Analyzing Internet Routing Security Using Model Checking

Adi Sosnovich[1], Orna Grumberg[1], Gabi Nakibly[2]

[1] Technion, Haifa, Israel

[2] National EW Research and Simulation Center, Rafael, Haifa, Israel

Routing on the Internet

- The Internet is composed of **Autonomous Systems (ASes)**
- Each AS is administered by a single entity

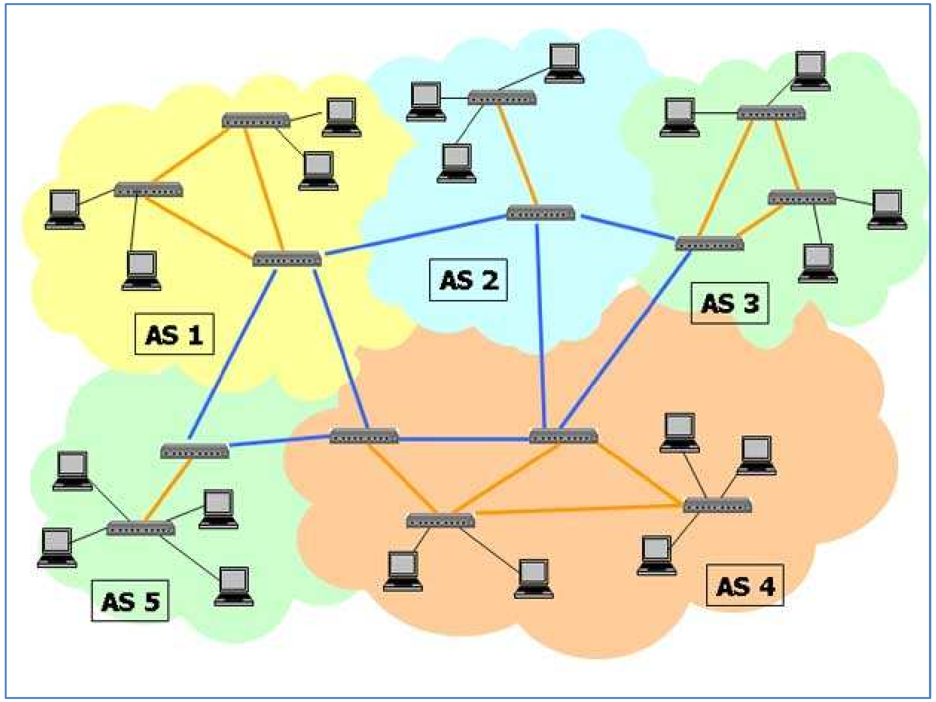

Inter-domain Routing

- Inter-domain routing determines through which ASes packets will traverse

- Routing on the AS level throughout the Internet is handled by a single routing protocol called the Border Gateway Protocol (BGP)

BGP Vulnerabilities

- The Internet is vulnerable to **traffic attraction attacks**

- A malicious AS can manipulate BGP to attract traffic to, or through, its AS

- Traffic attraction enables the AS to:
  - increase revenue from customers
  - drop, tamper or snoop on the packets

Example – Traffic Diversion

Source: http://research.dyn.com/2013/11/mitm-internet-hijacking/

Goals

- Reveal non-trivial scenarios of traffic attraction
- Provide insights to where and how BGP traffic attraction attacks are possible on the **Internet**

- Using techniques and tools from formal methods:
  - Model checking
    - To **automatically** find attraction scenarios or prove their absence
  - Reductions and abstractions
    - To handle the **full** Internet topology (~50,000 ASes)

The BGP Routing Protocol

- A **routing update** consists of a target network $\boldsymbol{n}$ and a path $\boldsymbol{\pi}$ of ASes

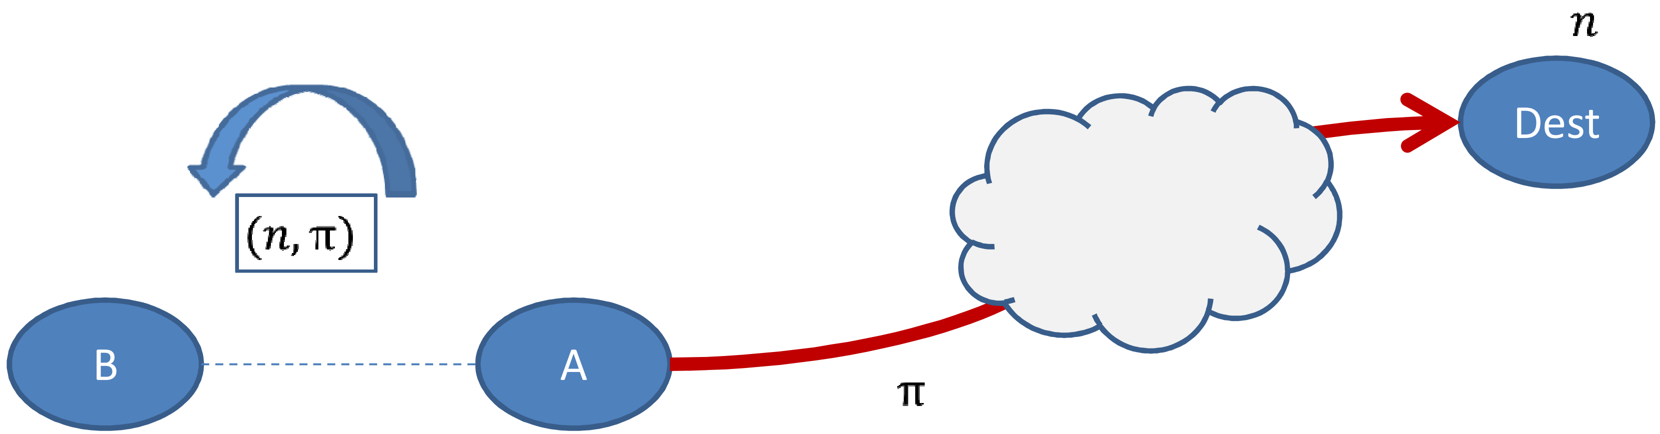

A announces to B that it is willing to carry packets destined to n from B, and the packets will traverse over the path π.

The BGP Routing Protocol

- Every AS **stores** the routes learned from its neighbors

- Each AS has a **local policy**:
  - If an AS has several routes to the same target network, it must choose its **most preferable** one [preference policy]
  - An AS can **propagate** its chosen path to a certain destination by prepending itself to that route and sending it to some of its **neighbors** [export policy]

- Theses policies are affected by business relationships between ASes

Business Relationships Between ASes

- **Customer-provider** : The customer **pays** its provider for connectivity

- **Peer-peer**: two ASes agree to transit each others traffic at **no cost**

Preference and Export Policies

- Normal Preference Policy:
  - Prefer routes announced by customers over routes announced by peers over routes announced by providers
  - Among the most preferable routes choose the shortest ones
  - If there is more then one such path, choose the one announced by the AS with lowest ASN
  - A path in which the AS itself already appears is rejected

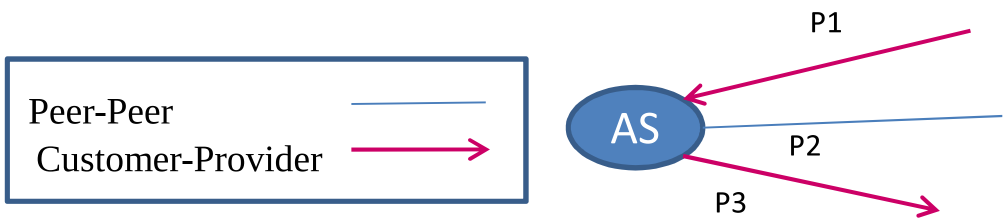

P1>P2>P3

Preference and Export Policies

- Normal Export Policy:
  - B will announce to A a route P via C if and only if at least one of A and C are customers of B

