

Breaking the Collusion Detection Mechanism of MorphMix

Parisa Tabriz and Nikita Borisov
University of Illinois, Urbana-Champaign

Outline

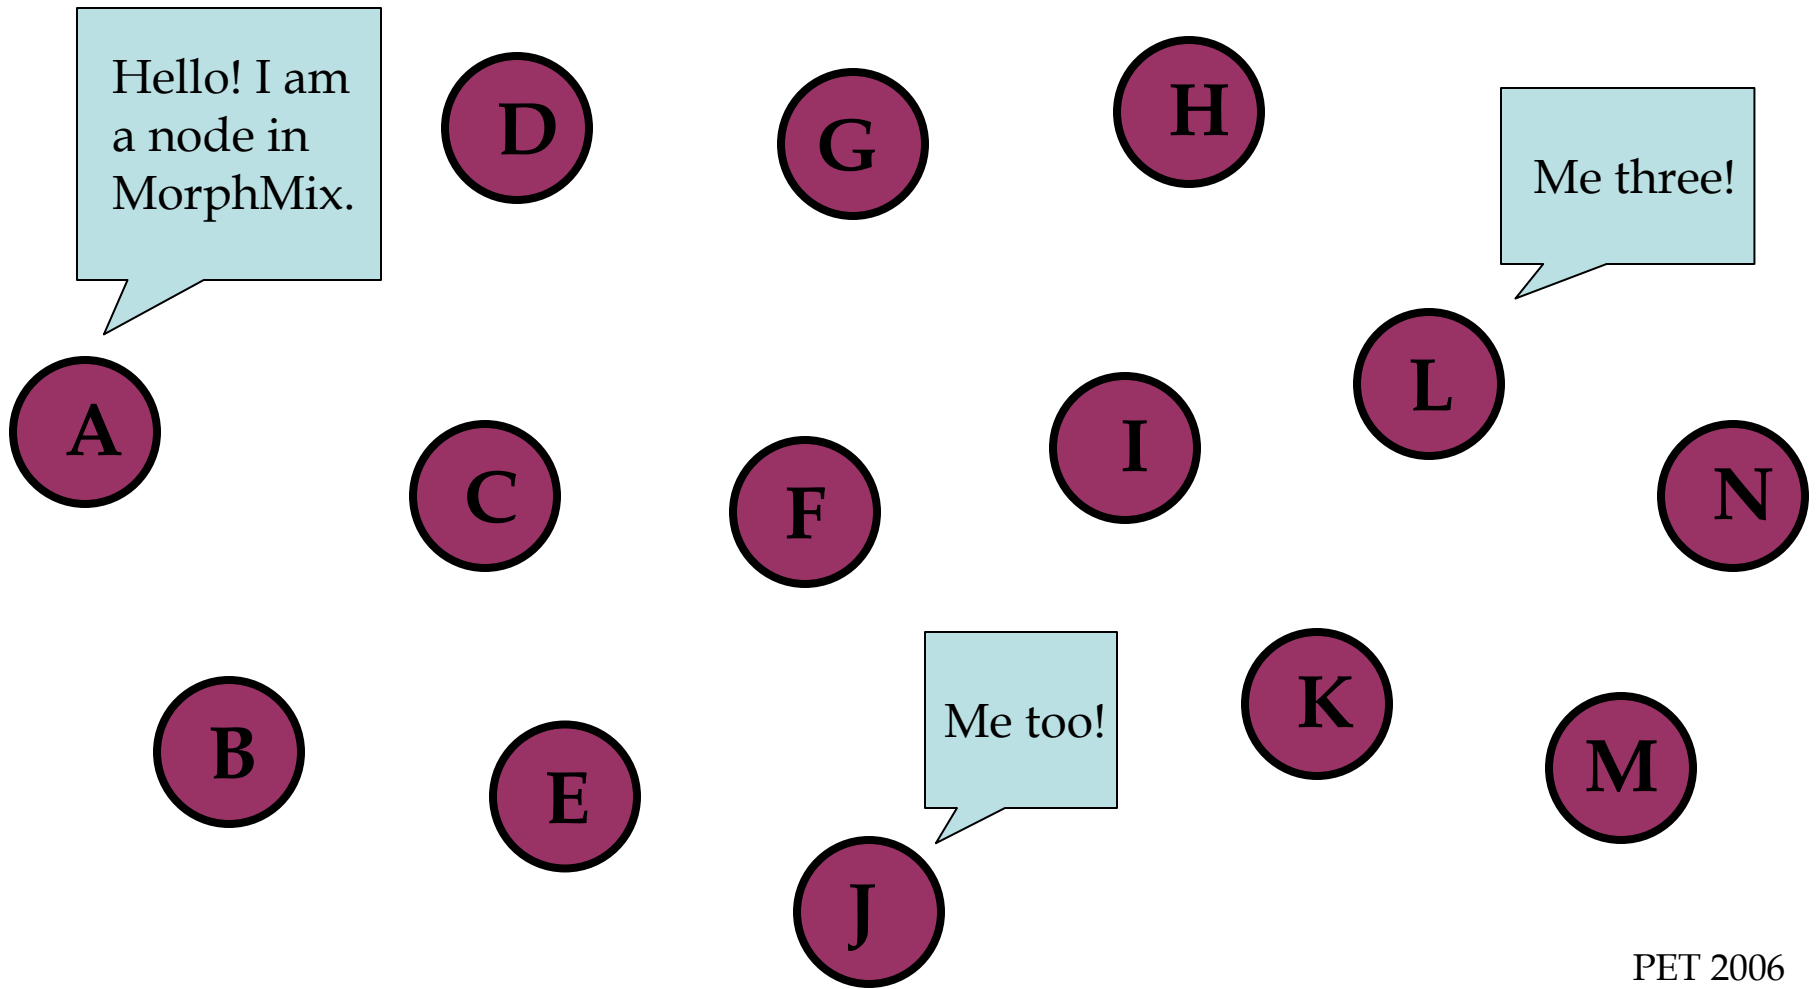
- I. The Scalability of MorphMix
- II. Collusion Detection in MorphMix
- III. Attacking the CDM
- IV. Results
- V. Countermeasures
- VI. Final Thoughts

