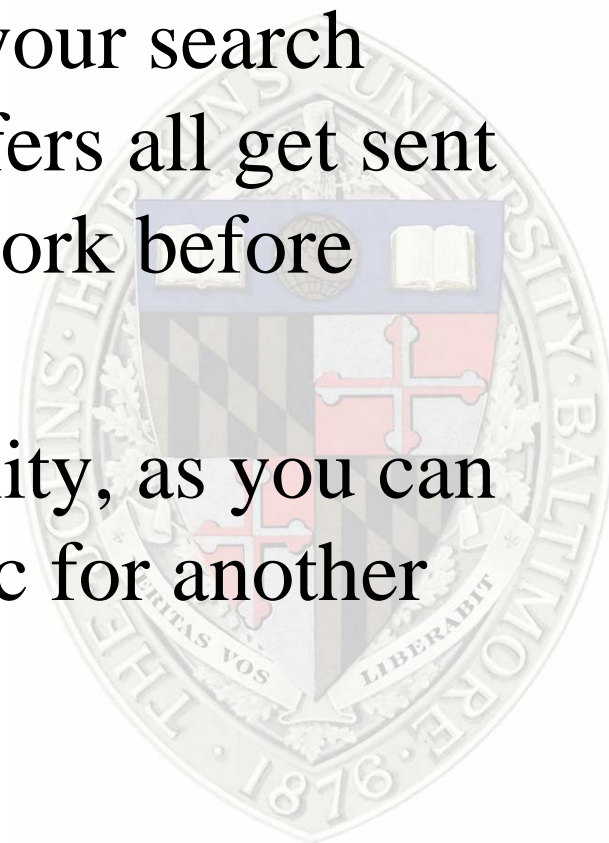
Pre-req's continued...

- People will give up bandwidth to download free stuff (duh).
- People will give up bandwidth to share files (as napster and others have shown)
- Peers are extremely sensitive to bandwidth usage (they're selfish).
- 4. It must be fast.



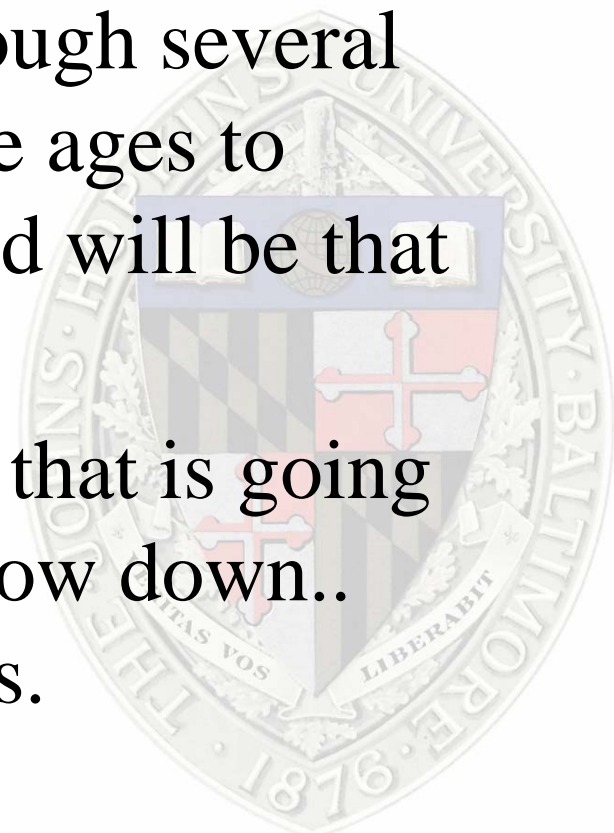
Initial Designs/MUTE

- We originally came up with an idea very similar to MUTE (mute-net.sourceforge.net).
- Simply put, MUTE insures that your search requests/responses and file transfers all get sent through other clients on the network before going to you.
- This gives you plausible deniability, as you can just claim to be forwarding traffic for another host.




