



AN ANALYSIS OF THE DEGRADATION OF ANONYMOUS PROTOCOLS

Matthew Wright*

Micah Adler*

Brian Levine*

Clay Shields†

* - Dept. of Computer Science, University of Massachusetts,
Amherst

† - Dept. of Computer Science, Georgetown University

[ANONYMOUS COMMUNICATIONS]

- Hiding your IP address on the Internet
- Prevent tracking, monitoring
- Protocols:
 - Crowds [Reiter, Rubin '98]
 - Onion Routing [Goldschlag, Reed, Syverson '98]
 - Mix-nets [Chaum '81] and variations
 - Hordes [Shields, Levine '00]

[DEGRADATION]

- Simple passive logging attack
- Malicious, collaborating members
- Reduce anonymity over time
 - Long-term or repeated sessions
- Quantify chance of correct ID

[PREVIOUS WORK]

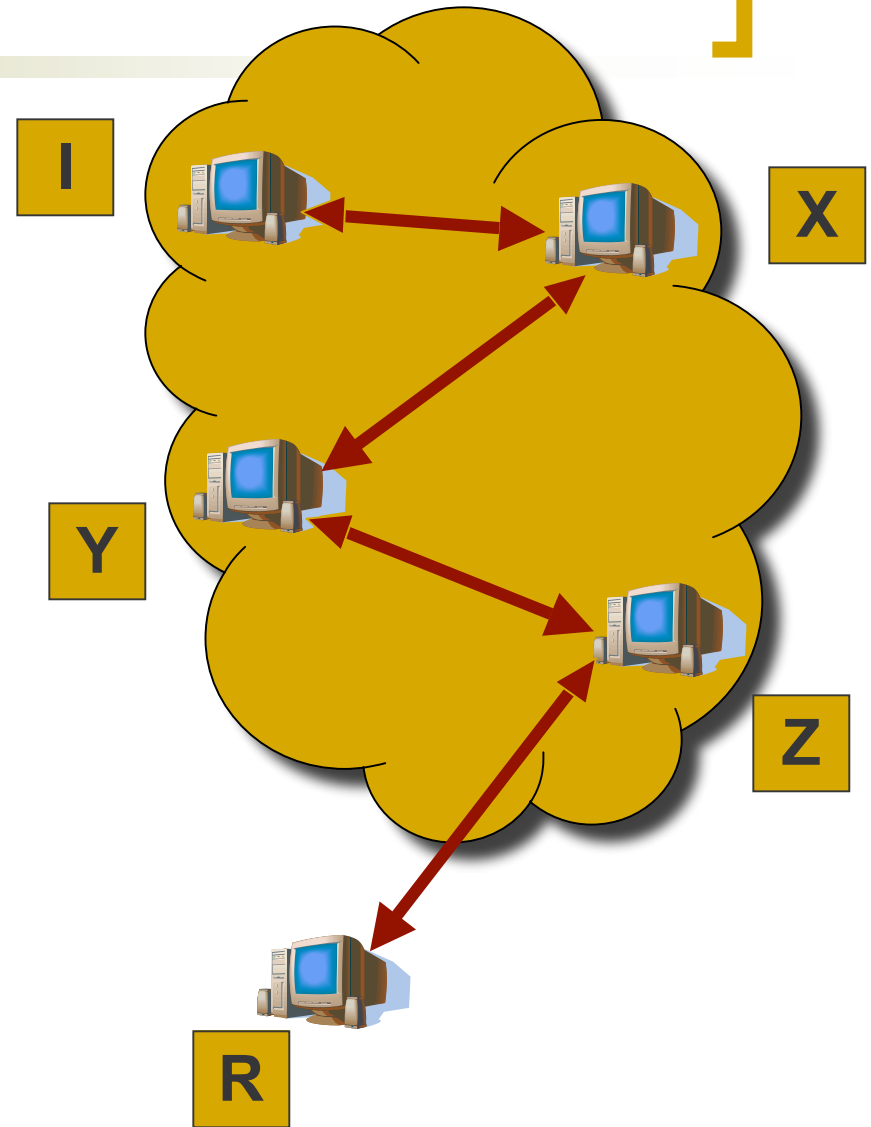
- Crowds [Reiter and Rubin, '98]
 - Identified the attack (against Crowds only)
 - Analysis: initiator will be identified
 - Defense: semi-static paths
- Towards an Analysis [Syverson et al, '00]
 - A single-round attack against Onion Routing
 - Probability of single-round identification

[OUR CONTRIBUTIONS]

- Can all existing protocols be attacked in this way?
 - Yes. We present a proof of this in the paper.
- How long does the attack take?
 - We present analytical results for several protocols
- Are there any defenses?
 - Yes, but there are tradeoffs

