



BGPpy: The BGP Python Security Simulator

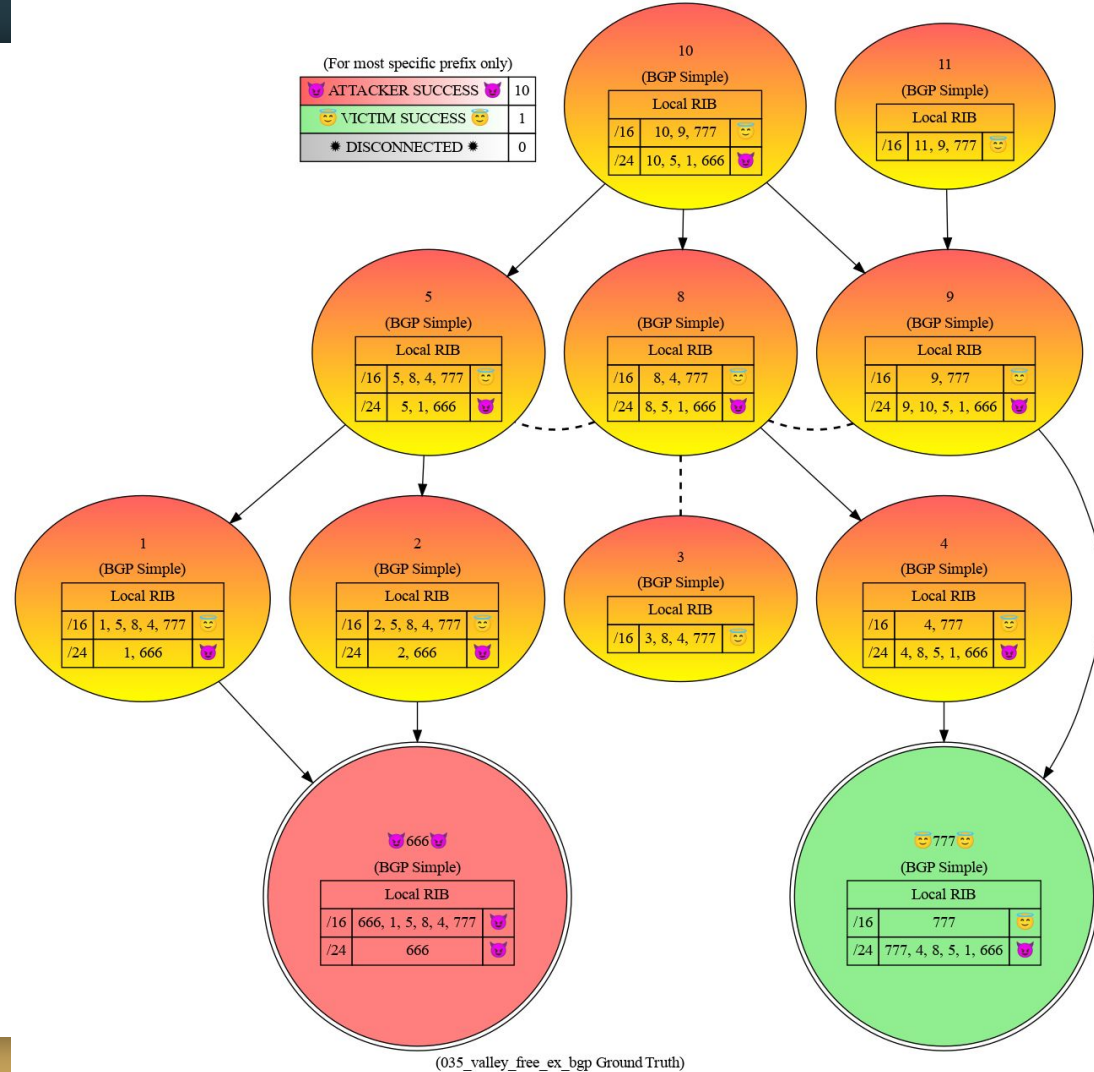
By: Justin Furuness, Cameron Morris, Reynaldo Morillo, Dr. Amir Herzberg, Dr. Bing Wang

(Affiliated with the University of Connecticut)