Anonymous Networking


Morph who?

- Tor [Dingledine'04]
 - We love you, but...
 - Not P2P.
 - Scalability limitations.
- Morphmix [Rennhard'02]
 - P2P anonymous networking overlay.
 - No centralized component.
 - Nodes only need a local view.
 - Collusion Detection Mechanism.


This is MorphMix. Now in a new shade of purple.




More MorphMix.



Tunnel Construction




Tunnel Construction

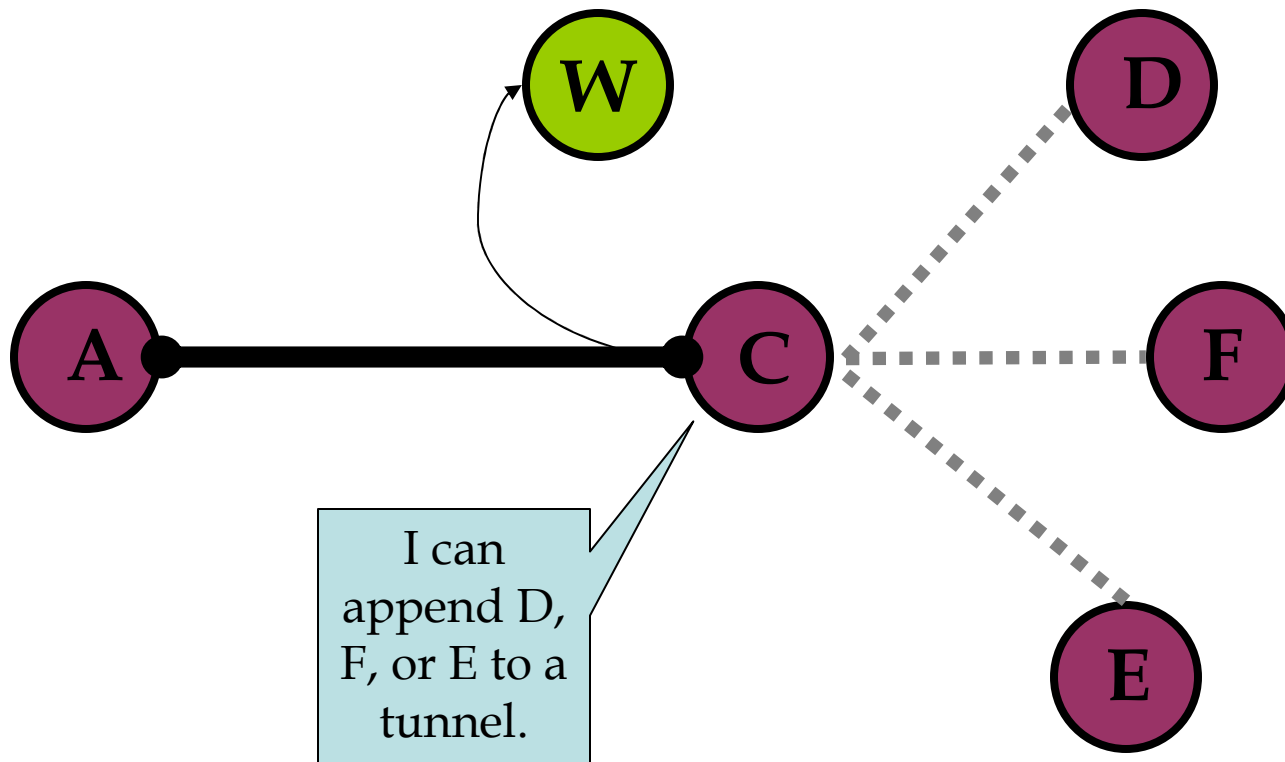


Tunnel Construction


Please
append a
node to
this tunnel
using
witness W.