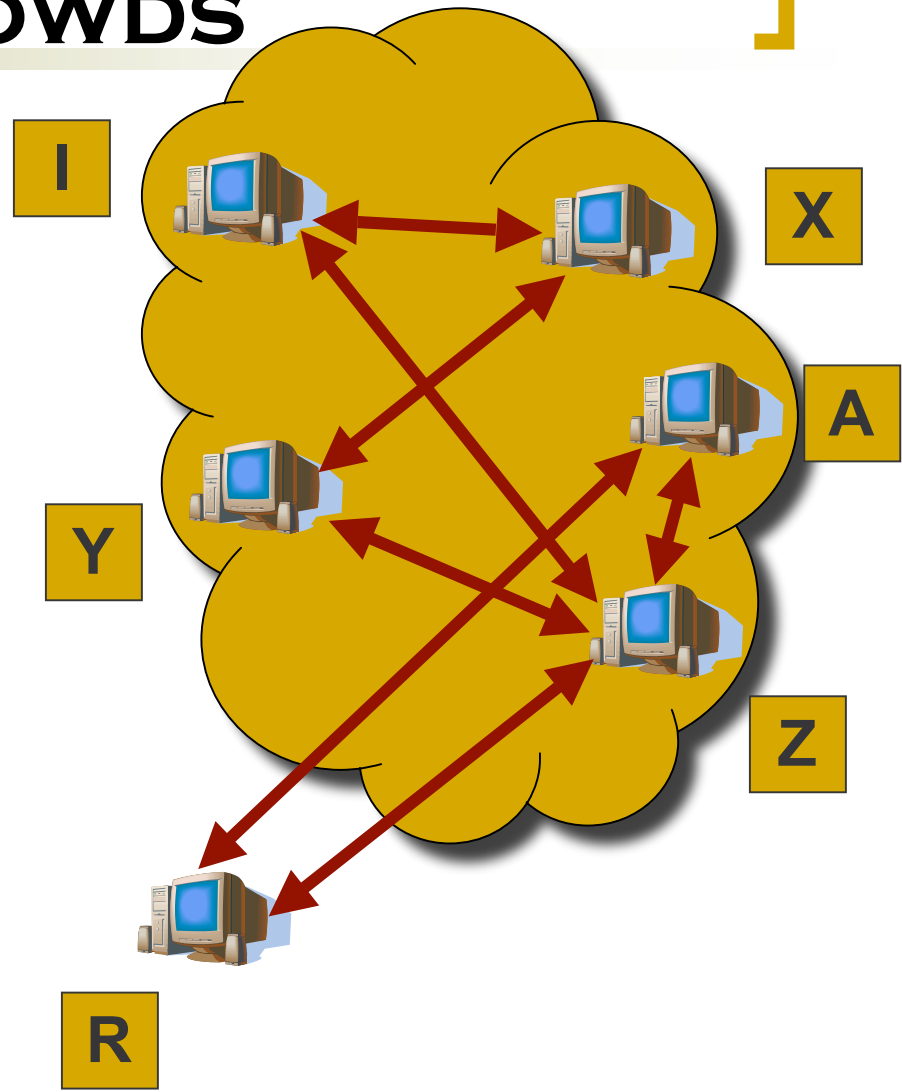
[CROWDS]

- Reiter, Rubin 1998
- P2P Anonymity
 - Both anonymous users and proxies
- Selecting paths
 - Weighted coin and spinner



ATTACKING CROWDS

- Attacker sees *session-identifying information*
 - Responder's IP address
 - Cookie, login name, specific content
- Paths change



ATTACKING CROWDS II

- Log the node before the attacker
 - Attacker is 1st proxy: Initiator with probability **1**
 - Attacker is not 1st: all nodes with probability **1/n**

node	I	X	Y	Z
times seen	41	18	24	17

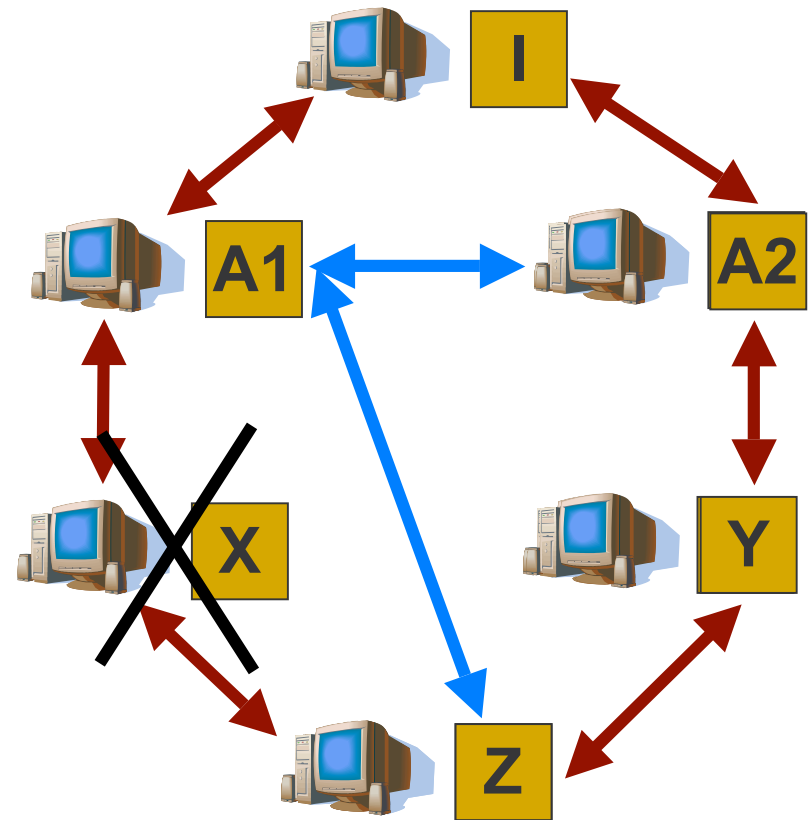
- Multiple attackers
 - **n** total nodes, **c** attackers
- **$O(n/c \log n)$** rounds