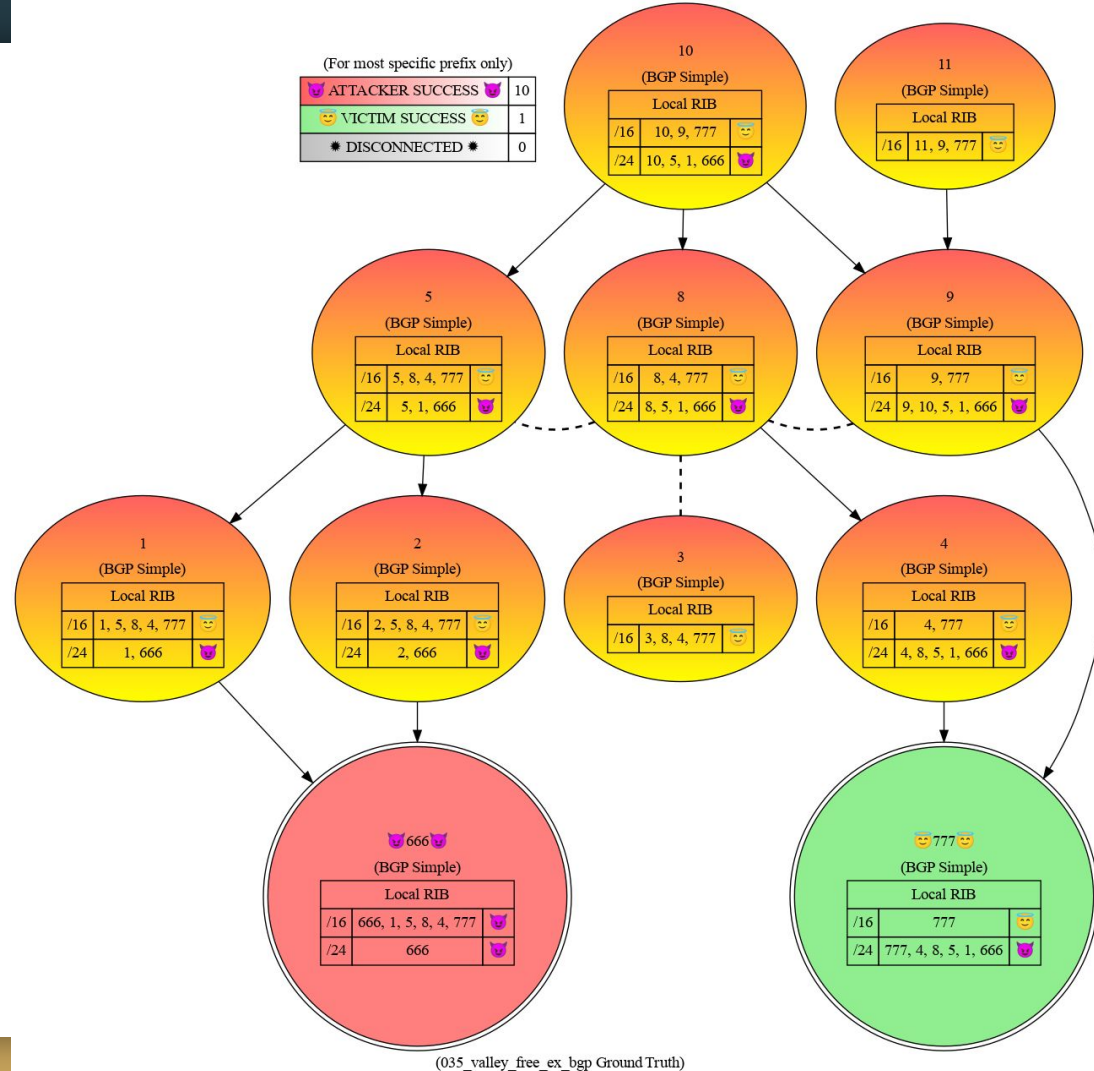
Background

- The internet is from a high level, a directed acyclic graph of over 70k nodes, known as Autonomous Systems, or ASes for short
 - These ASes route traffic for ISPs, CDNs, Organizations, etc
 - They can be identified by their AS Number, or ASN for short