Tunnel Construction



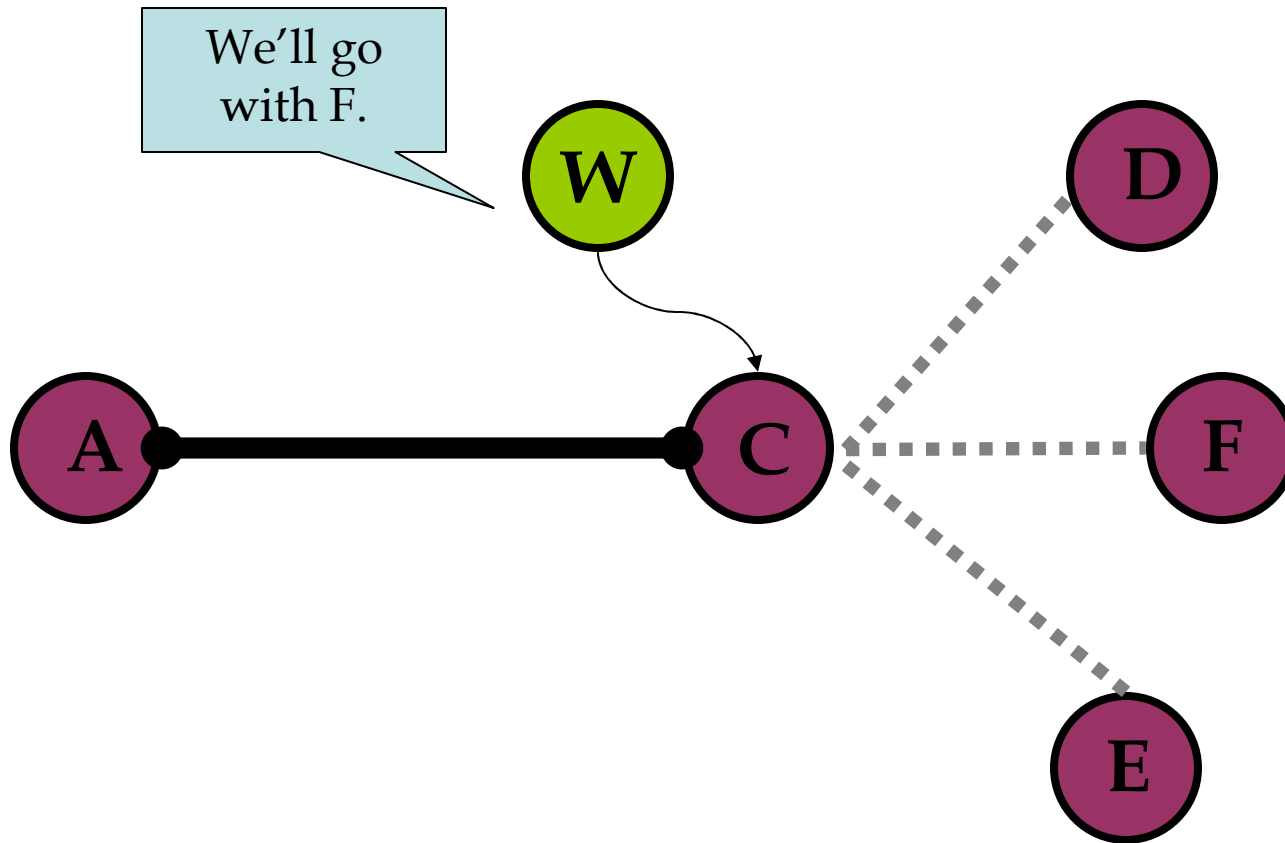
Tunnel Construction




Tunnel Construction




Tunnel Construction




Tunnel Construction




Tunnel Construction



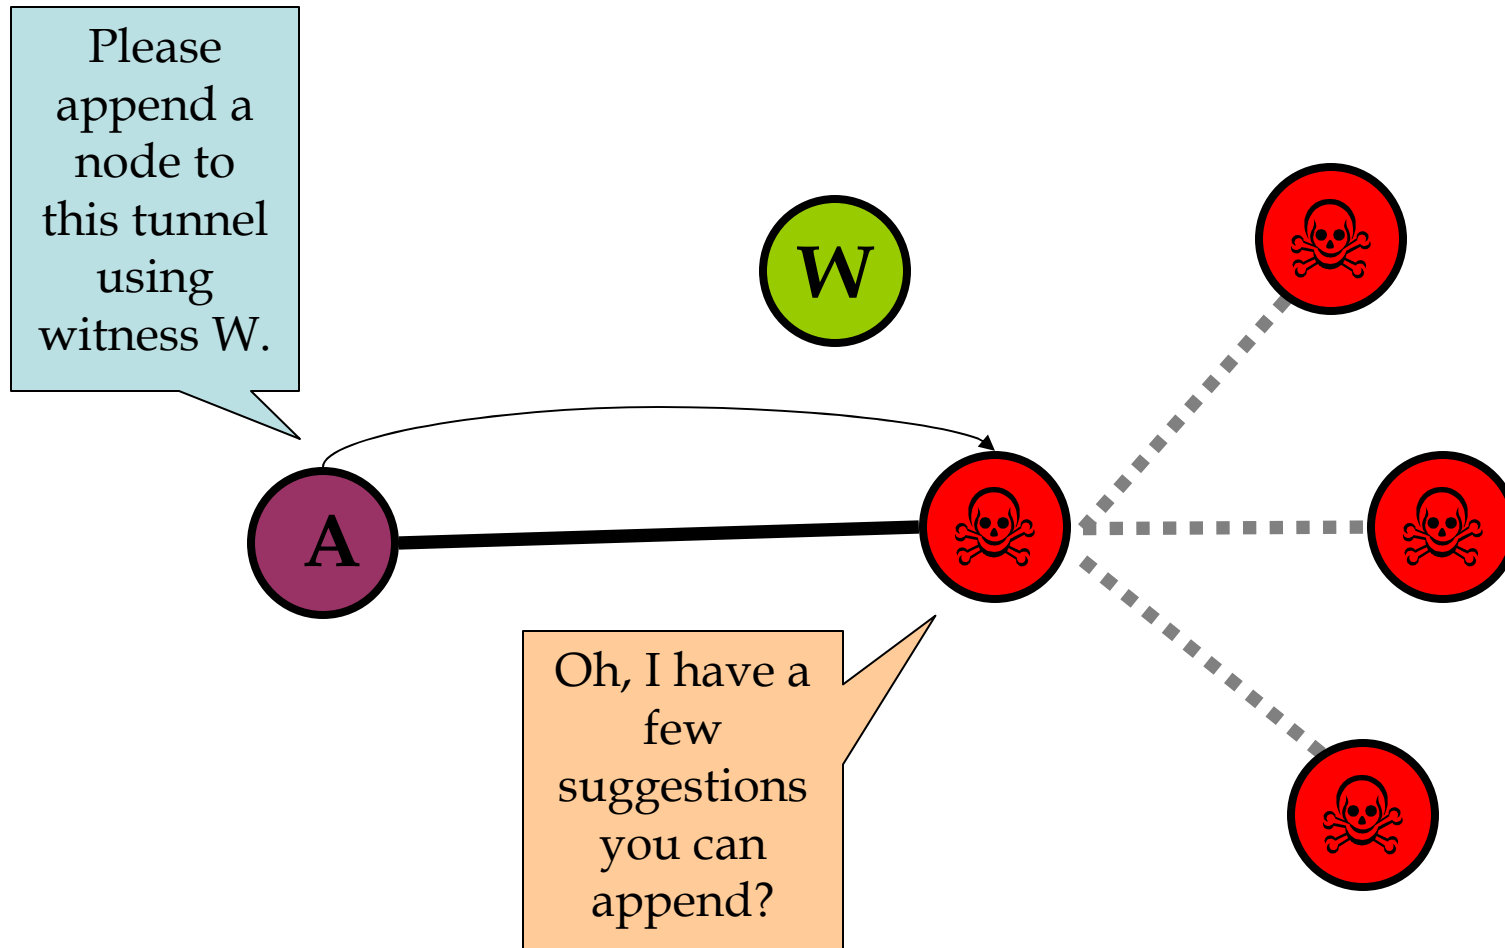
Tunnel Construction



Tunnel Construction



What can go wrong...



Collusion Detection Mechanism

- Assumes it is more difficult to control multiple IPs in many /16 subnets than many IPs in one /16 subnet.

Seems reasonable.

- Assume attackers will provide all or many malicious nodes in their selections.

Seems reasonable.

- Assume attackers will provide a random selection of malicious nodes in their selections.

We'll get to this later...