MUTE Problems

- .Avi files are large.
- Peers are selfish.
- Transferring a large movie through several peers will annoy them, and take ages to complete. Your download speed will be that of the slowest peer on the link.
- For every additional download that is going through a peer, all the others slow down..
Never ending cycle of slowness.




Other Schemes

- Crowds. – Anonymizes client.
- Hordes – Multicast – not realistic/uses too much relay bandwidth.

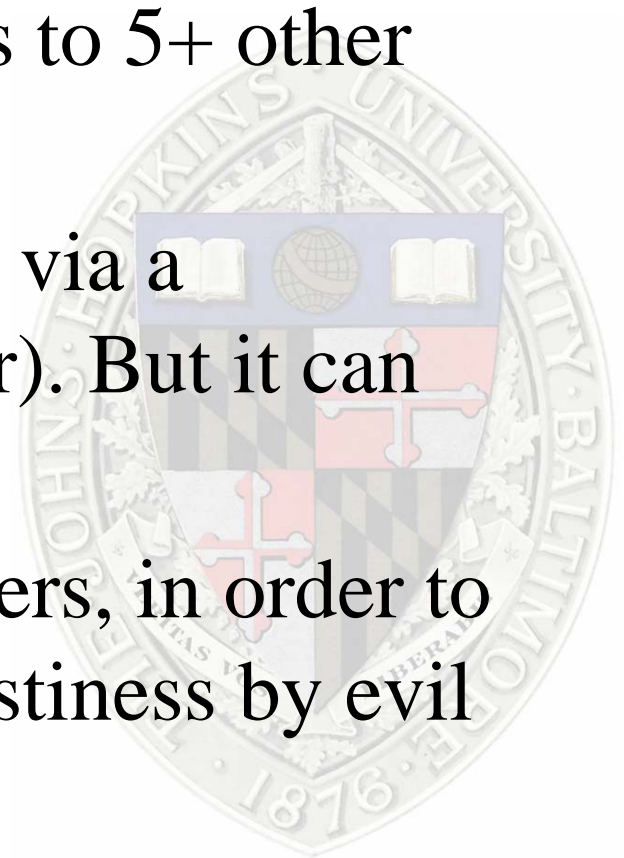


Mantis



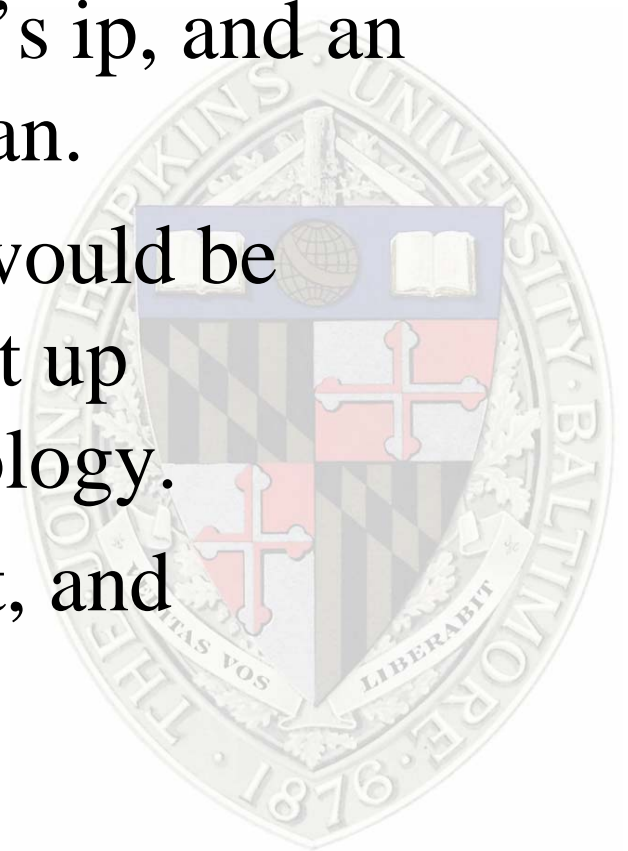
Mantis - Nodes

- Mantis nodes have 3 simultaneous roles: client, server and peer.
- Each node maintains connections to 5+ other nodes.
- Node discovery is typically done via a directory server (called a Blender). But it can be done in other ways
- Nodes should use multiple blenders, in order to avoid Sybil attacks, and other nastiness by evil Blenders.