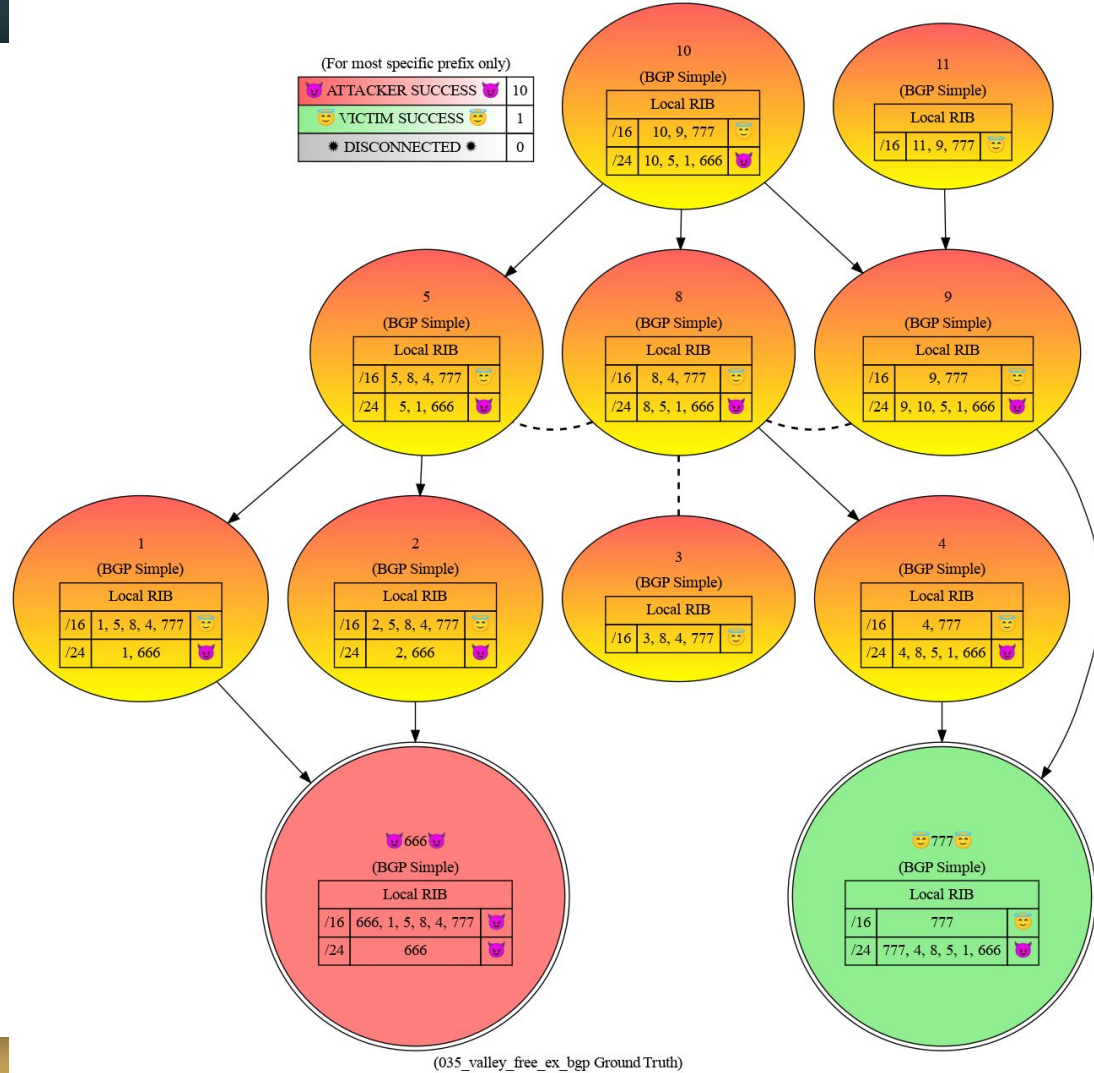
AS Relationships

- For our simulations, we focus on two of the most prevalent type of connections between ASes:
 - Provider -> Customer connections
 - Customers pay provider
 - Peer <-> Peer connections
 - Traffic flows freely
 - Sibling relationships not included




Announcements

- ASes forward traffic from one another. While announcements have many attributes, from a high level we can think of them as:
 - Prefix (containing a block of IP addresses), such as 1.2.3/24
 - Origin (the ASN that created the announcement)
 - AS Path (The path that the announcement took to reach it's destination)




Valley Free Routing

- Valley Free Routing is the default in our simulations, similar to prior works. This selects the best announcement in order to maximize profit.
 - customers > peers > providers
 - Shortest AS Path
 - Tiebreakers
 - Default to lowest ASN



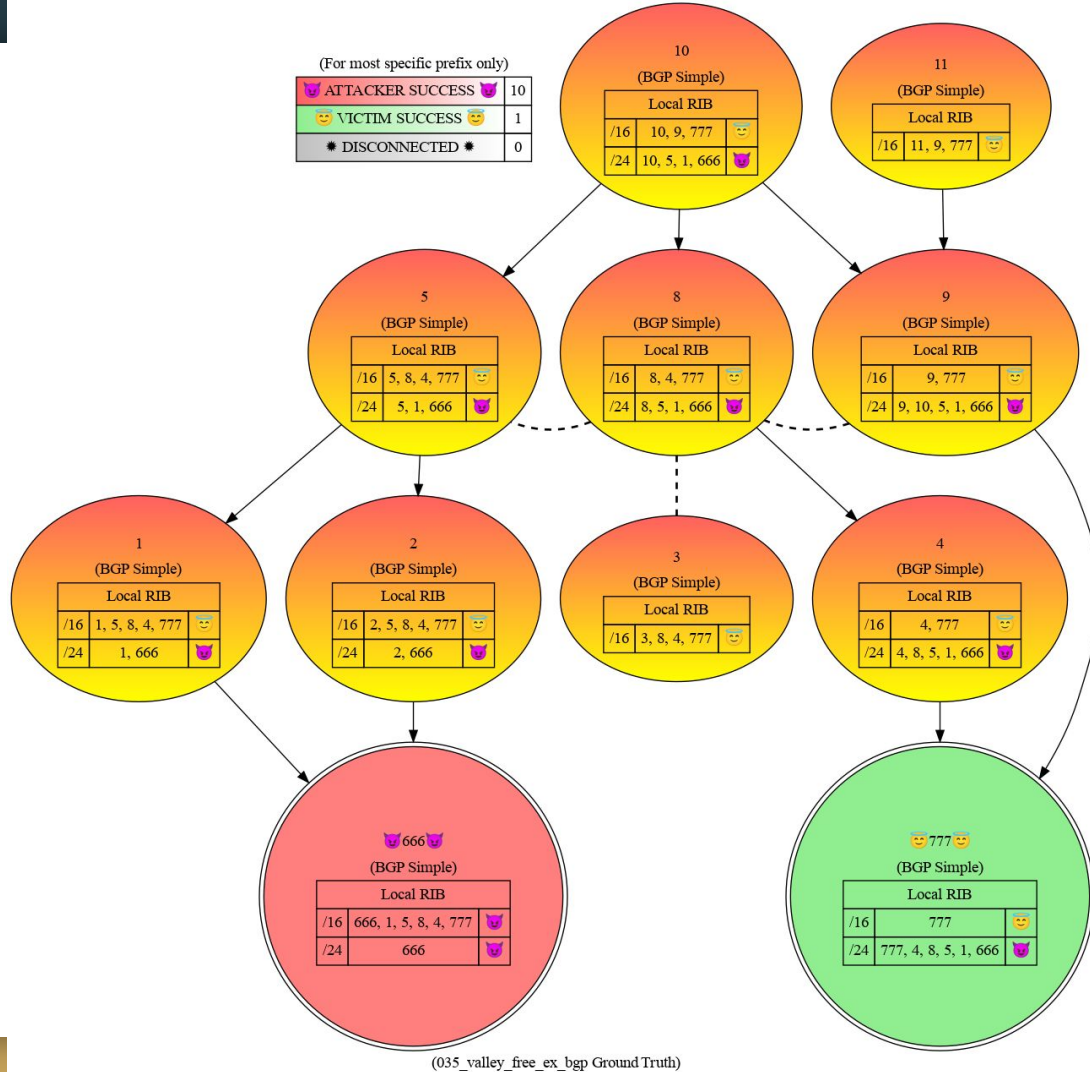
Export Policy

- Export Policy:
 - Announcements received from customers are sent to all
 - Announcements received from peers and providers are sent to customers