Collusion Detection Mechanism

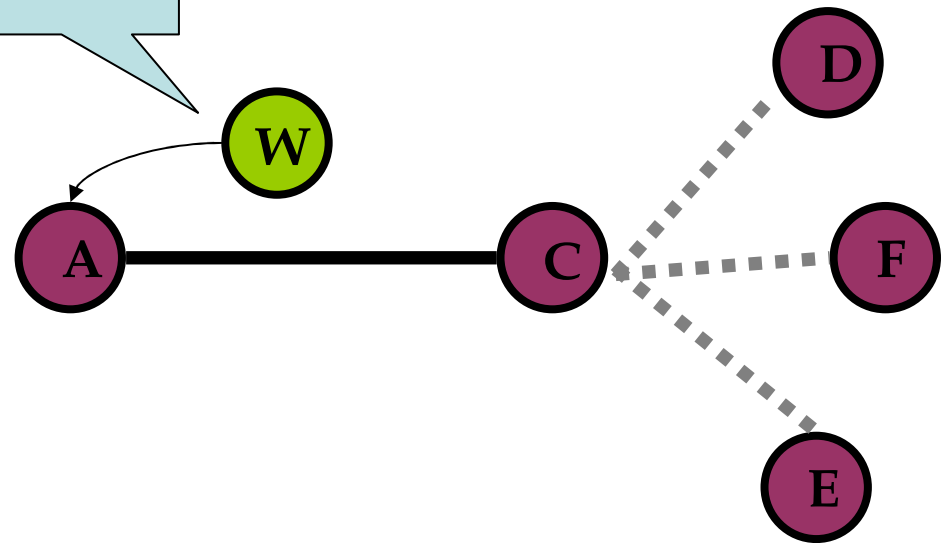
Definitions:

- Extended Selection (ES_L): 16-bit IP prefixes of the sending node and every node in the selection.
- Extended Selection List (L_{ES}): Stores recent extended selections.

Collusion Detection Mechanism

B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

C says you can append D, F, or E.



LE_S of Node A

Collusion Detection Mechanism

C	D	F	E
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

LE_S of Node A

1. Compare a new extended selection with every extended selection in a node's LES.
2. If they have at least one element in common, that extended selection's elements will contribute to the correlation.

Collusion Detection Mechanism

C	D	F	E
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

LE_S of Node A

3A. Count the number of elements that appear more than once and store this value in m .

$$m = 2$$

Collusion Detection Mechanism

C	D	F	E
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

3B. Count the number of elements that appear once and store this value in u .

$$u = 14$$

LE_S of Node A

Collusion Detection Mechanism

C	D	F	E
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

LE_S of Node A

4. The correlation, c , for the extended selection is: $\frac{m}{u}$
5. If c is beneath the correlation limit, accept the tunnel. Otherwise, reject the tunnel.

$$c = 1/7$$

MorphMix Attacker

- Attacker Goal:
 - Link connection initiator with an outgoing network stream.
 - Control at least the first intermediate and last node in a tunnel.
- Attacker Model:
 - Active, internal adversary.
 - Controls nodes in a percentage, C , of all unique subnets in MorphMix.
 - We assume C can range from 5% to 40%.

Attack Intuition

- Every node's correlation limit is based on local knowledge stored in its L_{ES} .
- To achieve low correlation, limit the overlap between an extended selection and elements in a target's L_{ES} .
- Since the organization and size of a node's L_{ES} is known, can model and manipulate any node's L_{ES} individually.