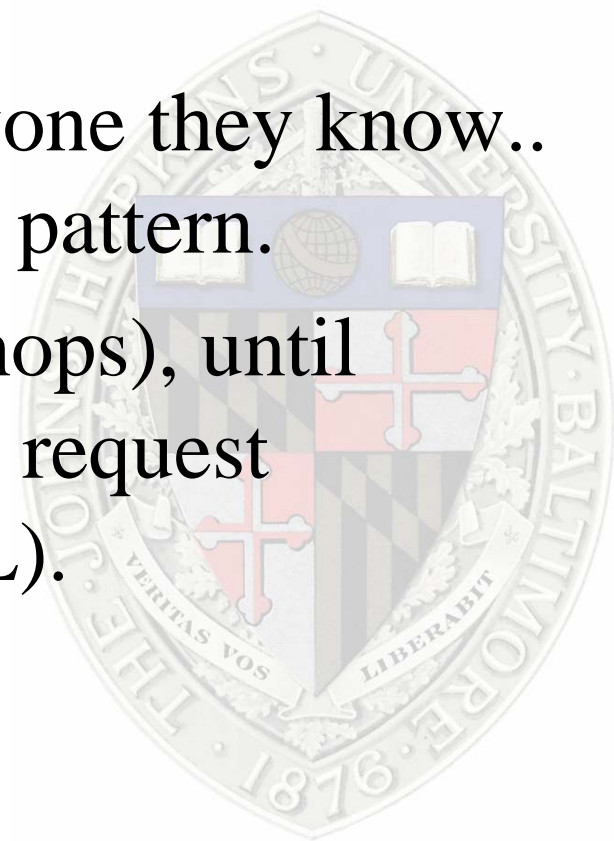
Blenders

- Nodes register themselves with one or more blenders.
- The blender only stores a node's ip, and an “accepting connections” boolean.
- Any more information, and it would be possible for the Blender to built up knowledge of the network topology.
- Nodes request a connection list, and connecto a random subset.




Mantis - searches

- When a node wishes to search for a file/service, it passes that search request on to its neighbors.
- They in turn pass it on to everyone they know.. Searches go out in a ripple like pattern.
- This repeats for a while.. (~ 5 hops), until nodes start dropping the search request (probabilistic dropping vs. TTL).