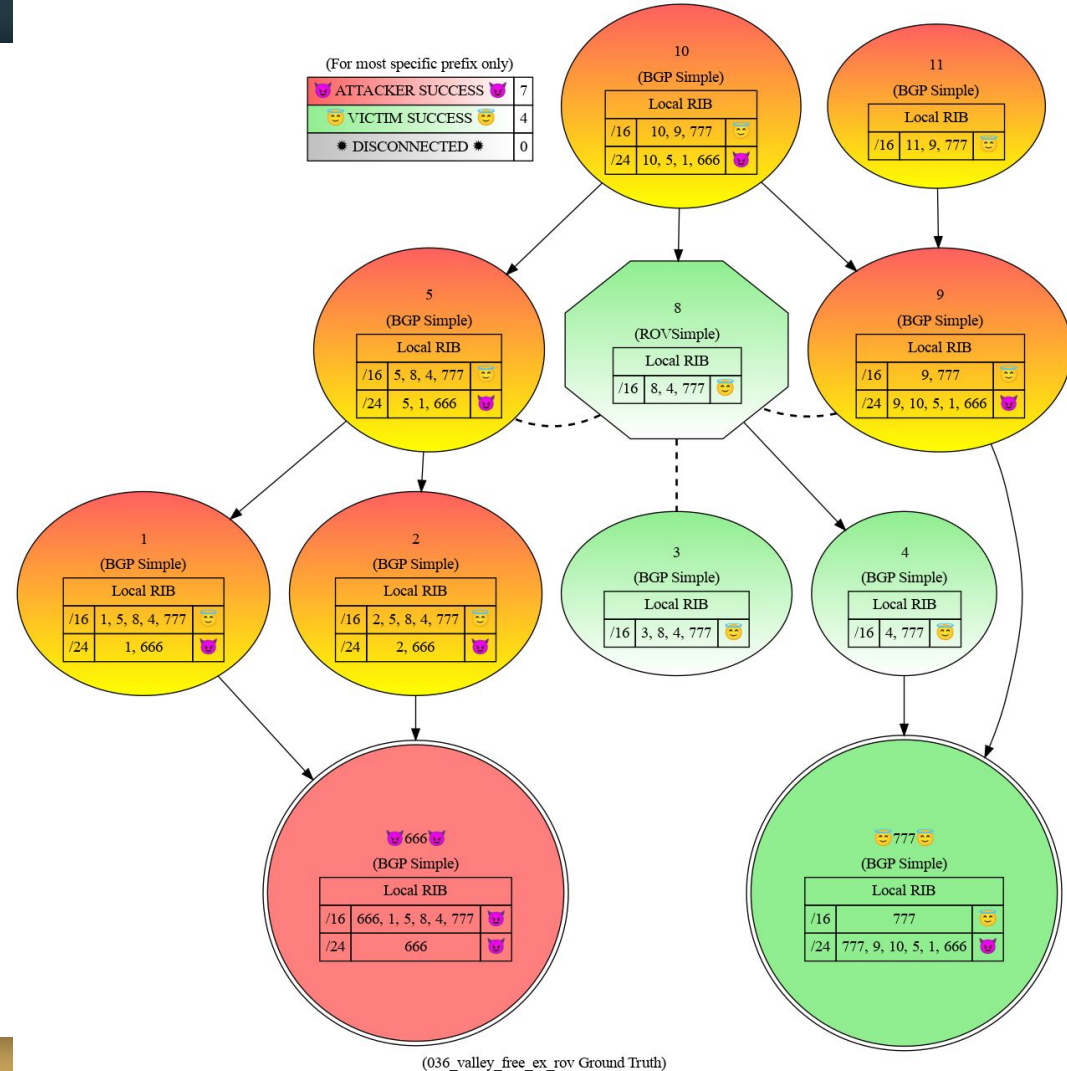
BGP Hijacking

- How does data plane traffic flow work if you have multiple prefixes for the same IP address?
 - Most specific prefix is chosen
 - This allows for BGP hijacks
 - This specific example is of a subprefix hijack



