Detecting Unsafe BGP Policies in a Flexible World

Debbie Perouli, Timothy G. Griffin, Olaf Maennel, Sonia Fahmy, Cristel Pelsser, Alexander Gurney, and Iain Phillips

ICNP

October 31, 2012

Balance Safety and Flexibility in Policy Based Routing

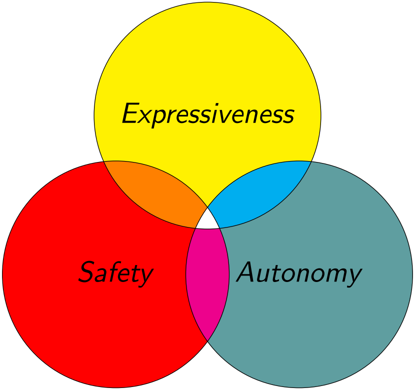

Expressiveness	ISPs innovate in policies as customer needs evolve
Safety	the protocol always converges to a unique routing solution
Autonomy	ISPs configure their network without global coordination

Autonomous and Safe: Prefer Customer over Peer

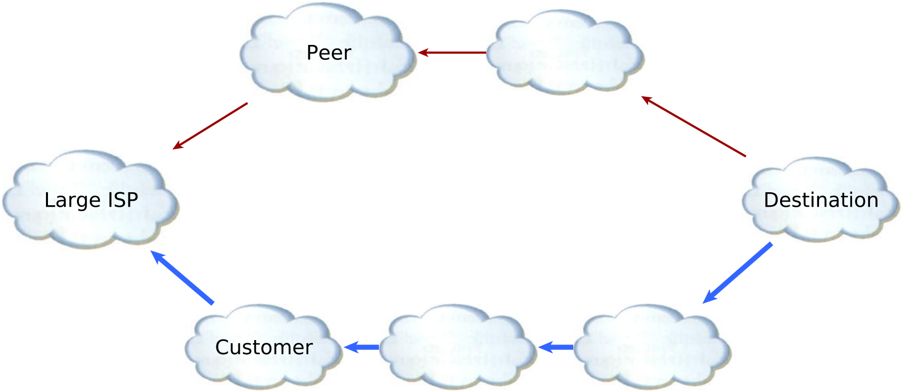

Preferred routing path for the Large ISP is in blue.

Gao and Rexford. Stable Internet Routing Without Global Coordination. SIGMETRICS 2000.

Relaxing(?) Safety: Prefer Peer to Avoid Specific AS

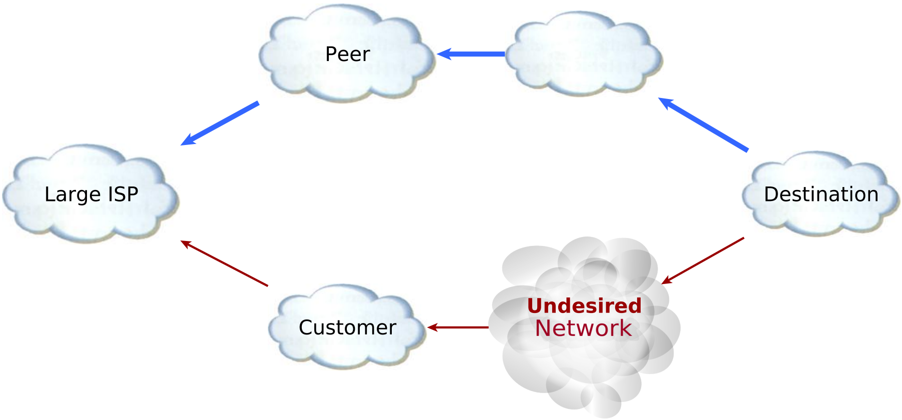

Preferred routing path for the Large ISP is in blue.