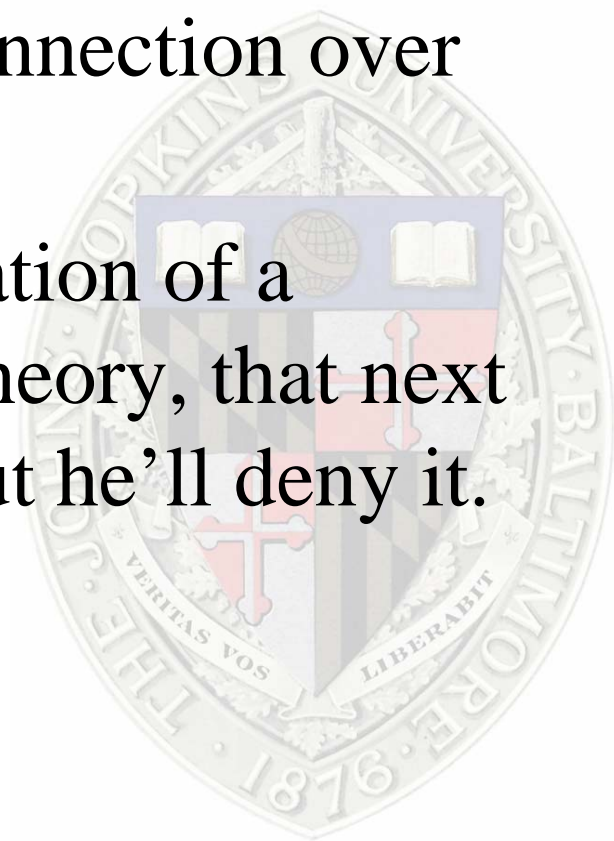
Mantis - searches

- When a node receives a search request, and it indeed has the file, it sends a response back up the tree.
- Search requests/responses carry unique ID's, which all nodes use to keep track of where to forward search replies, and other communications.




Mantis - Connections

- Using the combination of the uniquely generated search and response ID's, both parties are able to establish a connection over the link.
- No node knows the final destination of a message. Just the next hop. In theory, that next hop could be the destination. But he'll deny it.




File Transfers


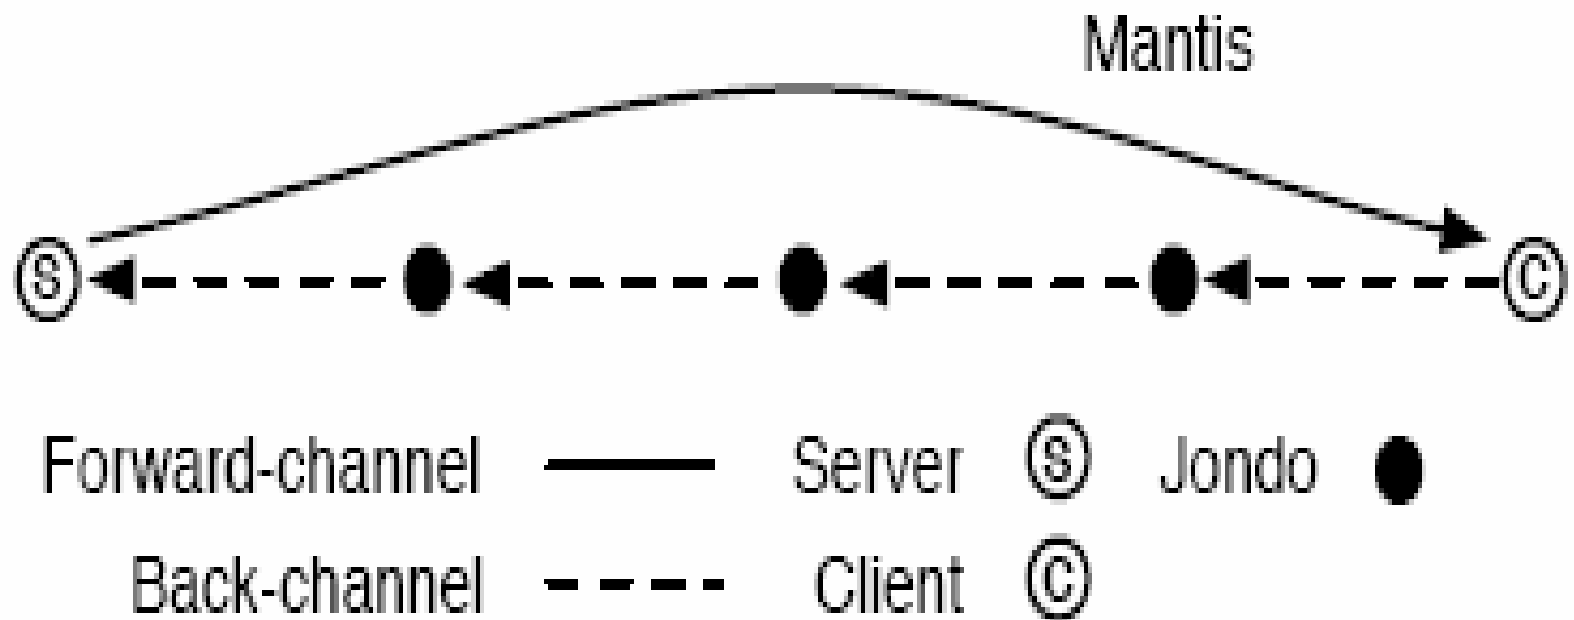
- A client has searched for something, gotten a few interesting responses, and now wishes to download one of the files.
- The client shares his IP and a listening port with the server.