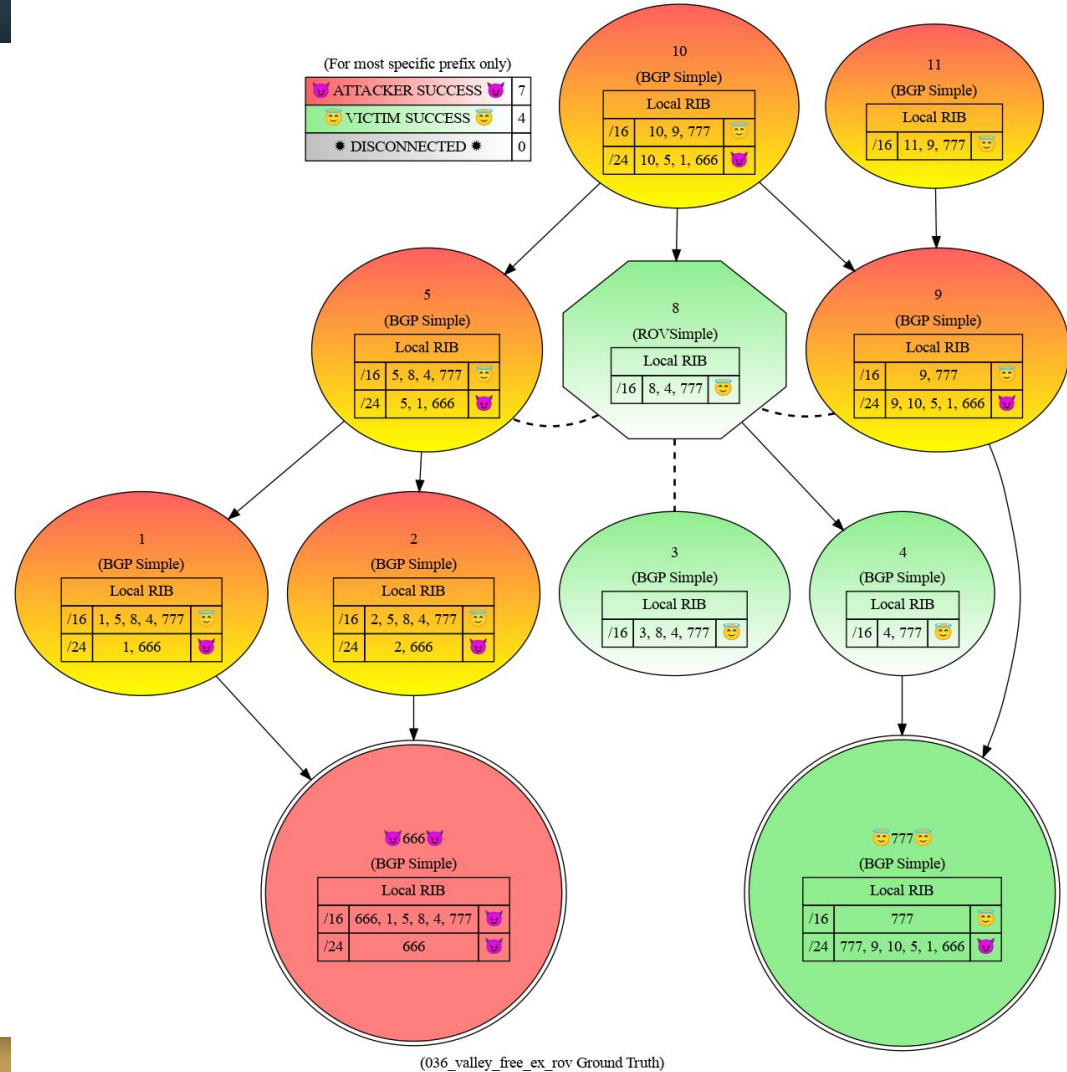
ROV

- From a very high level, Route Origin Validation (ROV) can declare announcements as valid or invalid
- Here AS 8 deploys ROV
- Since ROV declares the /24 subprefix invalid, AS 8 drops the /24 prefix. This protects:
 - AS 3
 - AS 8
 - AS 4