Relaxing Autonomy: Backup Policy Requires Coordination

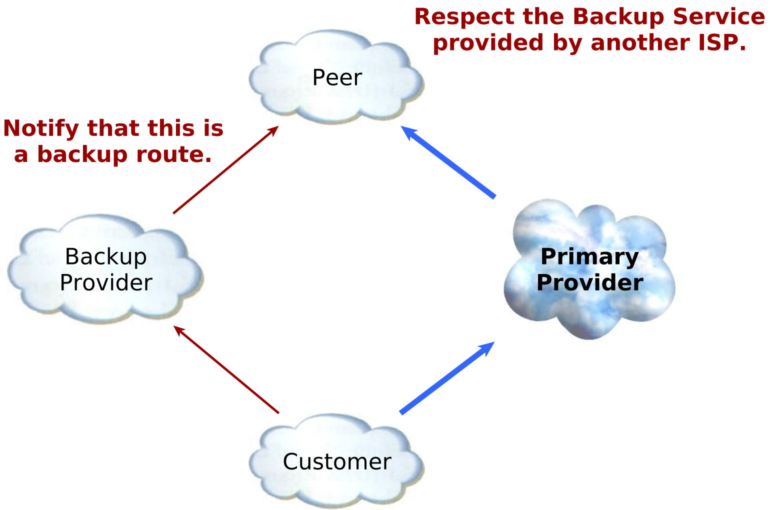

Gao and Rexford. Stable Internet Routing Without Global Coordination. SIGMETRICS 2000.

Griffin and Huston. BGP Wedgies. RFC 4264.

The Stable Paths Problem (SPP)

- Provides: a *sufficient condition for safety* (acyclicity of dispute digraph)

- Requires:
  - knowledge of all potential routing paths
    i.e. all paths permitted by the policies of each router
  - strict ordering of the potentially available paths of each router

Griffin, Shepherd, and Wilfong. Policy Disputes in Path-Vector Protocols. ICNP 1999.

The Stable Paths Problem (SPP)

- Provides: a *sufficient condition for safety* (acyclicity of dispute digraph)

- Requires:
  - knowledge of all potential routing paths
    i.e. all paths permitted by the policies of each router
    - Need for router configuration files, which ISPs consider proprietary.
    - In the worst case, path enumeration is an intractable problem.
  - strict ordering of the potentially available paths of each router
    - Requires a lot about the internals of an ISP, like IGP distances.
    - Depends on vendor specific details (e.g. tie break).
    - Including MED is computationally expensive, if not infeasible.

Extended SPP

- Provides: a *sufficient condition for safety*

- Requires:
  - knowledge of all potential routing paths
    i.e. all paths permitted by the policies of each router
  - strict ordering of the potentially available paths of each router

Enumerate All Paths Among Some ISPs Only

A small number of ISPs share their configurations with trusted third party.

Execute the BGP Decision Process Steps as Needed

Allow a router to equally prefer two paths, even if they do not share the next-hop.

Extended SPP

- Provides: a *sufficient condition for safety*
- Requires:
  - knowledge of all potential routing paths
    i.e. all paths permitted by the policies of each router
  - strict ordering of the potentially available paths of each router

Enumerate All Paths Among Some ISPs Only

A small number of ISPs share their configurations with trusted third party.

Execute the BGP Decision Process Steps as Needed

Allow a router to equally prefer two paths, even if they do not share the next-hop.

Contributions

- We define a new data structure, the Multipath Digraph ($\mathcal{MD}$).
- We prove the relationship $\mathcal{MD}$ has with the Paths Digraph[1,2] ($\mathcal{PD}$).
- We provide a methodology for ensuring BGP safety
  - assuming nothing about the **policies** ISPs use
  - assuming nothing about the **Internet graph** structure (hierarchical/flat)
  - requiring **no change to BGP**
  - detecting not only instability but also **multiple stable states**
  - relaxing SPP requirements so that router configuration information is used **only as needed**
  - pointing out safety risks when paths are only **partially known**

1: Gurney, Jia, Wang, and Loo. Partial Specification of Routing Configurations. WRiPE 2011.