Intelligent Selection Attack

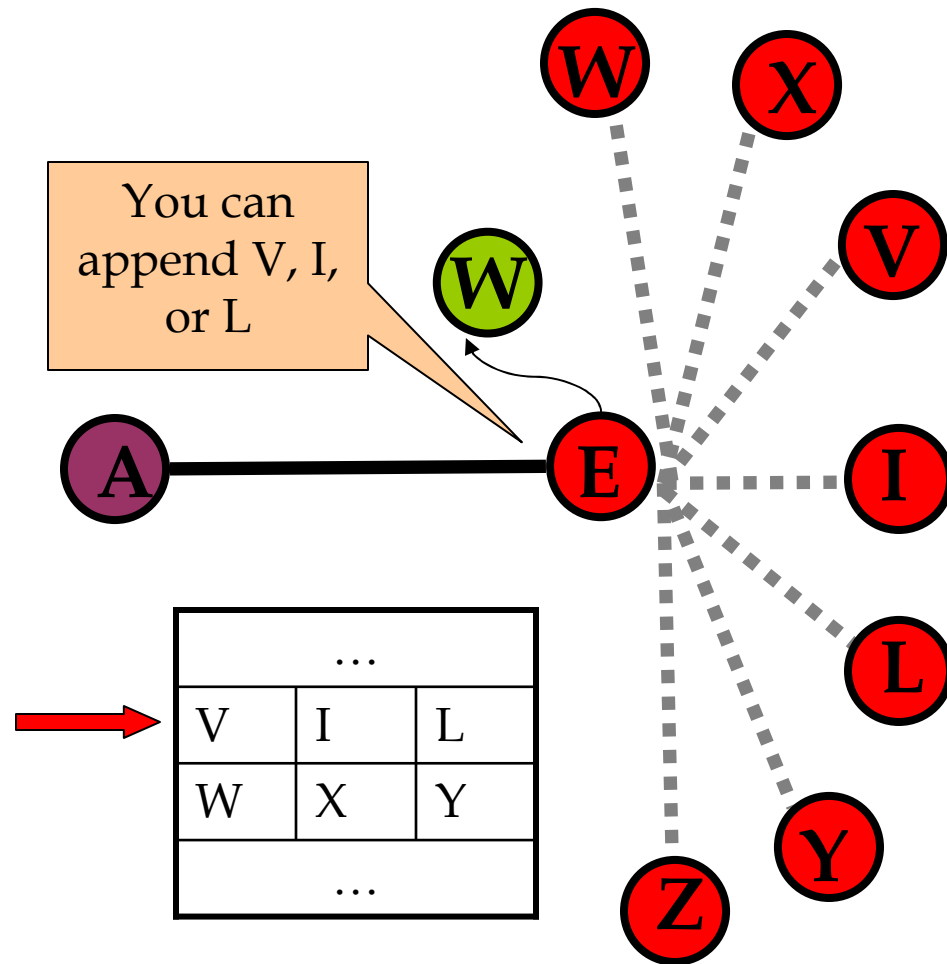
1. For every victim, v , construct a list of selections, S_v , comprised of only colluding nodes such that there is no overlap in node subnets between any selection entry in S_v .
2. Using a global pointer, p_g , maintain the position to a selection in S_v that is to be offered in the next attack attempt.
3. When v contacts any colluding node to be a first intermediate node and any subsequent node in a new anonymous tunnel, we offer the next selection to v .

Assumes that a node can determine when it is the first intermediate node in a tunnel.