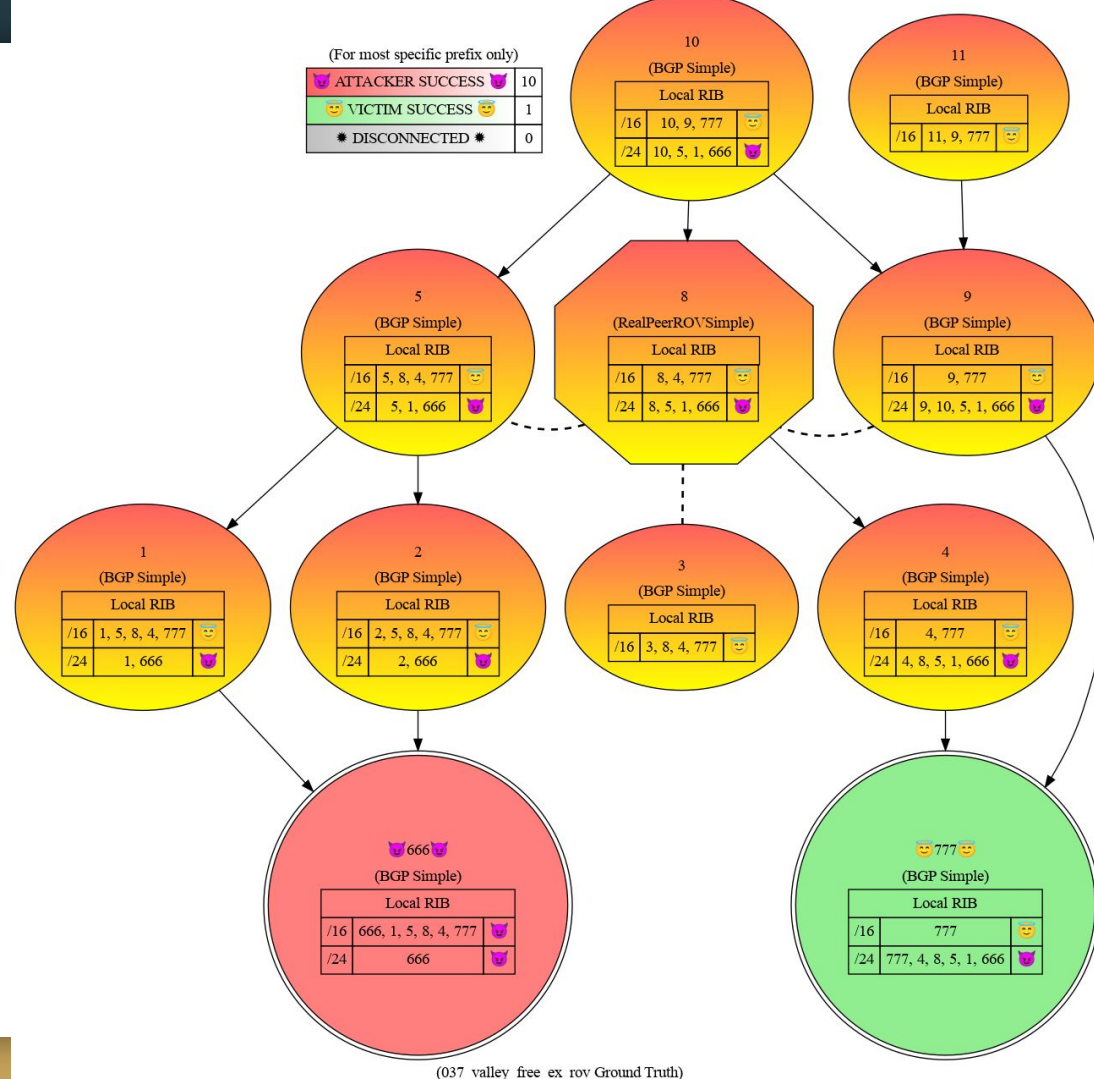
BGP_{Py}

- This is a system test included in BGP_{Py}
 - Easy to write
 - Easy to view
 - Easy to verify
- On a much larger scale, BGP_{Py} is open source and can simulate different attacks and defenses that are partially adopted
- The results can be compared with well defined benchmarks
- It's efficient enough to be run on a laptop and is easy to extend




PeerROV

- PeerROV is an ROV variant that only filters peers
- Here AS 8 deploys Peer ROV
- Peer ROV does not protect any ASes in this case



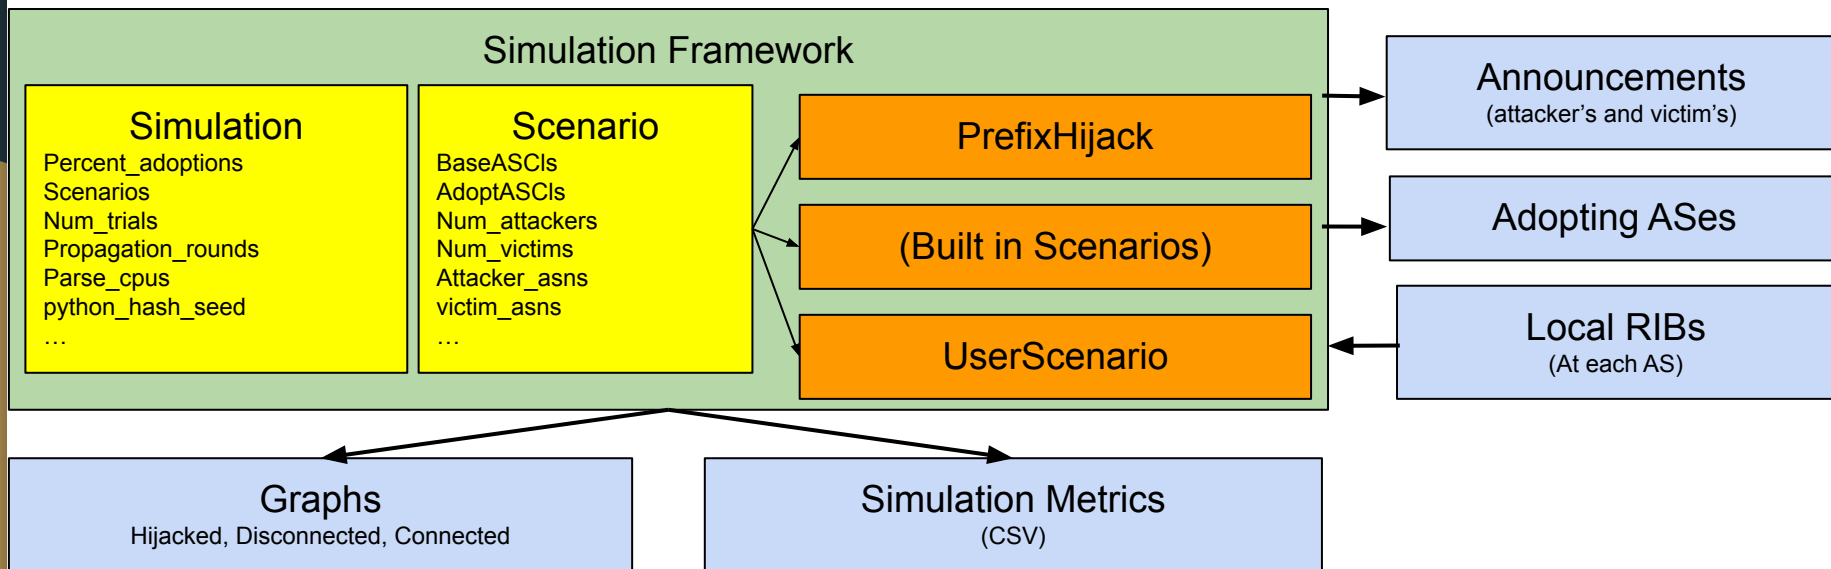
Results

- This is an example of a graph included by default in BGPy
- Here we are comparing ASes th adopting ROV vs PeerROV
- For example, at 40% adoption:
 - About 60% of ROV ASes are hijacked
 - Almost all PeerROV ASes are hijacked
- We also track
 - Percent Disconnected
 - Percent Successfully Connected