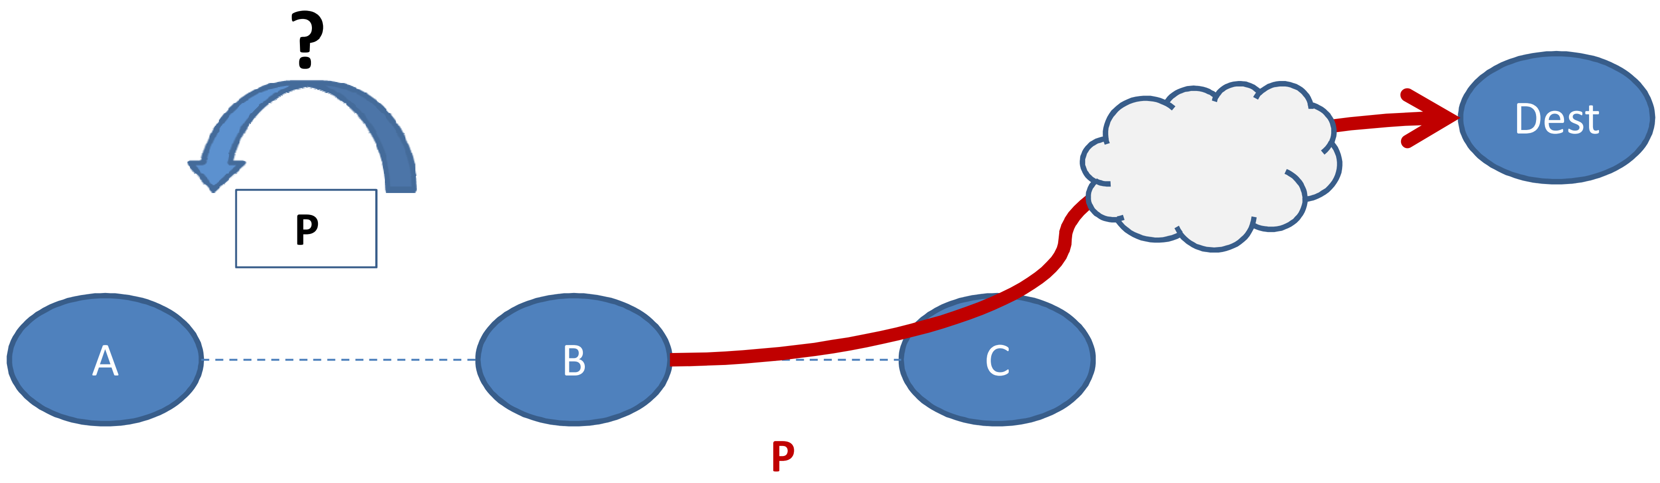

The BGP Routing Protocol - Example

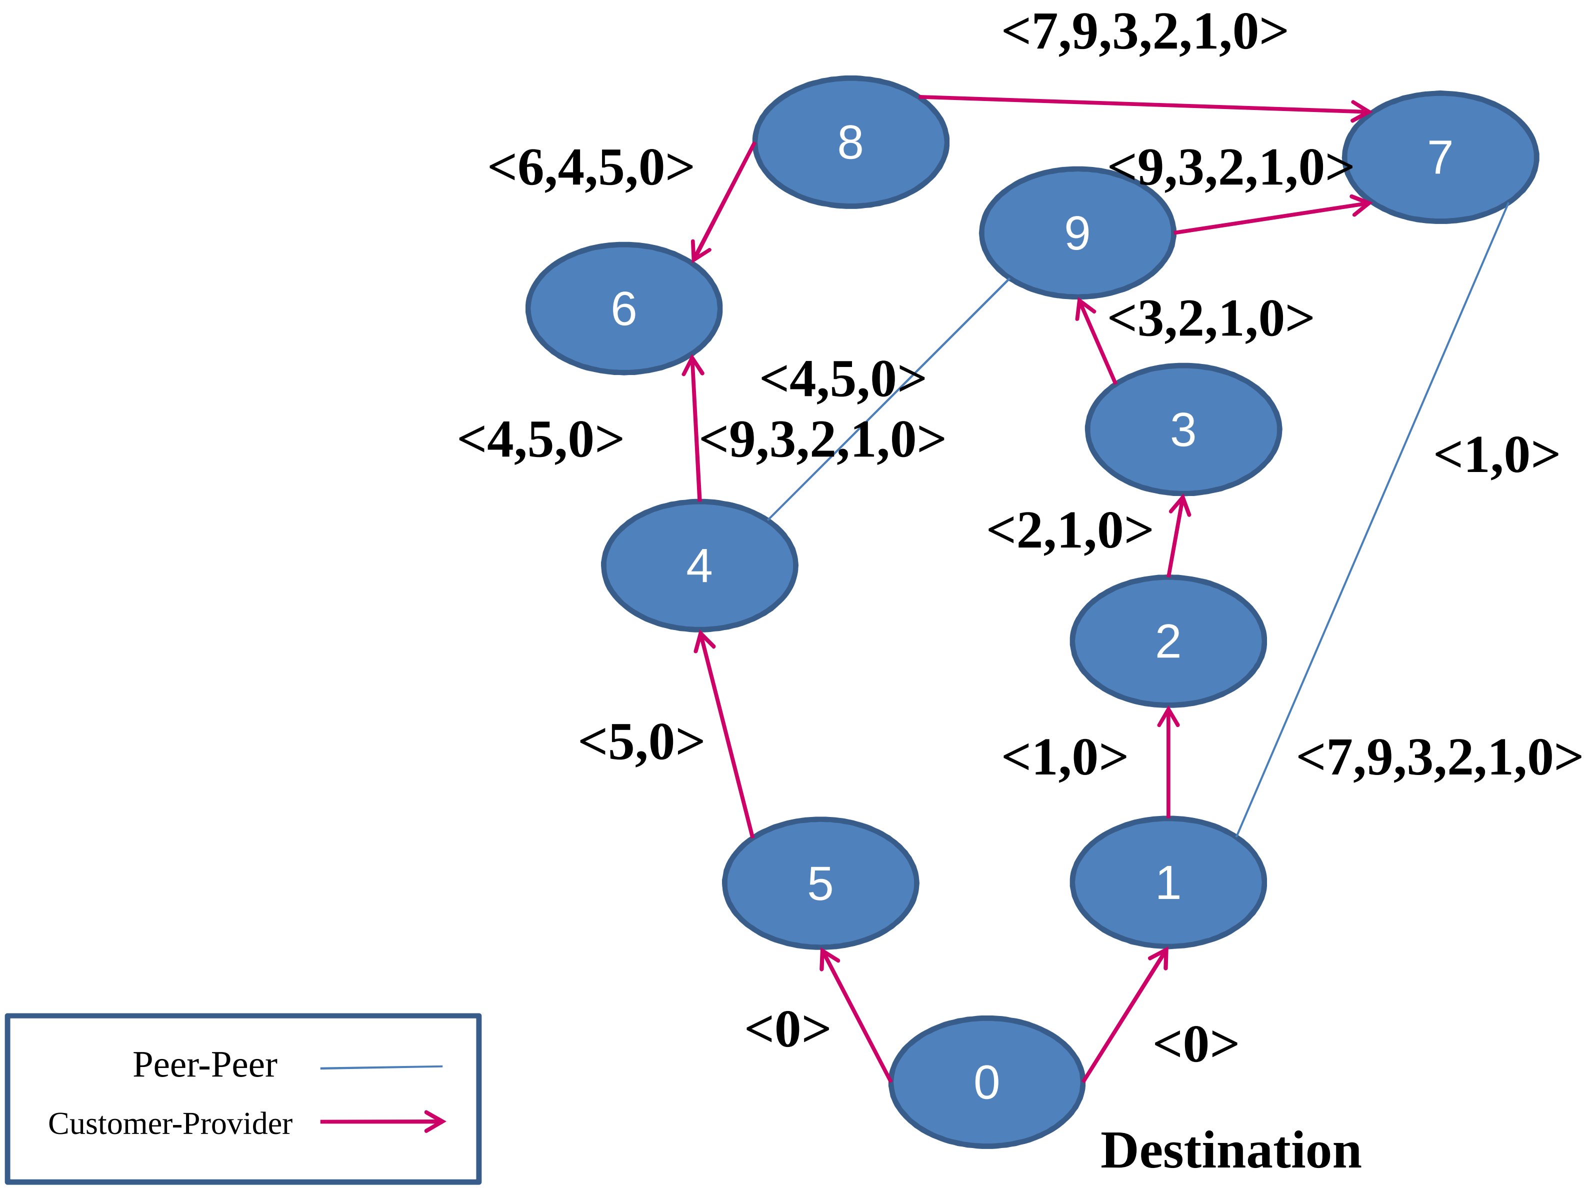

BGP Modeling

- **Network topology**
  - A graph of AS nodes with edges of type peer-peer or customer-provider
- **Dest** is a **single** predefined **destination** AS in which the target network resides
  - all ASes try to build routing paths to it
- **Attacker** is a predefined AS node representing a manipulator that can send false routing advertisements
  - Its goal is to achieve traffic attraction
  - It can send arbitrary paths and use arbitrary export policy

Types of traffic Attacks

- **Interception attacks**:
  - The traffic is diverted to the attacker's AS and then forwarded to its real destination
  - Allows the attacker to become a man-in-the-middle
- **Attraction attacks**:
  - The traffic is not forwarded to its real destination
  - Allows the attacker to impersonate the real destination or block access to it

Normal outcome

- **Normal outcome** is the final routing choices of all nodes when the attacker acts like a regular AS

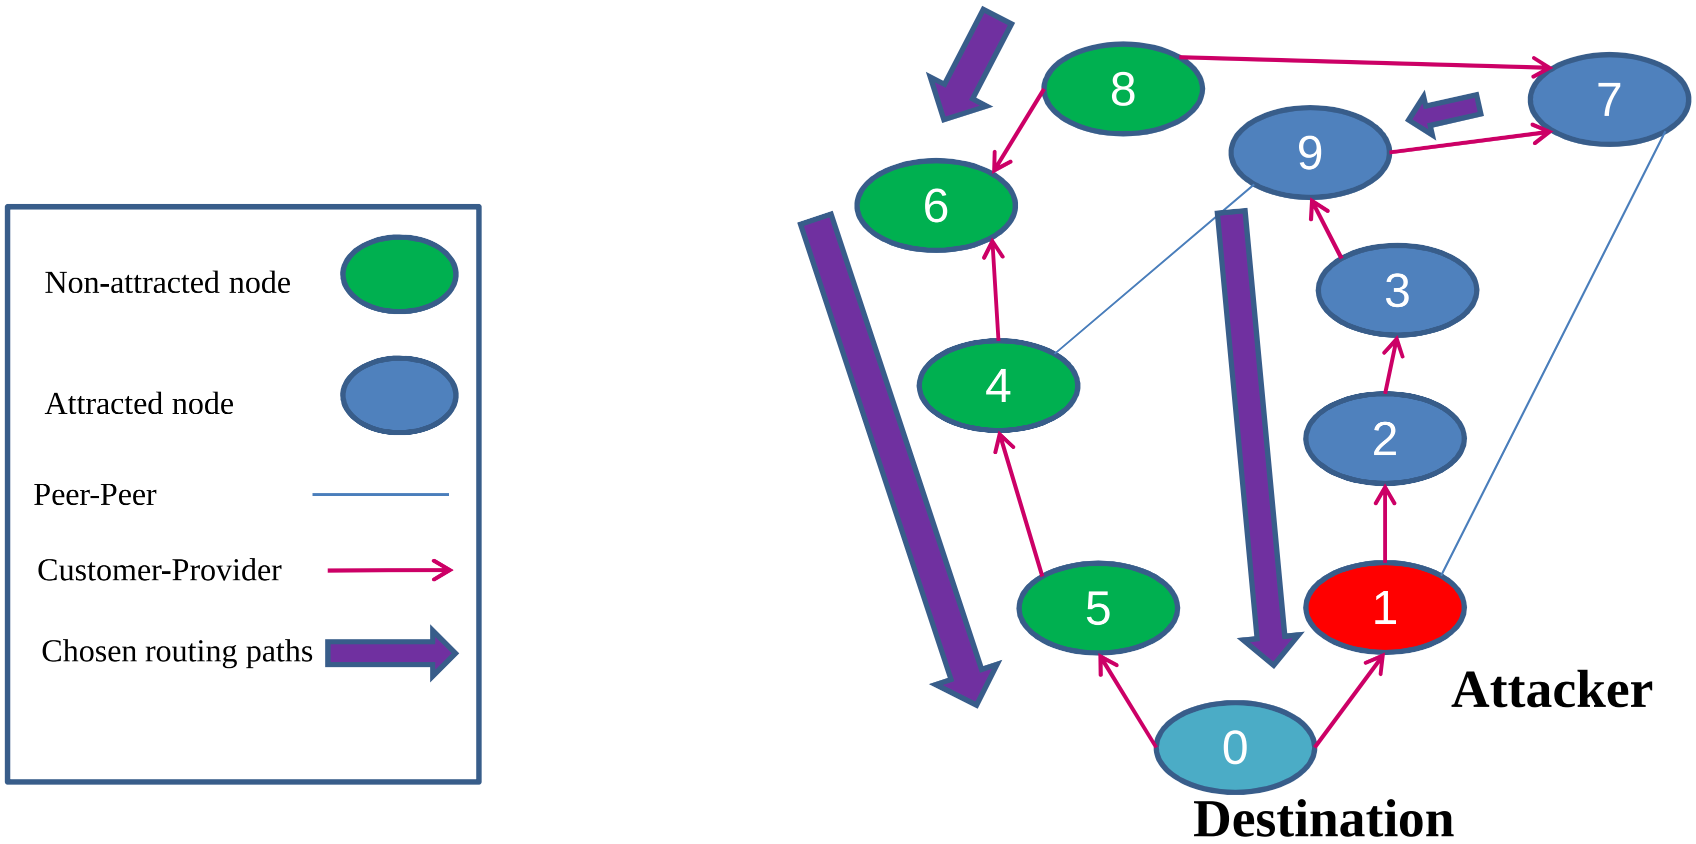

Trivial Attack Strategy

- In the trivial attack strategy the attacker **sends** a **false** advertisement to **all its neighbors** that the **target** network is located within its **own AS**