ATTACKING IN GENERAL

- Attack applies to any protocol for anonymity, provided that:
 - *Active sets* (i.e. paths) of proxies change
 - Uniformly random selection of active sets
 - There exists a position of attackers:
 - see the initiator send messages in the session
 - determine the session information
- Again, Initiator appears more often
 - Law of Large Numbers

[DC-NET]

- Chaum '85
- Ring-based configuration
 - 2 attackers = 2 virtual circles
 - Leave and re-join in a random place in the ring
 - Trap the initiator in $O(n \log n)$ rounds



OTHER PROTOCOLS

- Onion Routing
 - First and last node on path
 - Crowds-like attack: $O((n/c)^2 \log n)$
- Mix-Nets
 - Timing analysis stopped
 - Entire path required: $O((n/c)^L \log n)$
- DC-Net, complete graph configuration
 - To get any information that would distinguish the initiator, all $n-1$ nodes are required

[EXAMPLE ATTACK TIMES]

- Suppose:
 - 15 minute rounds (session traffic every round)
 - $n = 1000$, $c = 100$, $p_f = .75$, $L=3$
 - After each round begins, a random attacker leaves the protocol and rejoins
 - Confidence of .999
- Crowds: 550 rounds = 5.7 days
- Onion Routing: 5500 rounds = 57 days
 - OR with Crowds-like path lengths: 23 days
- Mix-Nets: 56000 rounds = 82 weeks

[DEFENSES]

- Variable path lengths
 - Security of shortest path length
 - Performance of longest path length
- Don't change paths
 - Selfish solution
 - Vulnerable to traceback
- Pick trusted paths
 - Decreases size of anonymity set

CONCLUSIONS

Protocol	Crowds	DC-Ring	Onion Routing	Mix-Nets
Rounds	$O(n/c \lg n)$	$O(n \lg n)$	$O((n/c)^2 \lg n)$	$O((n/c)^L \lg n)$

- Attack affects all known systems of anonymous communications
- Apparent trade-off between performance and security

[WHY RANDOM PLACEMENT?]

- Suppose deterministic placement
- Node ID or IP address
 - Forge ID or IP
 - Choose location in the ring
 - **$O(\log n)$**
- Last vacated position
 - Wait for prime location to open up

ATTACKING IN GENERAL

- Attack applies to any protocol for anonymity, provided that:
 - *Active sets* (i.e. paths) of proxies change
 - Uniformly random selection of active sets
 - There exists a position of attackers:
 - see the initiator send messages in the session
 - determine the session information
- Again, Initiator appears more often
 - Law of Large Numbers

[EACH ROUND]

- Total ordering on the packets, Π
- Initiator's transmission Π_I
 - detected by attackers
 - linked to the session
- Possible initiators
 - Clearly, the real initiator, I
 - Other nodes that send messages for I
- Equal probability of "detecting" any node
 - But initiator guaranteed to be sending

ONION ROUTING

- An Onion

- Initiator chooses path, and builds an onion
- Layered encrypted of data using the public key of each proxy in the path


$$\{Y, \{Z, \{R, \text{data}\}_{K_{Z+}}\}_{K_{Y+}}\}_{K_{X+}}$$

- Sending the onion

- I \rightarrow X: $\{Y, \{Z, \{R, \text{data}\}_{K_{Z+}}\}_{K_{Y+}}\}_{K_{X+}}$
- X \rightarrow Y: $\{Z, \{R, \text{data}\}_{K_{Z+}}\}_{K_{Y+}}$
- Y \rightarrow Z: $\{R, \text{data}\}_{K_{Z+}}$
- Z \rightarrow R: data

ATTACKING ONION ROUTING

- Attackers at the ends
 - Use simple timing analysis to determine that they're on the same path



A1	A2
3:12:20.92	3:12:20.98
3:12:37.21	3:12:37.28
3:12:49.36	3:12:49.41

- Crowds-like attack
- $O((n/c)^2 \log n)$ path resets

[SETUPS]

- Cover Broken
 - The FBI is tracking visitors to web site **X**
- The Setup
 - **S** initiates a connection through proxy **V** to **X**
 - Every path change goes first through **V**
- The Fall
 - It appears to the FBI that **V** is the initiator
- Not Foolproof
 - **V** keeps logs, cooperates with FBI, captures **S**

[MIX-NETS]

- Mixing
 - Reordering messages
 - Dummy messages
 - Delay
- Stops timing attacks
 - OR attack no longer works
 - Need the entire path to trace the message
- **$O((n/c)^L \log n)$** path resets