And now for the fun bit!


- The server now sends the file to the client over a direct UDP channel to the client (i.e not passed through all the other nodes).
- He spoofs his source address when doing so.
- Control data (to arrange for retransmission of lost packets, etc) is sent over the anonymous back channel of other nodes.

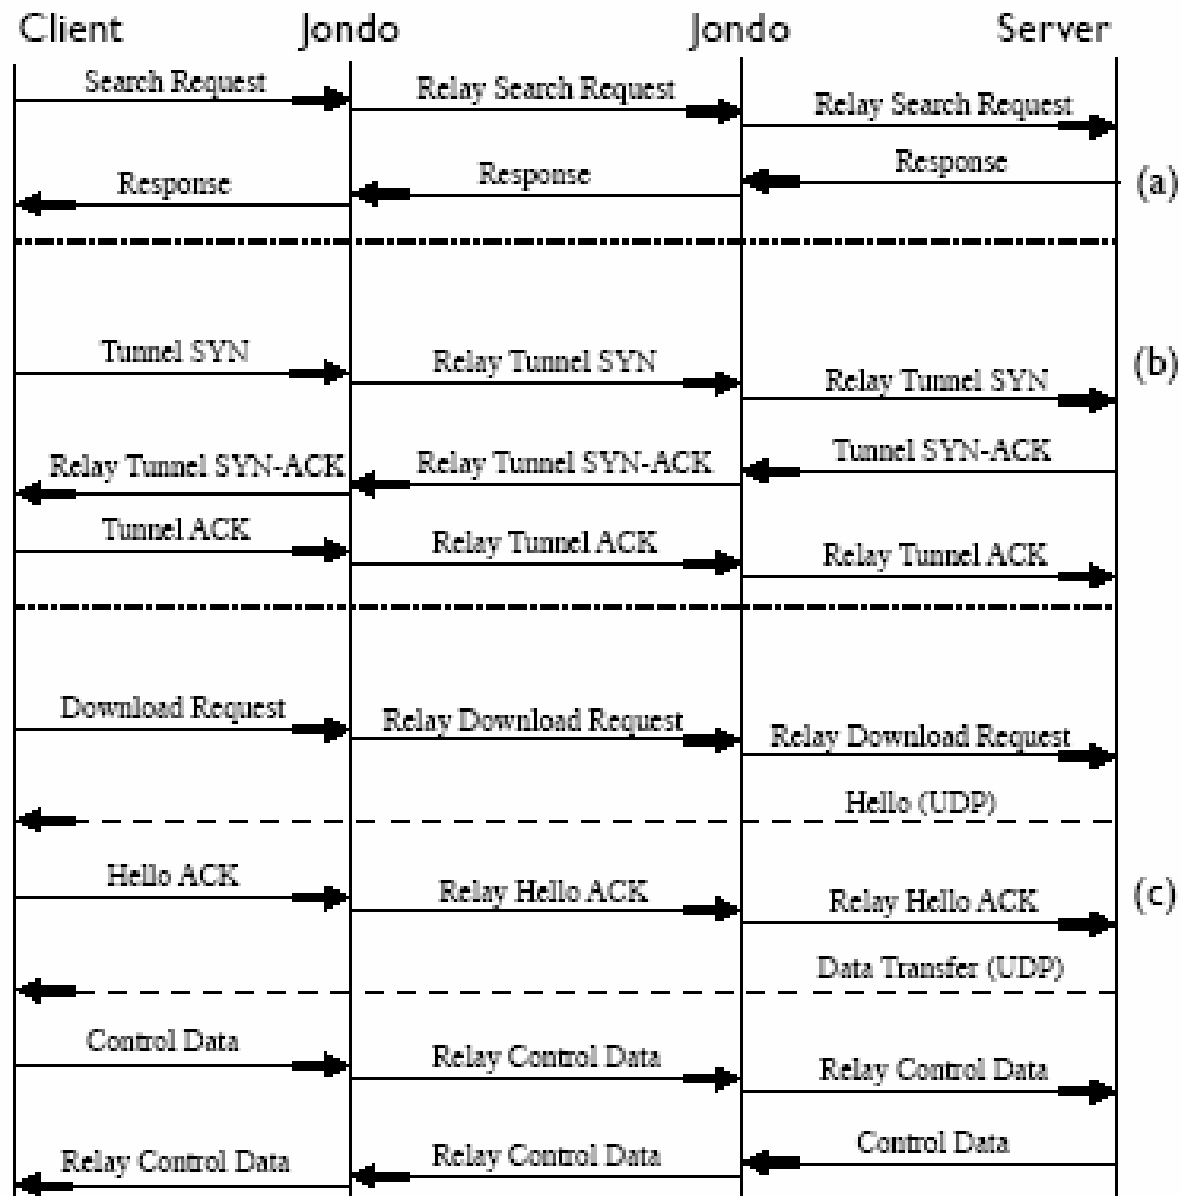




More of the fun bit.

- Jondo nodes only have to pass search requests/responses and control data.
- 2 DSL speed hosts can communicate through a sea of modem users, and still transfer a file at high speed.
- The server remains anonymous.