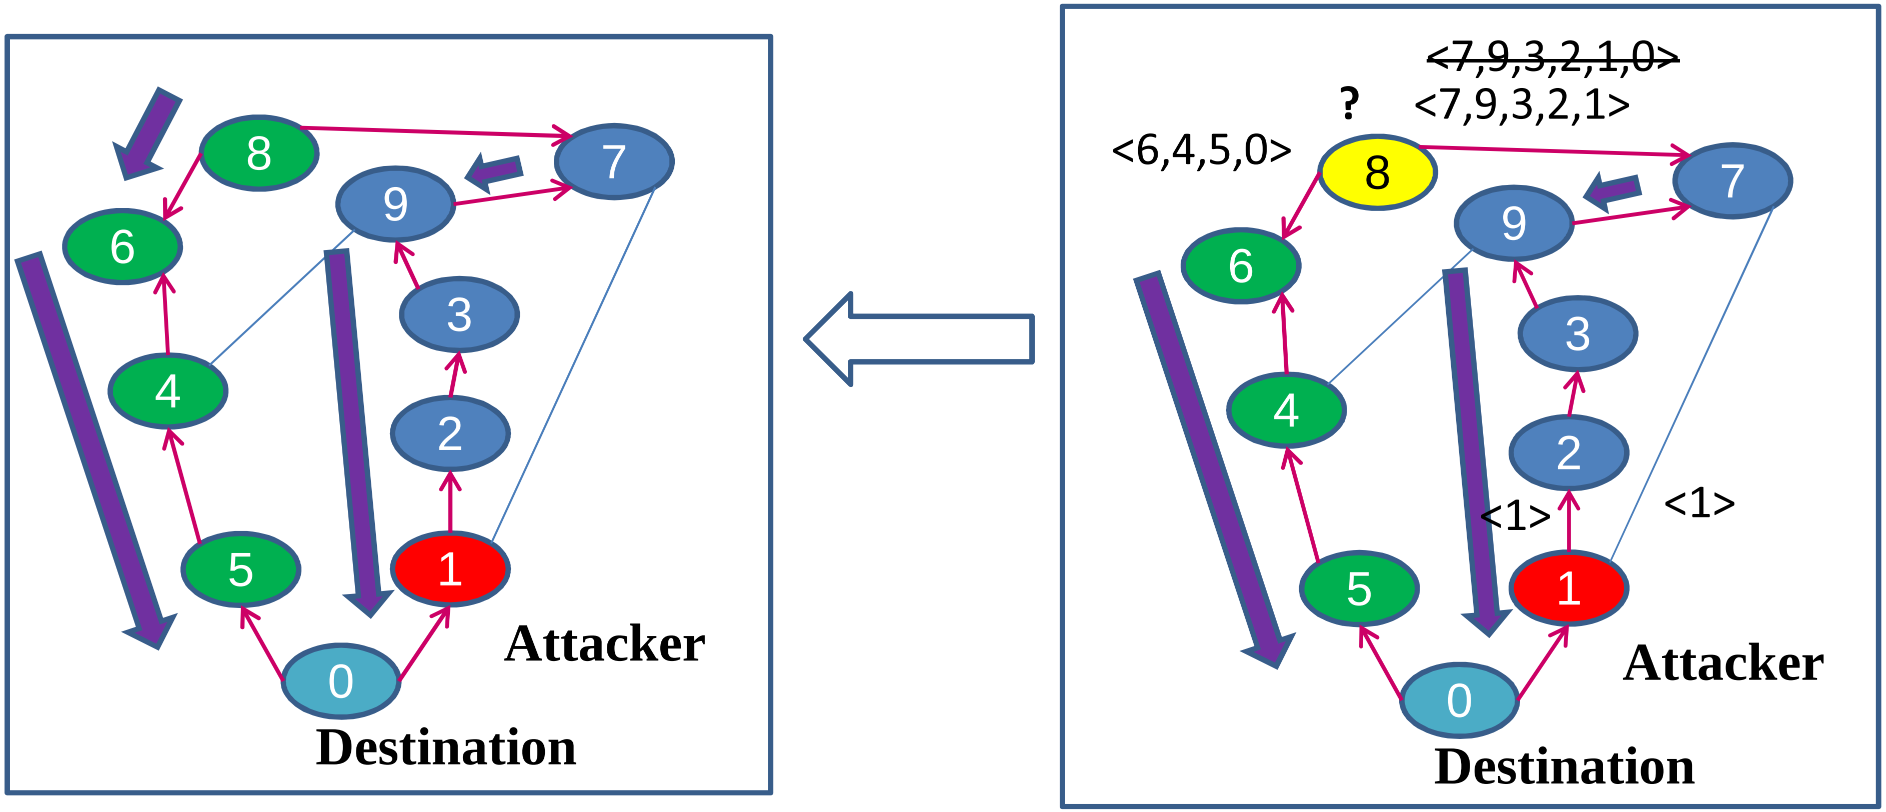

Specification

- We search for non-trivial attack strategies
- We search for attacks that manage to gain new attraction/interception
- We **specify** when an attack is successful based on a **comparison** to other BGP runs: a normal run and a run with a trivial attacker
- **If the attacker can attract (or intercept) traffic from some victim, while it fails to do so in the normal run and in the trivial attack, the attraction (or interception) specification is satisfied**

Reductions of a BGP Network

- To find traffic attraction scenarios or prove their absence we use **model checking**
- Applying model checking on the full Internet topology (~50,000 ASes) is **infeasible**
- We develop reductions to obtain a manageable sized **fragment** of the large network

Network Reduction – First Attempt

- Pick an **arbitrary** sub-network from the Internet
- Problem:
  - If some attraction scenario is **found**, it is **not guaranteed to be preserved** in the context of the full Internet topology
  - ASes outside of the sub-network may interfere and affect the routing choices of ASes within that sub-network

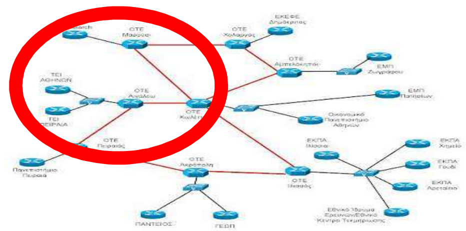

- Solution:
  - Find an isolated sub-network that is not affected by ASes outside, by using **valid paths**

Valid Paths

- A path $(n_1, \dots, n_k)$ in the BGP network is **valid** if :
  - $n_1 \in \{Attacker, Dest\}$
  - For each n_i with $1 < i < k$:
    - $n_i \notin \{Attacker, Dest\}$
    - At least one of n_{i-1}, n_{i+1} is a customer of n_i
  - No node is repeated on the path

Valid Paths Examples

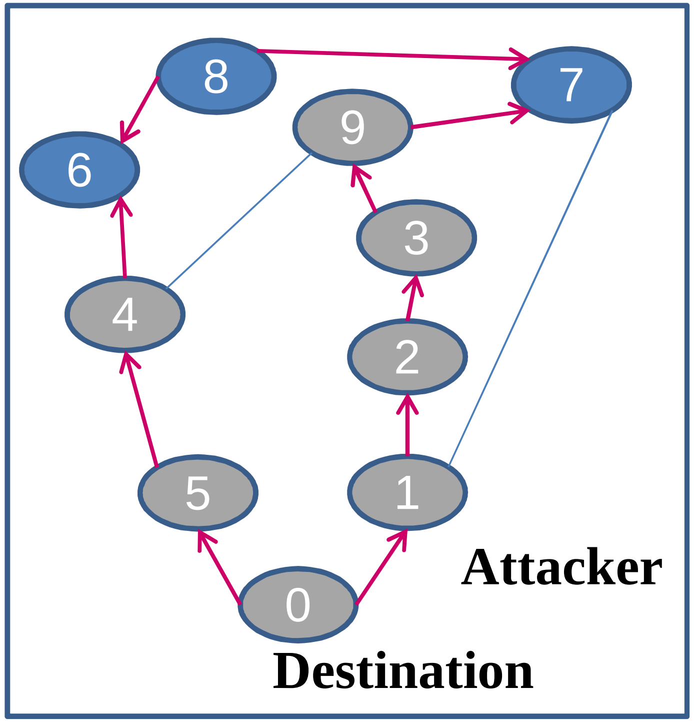

(0,5,4,9,3,2,1)