BGPpy: The Simulation Framework

The Simulation Framework is a wrapper around the Simulation Engine that facilitates the comparison of multiple security policies against attack scenarios at partial adoptions



BGPpy: The Simulation Framework (PseudoCode)

```
1 engine = SimulationEngine()
2 metric_tracker = MetricTracker()
3 for trial in range(num_trials):
4     for scenario in scenarios:
5         # Select attacker(s) and victim(s), kept consistent across scenarios
6         scenario.select_attacker_and_victim(engine)
7         # Select Adopting ASes, kept consistent across scenarios
8         # Also sets Adopting ASes, which varies from scenario to scenario
9         scenario.set_adopting_ases(engine)
10        # Seeds the attacker(s) and victim(s) announcement in the engine
11        scenario.seed_attacker_victim_announcements(engine)
12        # Propagates announcements throughout the AS topology
13        engine.run()
14        # Records metrics for graphing later
15        metric_tracker.analyze(engine)
16        # Remove announcements from the graph
17        engine.clear()
```



```

1 class SubprefixHijack(Scenario):
2     def _get_announcements(self, *args, **kwargs):
3         anns = list()
4         anns.append(
5             self.scenario_config.AnnCls(
6                 prefix=Prefixes.PREFIX.value,
7                 as_path=(victim_asn,),
8                 roa_valid_length=True,
9                 roa_origin=victim_asn,
10            )
11        )
12
13        anns.append(
14            self.scenario_config.AnnCls(
15                prefix=Prefixes.SUBPREFIX.value,
16                as_path=(attacker_asn,),
17                roa_valid_length=False,
18                roa_origin=victim_asn,
19            )
20        )
21        return tuple(anns)

```

BGPpy: Scenario

The Scenario controls attacker and defender strategies, including:

- Which ASes can be chosen as attacker/victim
- What announcements those ASes announce
- What routing policies are deployed across the adopting ASes
- Etc

BGPy: Simulation

The Simulation object controls all aspects of the simulation

- Trials, CPUs, etc
- Percent Adoptions
- Scenarios

```
1  sim = Simulation(  
2  percent_adoptions=(  
3      SpecialPercentAdoptions.ONLY_ONE,  
4      0.1,  
5      0.2,  
6      0.4,  
7      0.8,  
8      SpecialPercentAdoptions.ALL_BUT_ONE,  
9  ),  
10 scenario_configs=(  
11     ScenarioConfig(ScenarioCls=SubprefixHijack, AdoptASCls=R0VSimpleAS),  
12     ScenarioConfig(ScenarioCls=SubprefixHijack, AdoptASCls=PeerR0VSimpleAS),  
13 ),  
14 output_dir=Path("~/Desktop/paper_graphs").expanduser(),  
15 num_trials=1000,  
16 parse_cpus=12,  
17 )  
18 sim.run()
```

BGP_{Py}: ROV AS

ROV AS is an example of how you would subclass the default BGP AS

- Most extensions are fairly simple
- Here we merely extend the valid announcement check (which checks for loops in the AS path)

```
1 ✓ class ROVAS(BGPAS):  
2 ✓     def _valid_ann(self, ann: Ann) -> bool:  
3         # If ROA is invalid, ROV says announcement is invalid  
4 ✓         if ann.invalid_by_roa:  
5             return False  
6         # If ROA is valid, determine validity with BGP  
7 ✓         else:  
8             return super(ROVAS, self)._valid_ann(ann)
```

BGPpy: Peer ROV AS

Peer ROV is the same as ROV, except it only filters announcements from peers

Again fairly straightforward in BGPpy

```
1 class PeerROVAS(BGPAS):
2     def _valid_ann(self, ann: Ann) -> bool:
3         # If ROA is invalid, ROV says announcement is invalid
4         # For PeerROV, only filter by Peers
5         if ann.invalid_by_roa and ann.recv_relationship == Relationships.PEERS:
6             return False
7         # If ROA is valid, determine validity with BGP
8         else:
9             return super(PeerROVAS, self)._valid_ann(ann)
```