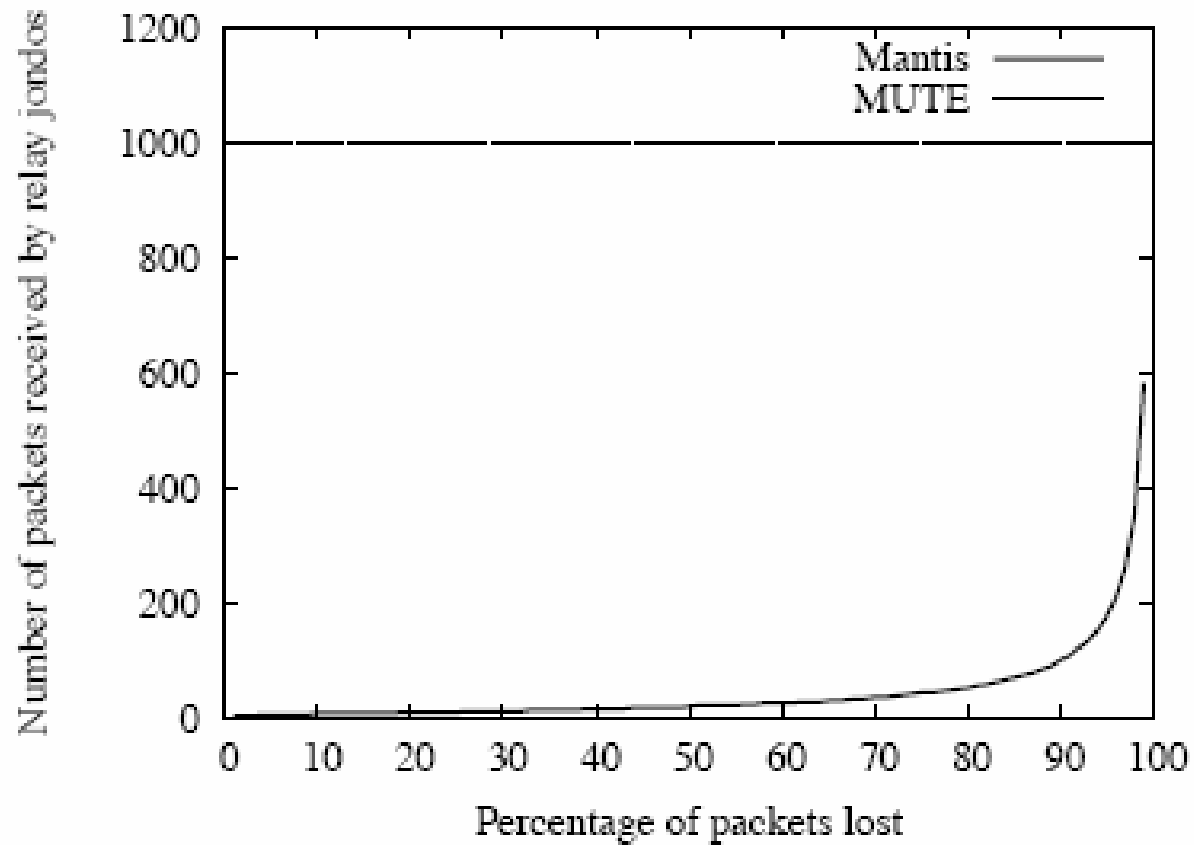
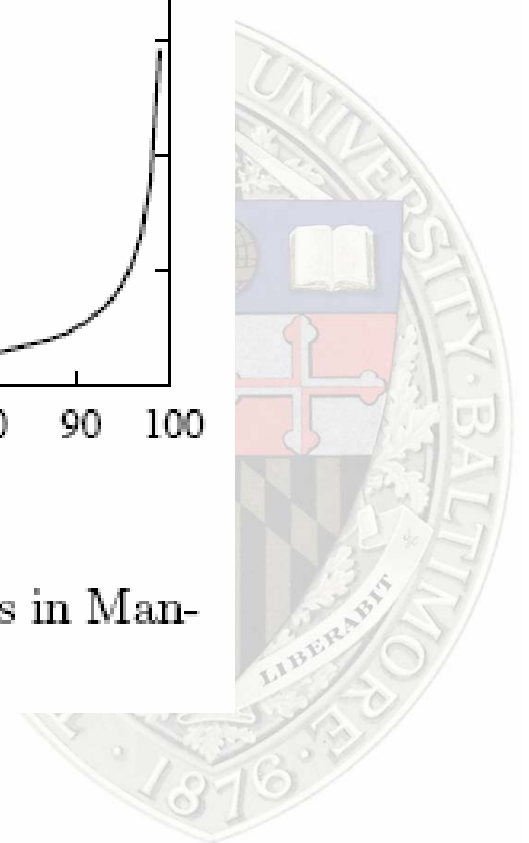



Figure 3: Packets