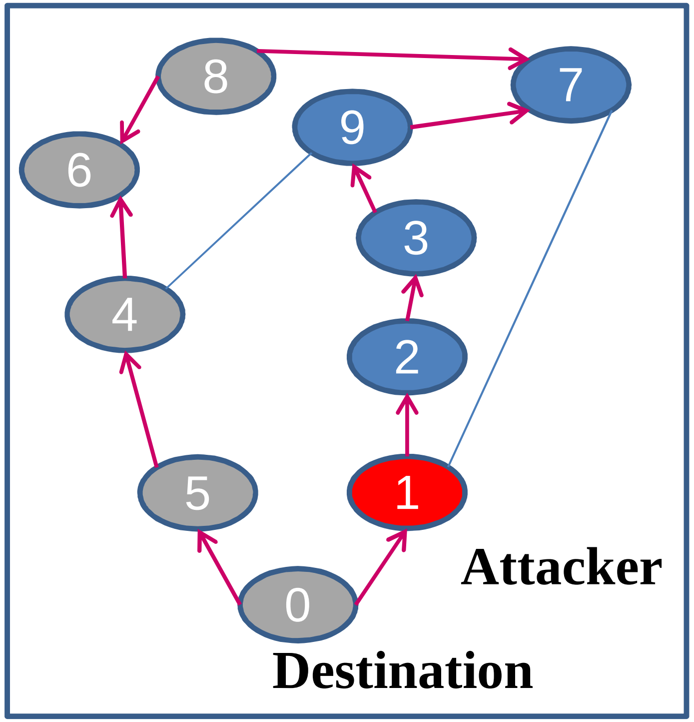

(0,5,4,6,8)

Export actions of regular nodes is performed only along valid paths

Self-contained Fragments

- Let **S** be a sub-network of a BGP network
- **S** is a **self-contained fragment** of a BGP network if for every node $n \notin S$, there is no BGP run in which an export action from n to some $n' \in S$ is performed

- **Nodes outside of S cannot change routing decisions of nodes in S**

Self-contained Fragments

- Lemma:
  - Let N be a (large) BGP network and let S be a self-contained fragment of N
  - Then, any traffic attack found on S can occur on N as well
  - Moreover, if we obtain a proof that an attacker cannot attract traffic from some victim within S, then the proof applies for N as well

Extracting Self-contained Fragments

- Initially, {Dest, Attacker} and their neighbors are in **S**
- **A node $c \notin S$ is added to S if:**
  - **c is a neighbor of some $n \in S$**
  - **c is on a valid path from some originator (Dest/Attacker) to n**