Intelligent Selection Attack

E	V	I	L
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E
Q	G	P	T

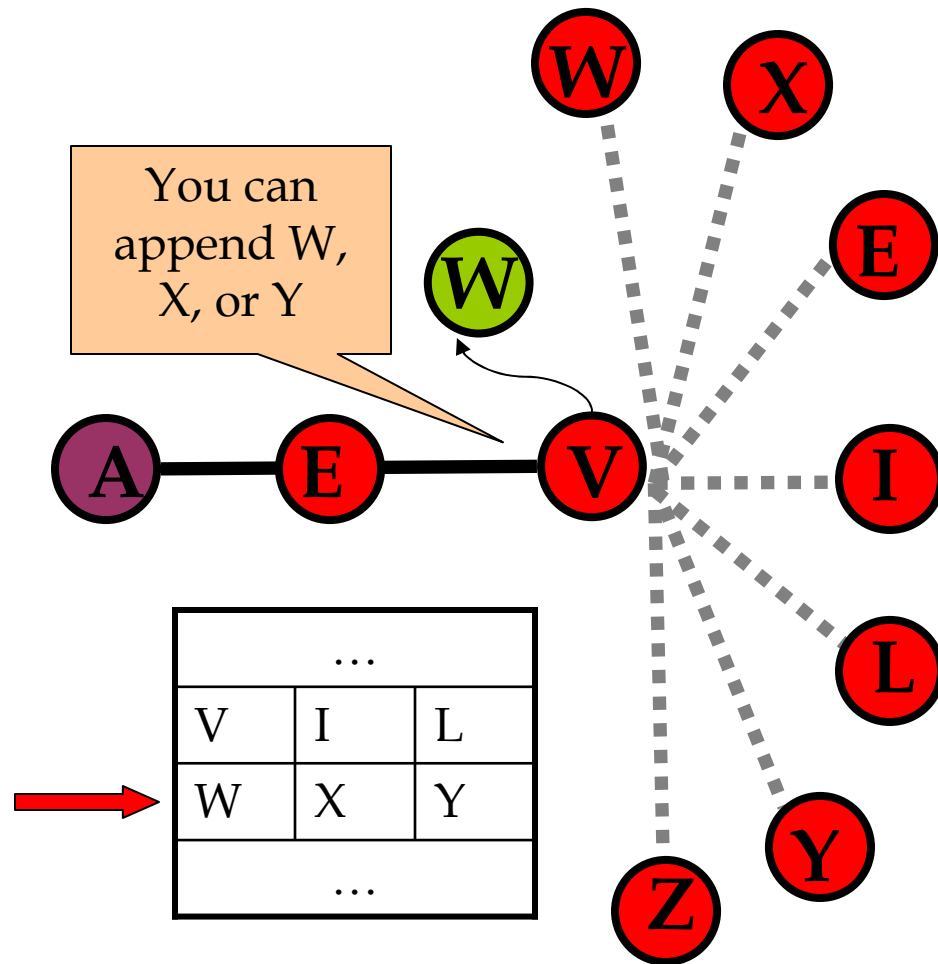
LE_S of Node A



Intelligent Selection Attack

V	W	X	Y
E	V	I	L
B	F	J	K
L	J	I	M
C	L	O	P
M	E	X	J
N	Q	R	E

LE_S of Node A



Intelligent Selection Attack (Revised)

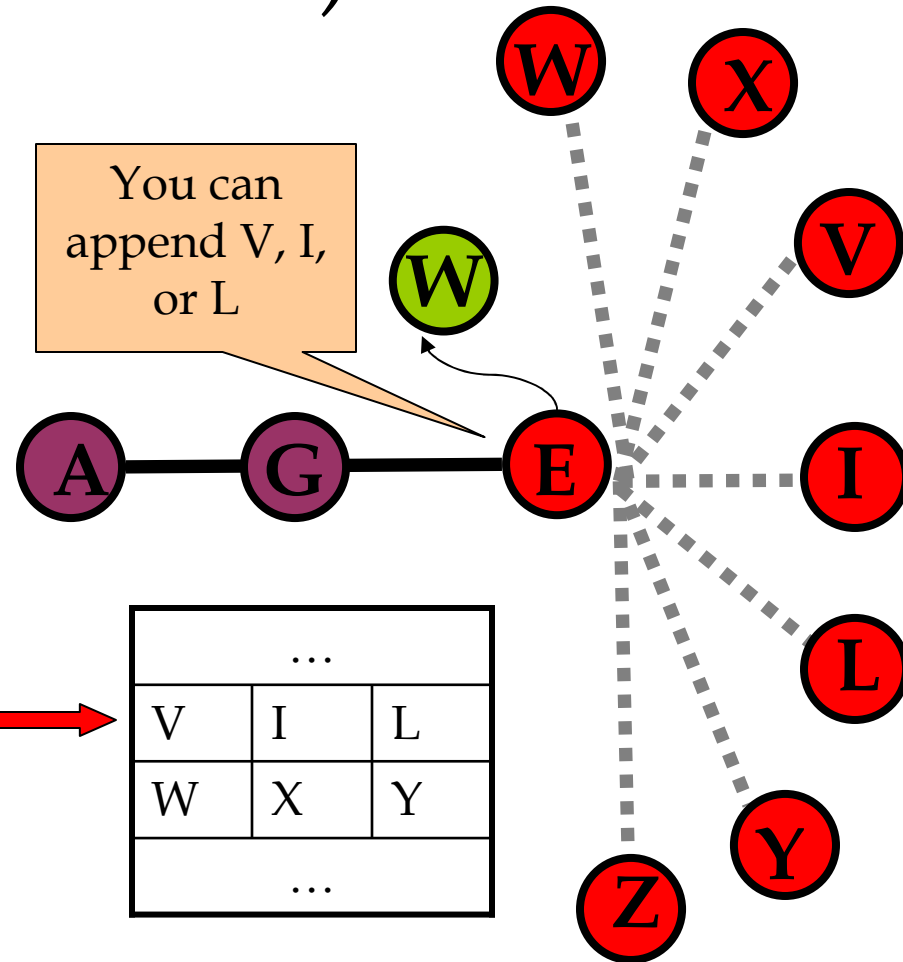
- Whenever v requests a selection from a colluding node, we begin the attack by assigning a local pointer, p_l , to the selection referenced by p_g and offer that selection to v . We cannot verify if v is the tunnel initiator or some other node, v_0 .
- For every successive selection request, we increment p_l in S_v and offer the new selection that p_l points to.
- After the tunnel is created, determine if the tunnel initiator was v .
- If the tunnel was initiated by v , we update p_g to hold the value referenced by p_l . Otherwise, p_g maintains its original value.