Forwarded by Relay Jondos in Mantis vs. MUTE



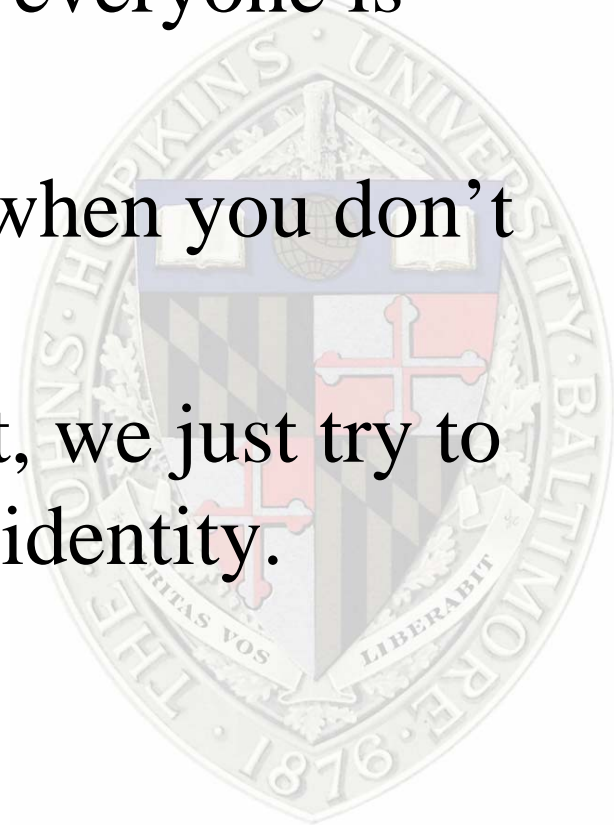
Issues?