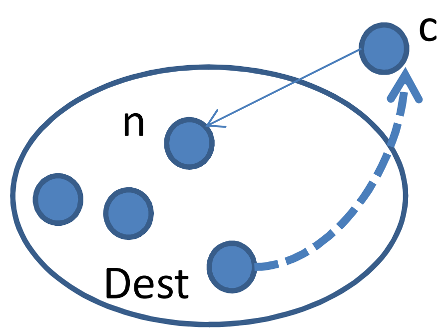

Extracting Self-contained Fragments

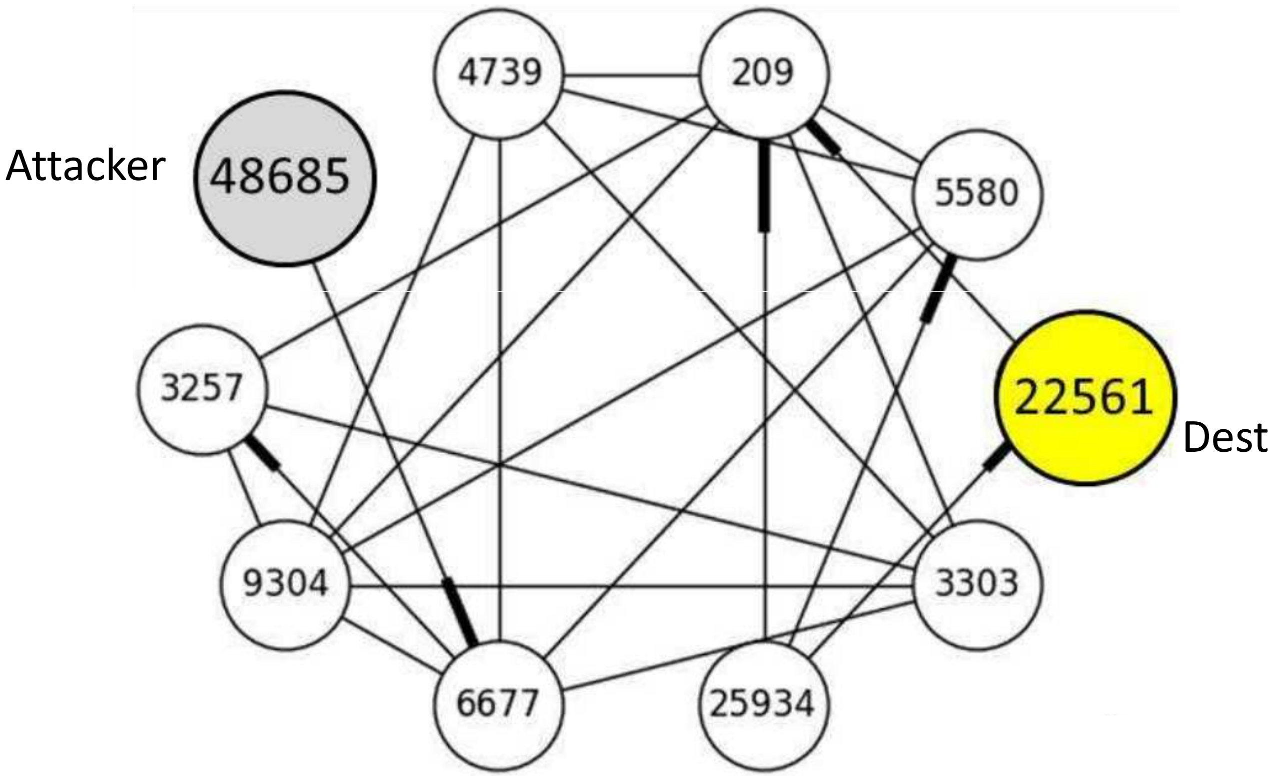

Extracting Self-contained Fragments

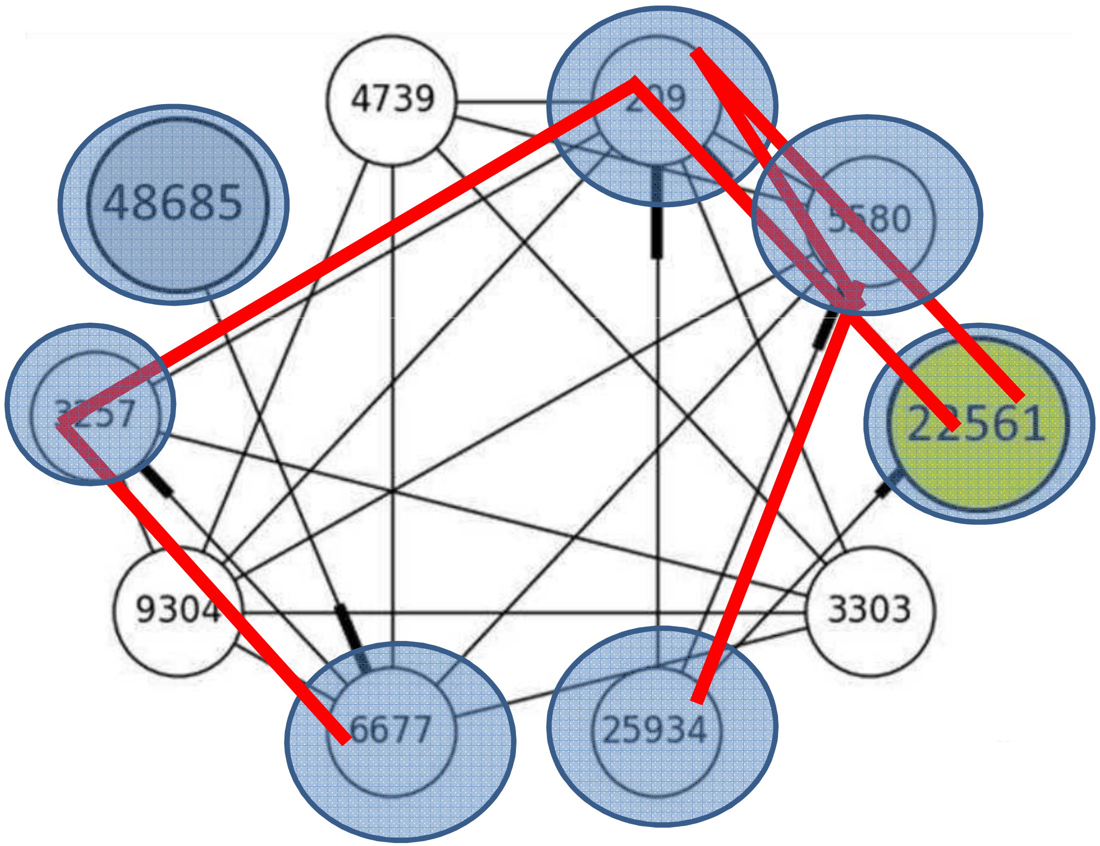

Extracting Self-contained Fragments

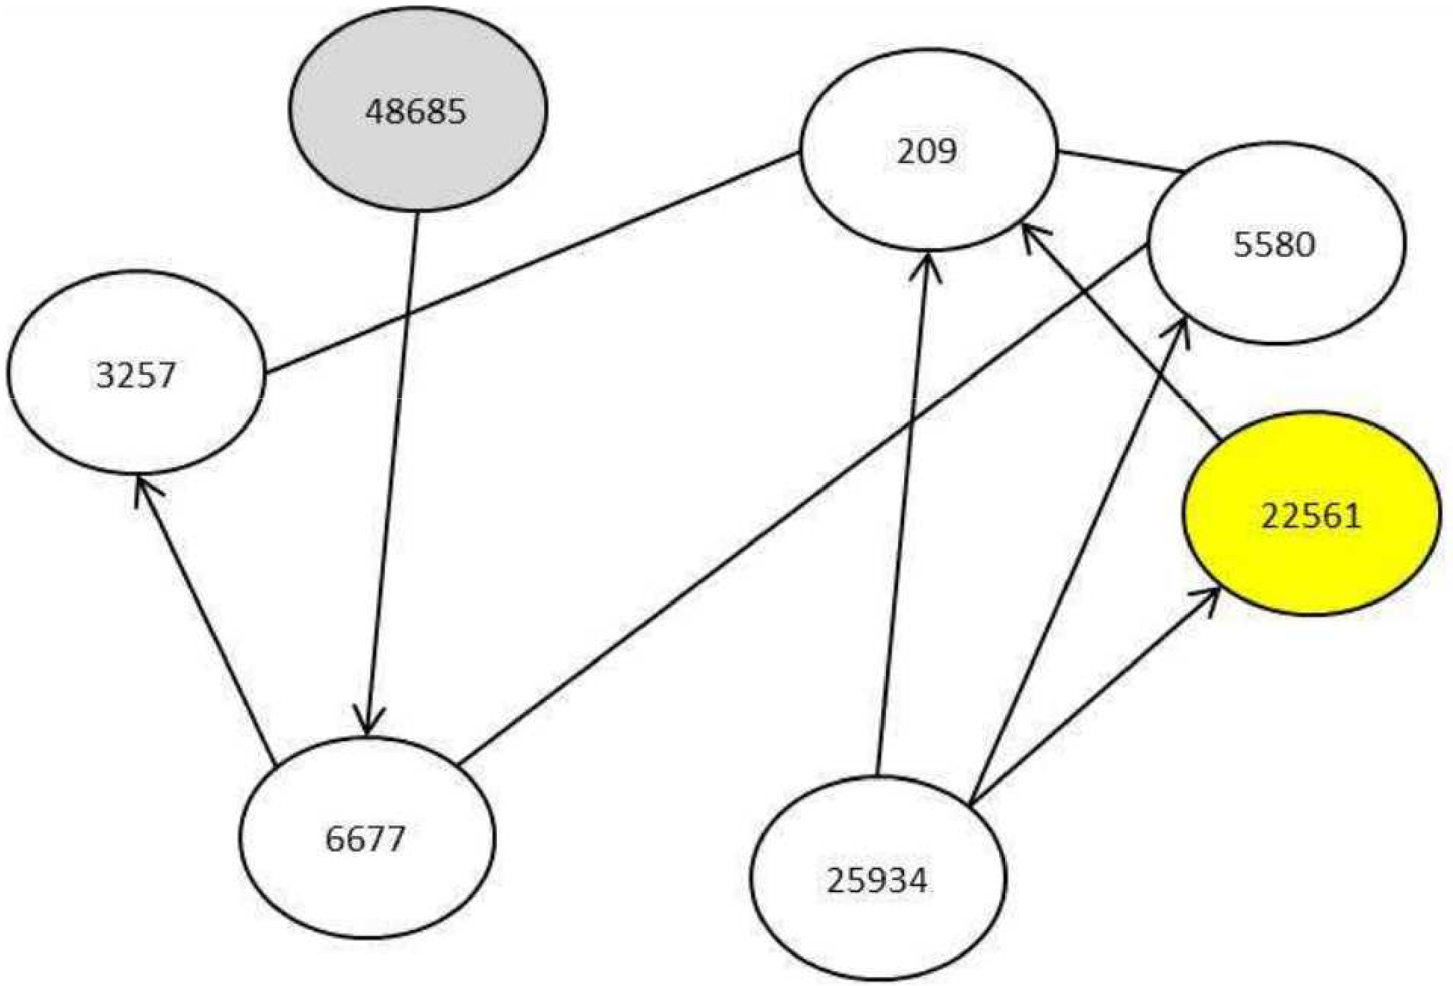

Definite Routing Choices

- Identifying nodes that **never route via the attacker**
- A node has a **definite routing choice** if its **chosen path is via the destination and not via the attacker for every possible run**, regardless of the attacker's actions

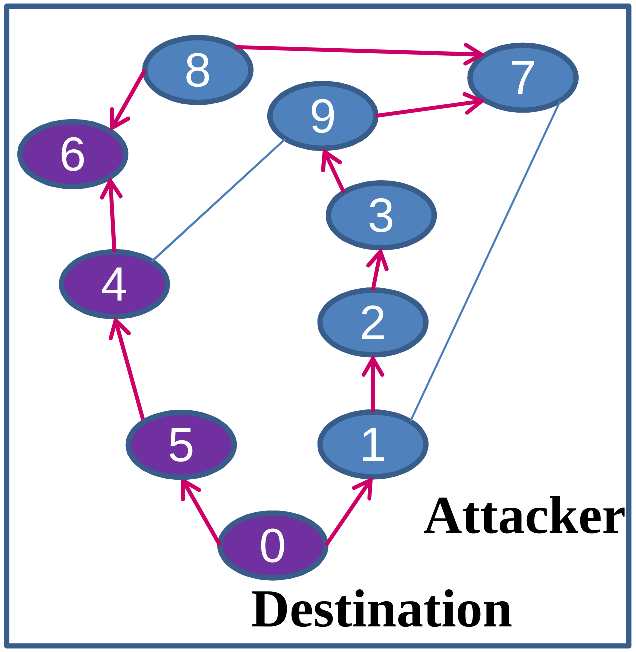

Definite Routing Choices Reduction

- A node with a definite routing choice can be **eliminated** from the network
- Its export actions to non-eliminated neighbors are known
- After elimination, the model’s initial configuration is updated : the results of the exports actions are already in the queues of the appropriate neighbors