BGPpy: Announcement

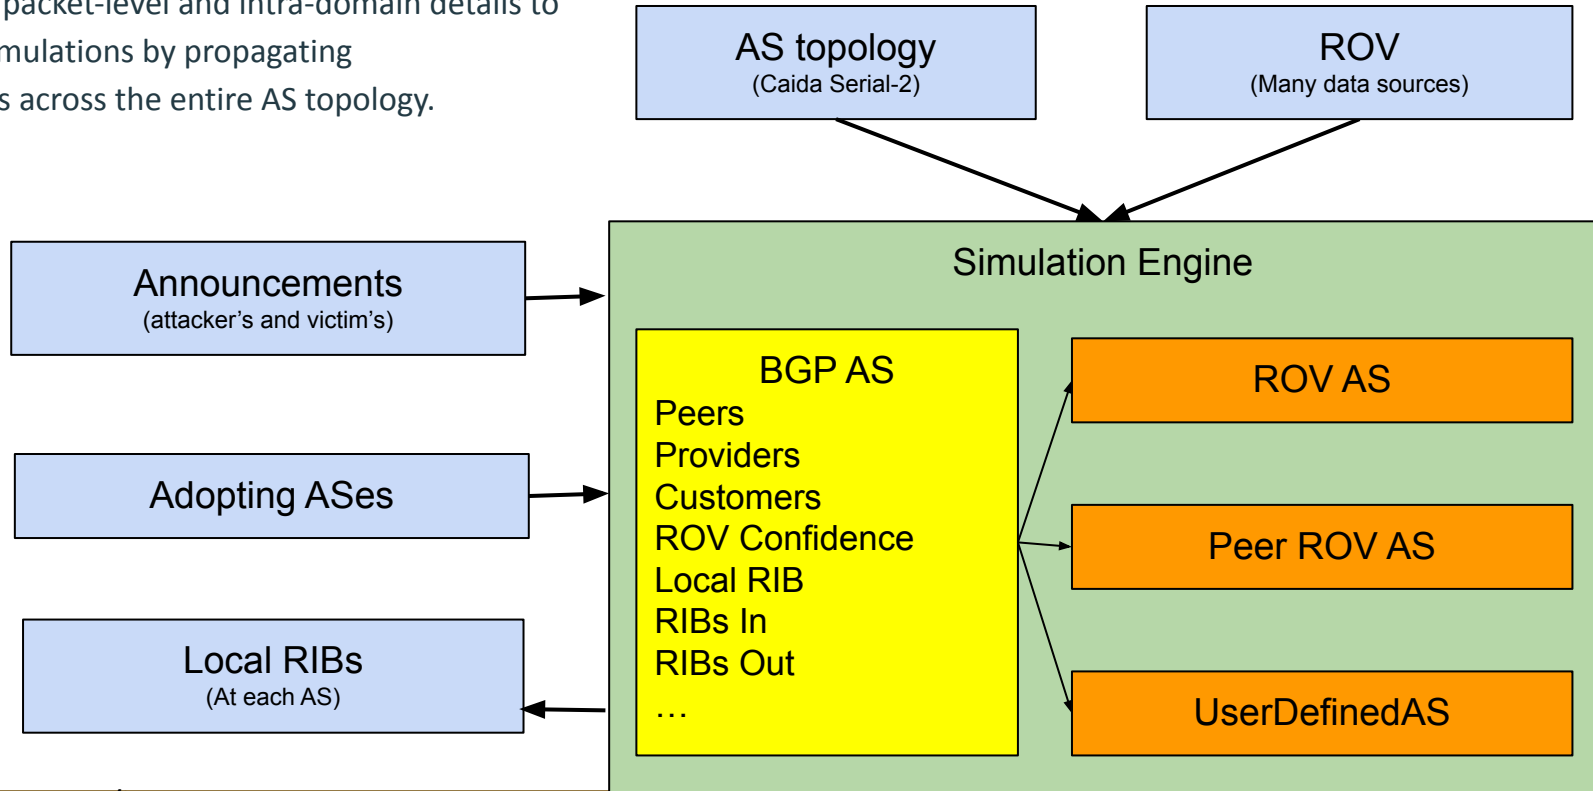
Sometimes it's desired to add extra attributes to announcements

Here is an example for ROV++, where we added a few attributes

```
1 @dataclass(frozen=True, slots=True)
2 class ROVPPAnn(Announcement):
3     holes: tuple[str] = ()
4     blackhole: bool = False
5     # V3 attributes
6     preventive: bool = False
7     attacker_on_route: bool = False
```


BGPpy: The Simulation Engine

Abstracts away packet-level and intra-domain details to perform BGP simulations by propagating announcements across the entire AS topology.



BGPpy: The Simulation Engine (ROV Data)

- Optionally, real world ROV data can be used in the simulations
- Data Sources included by default:
 - Rov.rpki.net
 - Isbgpsafeyet.com
 - Revisiting RPKI Route Origin Validation on the Data Plane

BGPpy: The Simulation Engine (Propagation)

- First, announcements are inserted at the victim(s) and attacker(s)
- Then, announcements are propagated throughout the AS topology:
 - First from customers to providers
 - Second from peer to peer
 - Third from providers to customers
- We converge in $O(E)$ time in a single round of propagation
 - This allows us to write BGPpy in Python and still be able to run it on a laptop
 - If users would like to propagate for more than 1 round, they can set the number of rounds in a parameter

BGPpy: Future Work

- We just released a new version with some improvements, and will continue to iterate
- The next big leap forward will be likely using rust bindings with PyO3 for massive speed improvements
- More real world data can be added to make simulations more realistic
 - IXP data
 - BGP Communities
 - AS Blacklists
 - etc

Thank you! Questions?

- Contact Info:
 - Justin Furuness: jfuruness@gmail.com
 - Cameron Morris: cameron.morris@uconn.edu
 - Reynaldo Morillo: reynaldo.morillo@uconn.edu
 - Dr. Amir Herzberg: amir.herzberg@gmail.com
 - Dr. Bing Wang: bing@uconn.edu
- Link to BGPy: github.com/jfuruness/bgpy