- Q: Can't middlemen learn of the client's IP/port.
- A: Not usually. Communication between nodes is encrypted at every hop, with an additional layer of encryption between client and server.




Issues

- A middle-node can man in the middle the connection, but there is nothing we can do to stop this. PKI's don't work when everyone is anonymous.
- You can't authenticate someone when you don't know who they are.
- We don't aim to protect the client, we just try to make it a bit difficult to learn his identity.



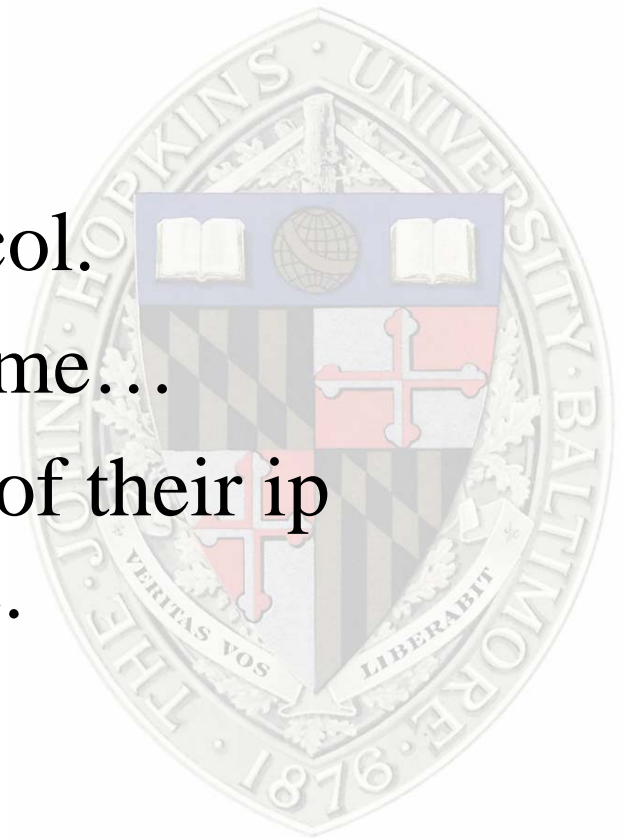
Issues

- Q: Are you subject to Timing Attacks?
- A: We hope not. Nodes use random-ish delays before returning data, which allows them to claim that they are just forwarding a message for someone else, and are not the originator.




Issues?

- Q: If ISP's start egress filtering, and don't allow spoofed packets out, aren't you SOL?
- A: Good Question!
- Sliding Egress detection protocol.
- Revealing info one octet at a time...
- Servers can choose how much of their ip they are willing to make public.




Egress issues

- With good egress in place, you can't pretend to be an AOL user, but you can at least hide amongst the other users on your ISP's network.
- If the RIAA forces ISP's to implement widespread egress filtering, we all win anyway, as DDoS will be made much harder.



Issues...

- Reconnecting?
- As the size of a file increases, as well as the number of nodes between the client and server, the chance that some node is going to logout increases.
- We introduce a way for the server to reconnect anonymously to the client.



Show me the code!

- A technical report describing our scheme is available at:
- <http://spar.isi.jhu.edu/~chris/>
- Code has been written in a network simulator. We're currently porting it to the 'real world', and hope to release it in a few months.
- Other Questions?