The BGP-SA Method

- We use reductions and model checking to apply a formal **BGP security analysis** of traffic attraction attacks on the Internet

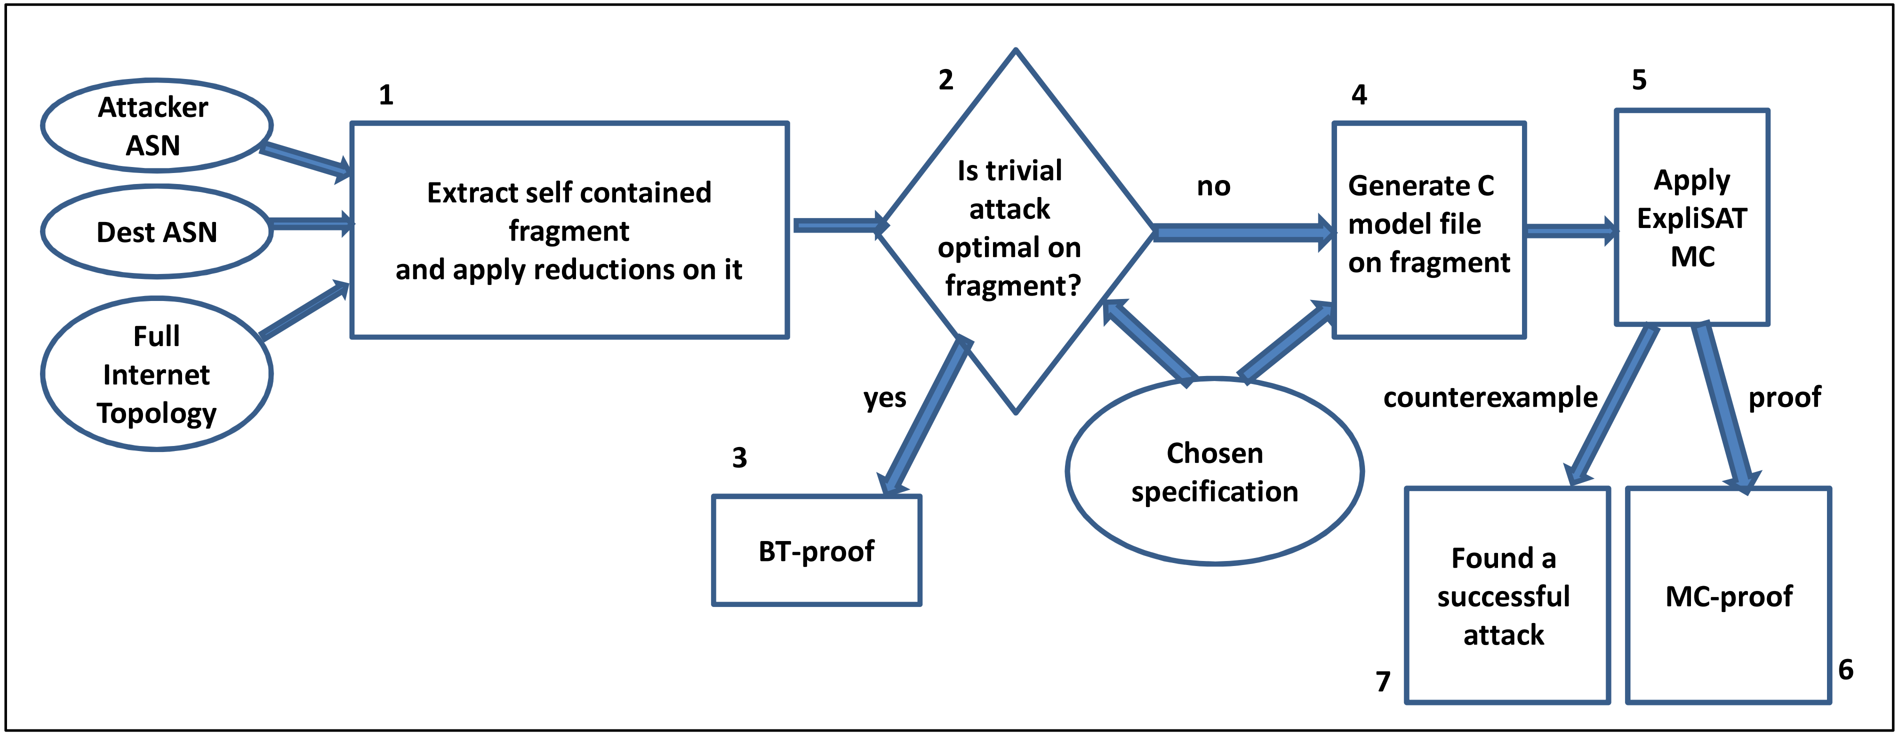

Trivial Attack Simulation

- We run on the reduced fragment a simulation of the trivial attack
- If **all nodes are attracted** then the trivial attack is **optimal within the fragment**, and there cannot be found a non-trivial strategy to gain new attraction
- This is considered as Best Trivial attraction proof (**BT-proof**)

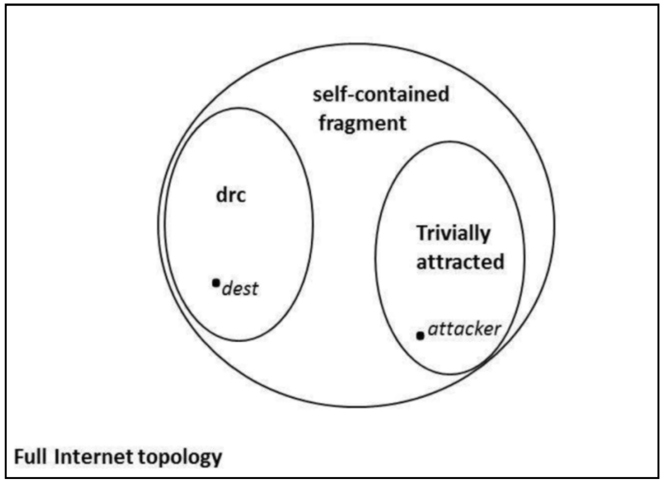

Safe Nodes

- We identify safe nodes, that cannot be attracted by the attacker:
  - Nodes that have a definite routing choice
  - Nodes for which the model checker provides a proof that there is no attacker's strategy that can attract them

Related Work

- Goldberg, Sharon, et al. "How secure are secure interdomain routing protocols." *ACM SIGCOMM Computer Communication Review* 41.4 (2011): 87-98. [Goldberg 2011]

  - Demonstrates non-trivial and non-intuitive attack strategies
  - Gives anecdotal evidence, obtained manually, for each attack strategy in specific parts of the Internet

Example of a non-trivial interception scenario

- [Goldberg 2011] showed a non-trivial interception scenario on a variation of the network below
- In that scenario, **the attacker does not export a path to AS2**