Intelligent Selection Attack (Revised)

B	E	A	D
F	O	L	M
J	K	M	N
I	T	A	L
G	E	R	M
C	U	T	W



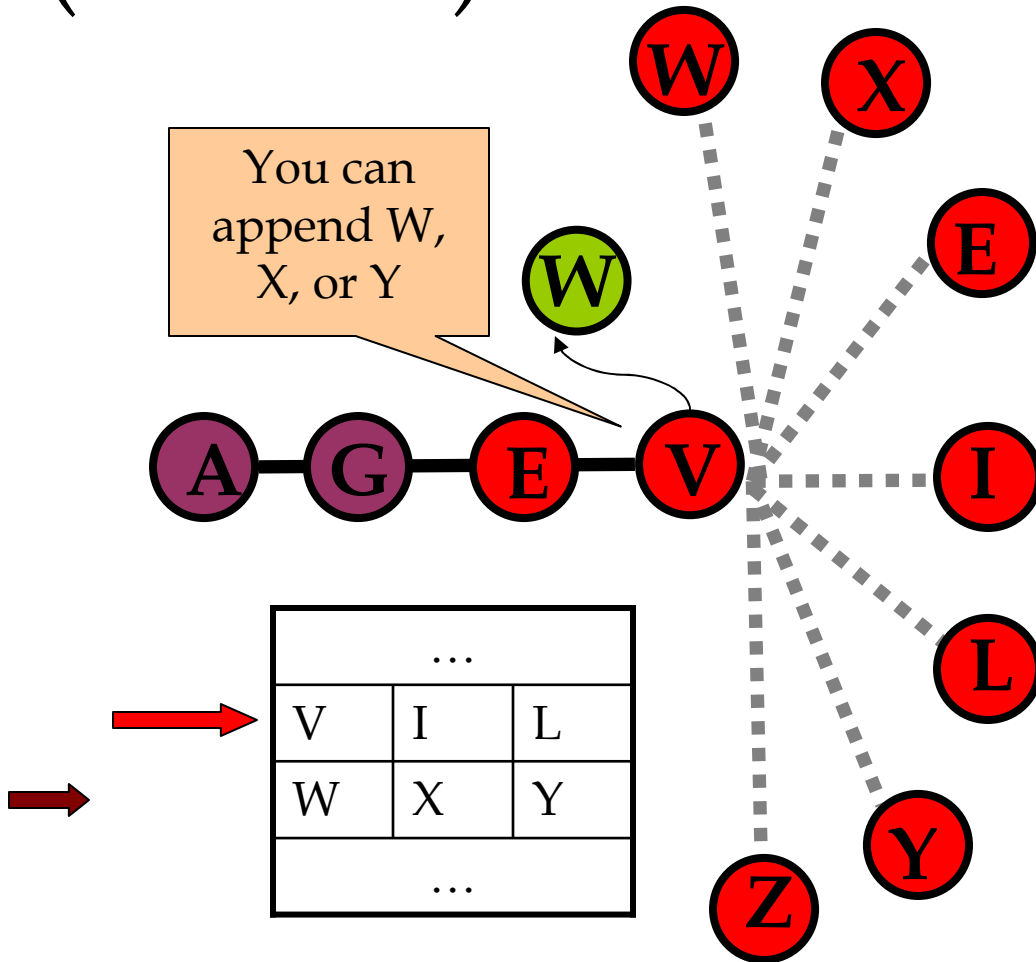
...		
V	I	L
W	X	Y
...		



LE_s of Node G

Intelligent Selection Attack (Revised)


B	E	A	D
F	O	L	M
J	K	M	N
I	T	A	L
G	E	R	M
C	U	T	W




LE_s of Node G

Intelligent Selection Attack (Revised)

B	E	A	D
F	O	L	M
J	K	M	N
I	T	A	L
G	E	R	M
C	U	T	W



DAMN!
Back to the
beginning.



LE_S of Node G

Simulation

- 10,000 MorphMix nodes
- 5,000 tunnels (~week usage)
- 1184 entries in the extended selection list
- 14 nodes in a selection
- 5 nodes to an anonymous tunnel

And a partridge in a pear tree...

Simulation Results

C	Honest Tunnels	Malicious Tunnels	Percentage Compromised
5%	3337.9	6.8	.2%
10%	2951.4	33.8	1.1%
15%	2283.2	470.1	17.1%
20%	1930.0	860.4	30.8%
30%	1171.5	1385.0	54.2%
40%	450.9	1847.5	80.4%

Countermeasures

- Increase the size of L_{ES} and selection size.
 - Doesn't prevent attack.
 - Negative impact on performance.
- Use variable length tunnels.
 - Limited to a small range of realistic values, so an attacker can estimate distribution of tunnel lengths.
 - Attacker *may* waste some selections, but attack is not eliminated.

General Limitation

- The CDM only considers a node's local knowledge when detecting collusive behavior.
- A strong adversary can model and manipulate this local knowledge to build malicious tunnels.
- An effective CDM requires a more global perspective of the network.

Lessons Learned

- Collusion detection is a difficult task in anonymous networking.
- There needs to be a middle ground between full global knowledge and limited scalability and good scalability but local knowledge only.

Questions