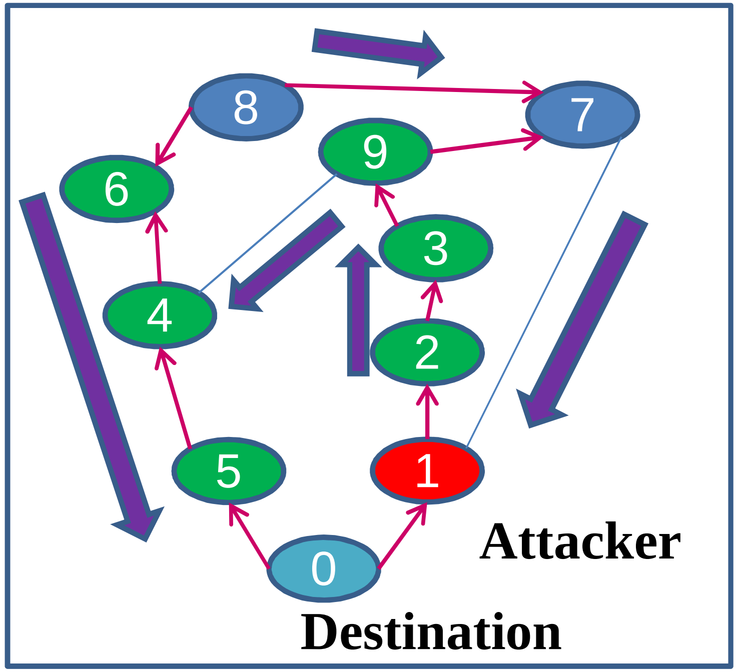

New attacker's strategy – new attraction

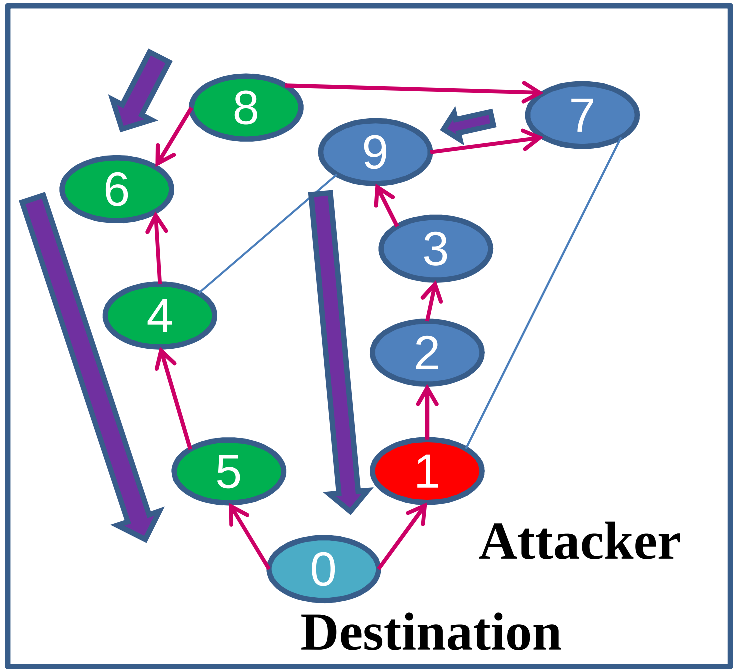

Normal outcome and trivial attack

Applying model checking to find non-trivial interception scenarios

- The model checker found a scenario with greater attraction
- In the found scenario, the attacker exports a path to AS2 that creates a loop at AS9 : <1,9>, causing only AS9 to reject the path <3,2,1,9>

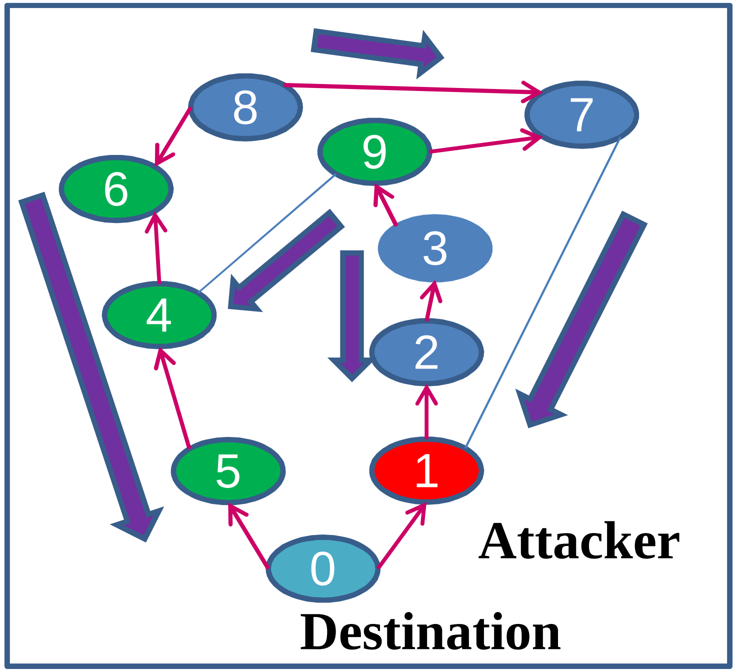

The newer strategy found by MC

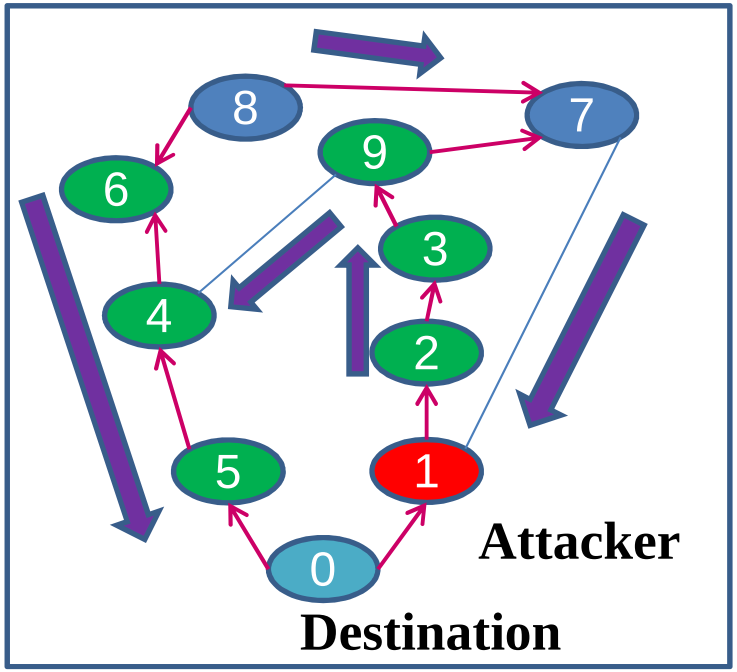

New attacker’s strategy – new attraction

Results on Internet Fragments

	Fragment size (#nodes)	Reduced size (#nodes)	Trivial attraction (#nodes)	Specification	Result	Time (min)	Dest ASN	Attacker ASN
1	16	11	9	attraction	BT proof	-	31132	16987
2	17	6	4	attraction	BT proof	-	9314	7772
3	22	10	8	attraction	BT proof	-	11669	36291
4	29	9	5	attraction	MC proof	1.5	29117	15137
5	15	13	10	attraction	MC proof	1	12431	18491
6	36	18	7	attraction	MC proof	17	19969	13537
7	69	27	17	attraction	MC proof	340	8296	20091
8	15	13	invalid	interception	counterexample	0.1	12431	18491
9	28	10	invalid	interception	counterexample	0.5	19361	32977
10	80	48	invalid	interception	counterexample	13	9218	43571
11	81	31	invalid	interception	counterexample	9	37177	40473
12	114	30	invalid	interception	counterexample	18	36040	29386
13	71	68	65	interception	N/A	>12h	30894	1290

Conclusion

- The Internet is vulnerable to traffic attraction attacks

- We developed automatic analysis that can reveal possible attraction scenarios on the Internet and prove that certain scenarios are not possible

- Our method is based on useful reductions that enable the automatic analysis