2: Sobrinho. Network Routing with Path Vector Protocols: Theory and Applications. SIGCOMM 2003.

Strict SPP Example

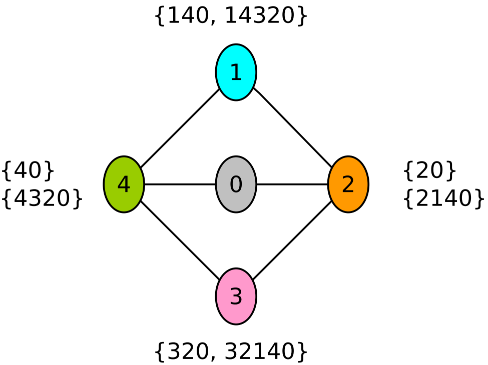

Network Topology

Node 0	Destination
$\{p_1,\ p_2\}$	p_1, p_2 equally preferred
$\{p_1\}$ $\{p_2\}$	p_1 preferred over p_2

Strict SPP Example

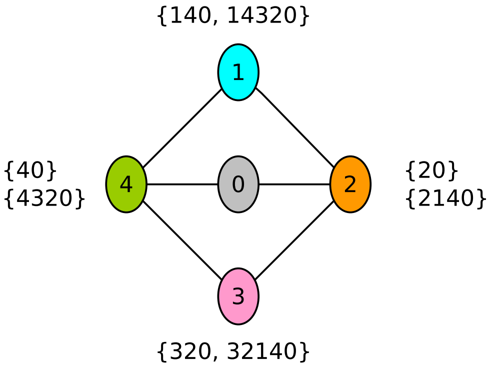

Network Topology

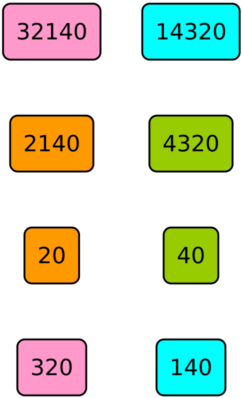

$\mathcal{PD}$

Strict SPP Example

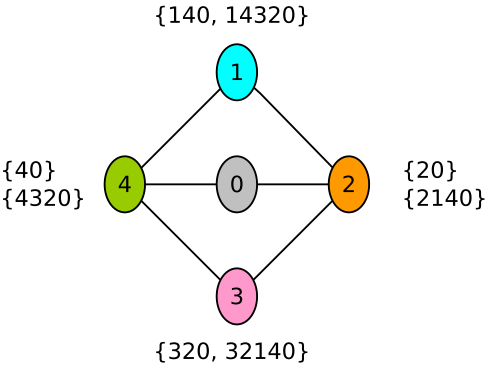

Network Topology

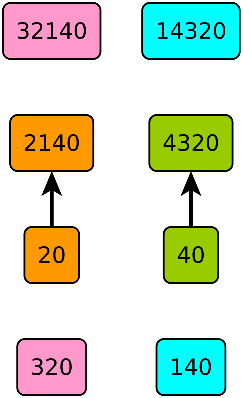

$\mathcal{PD}$

Strict SPP Example

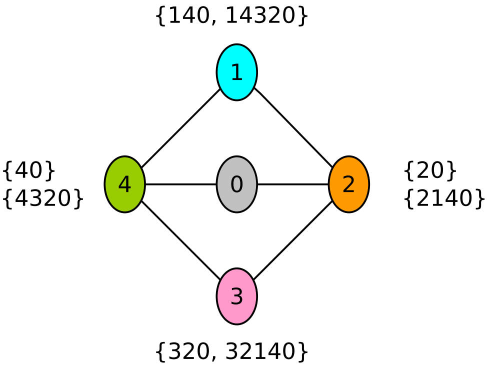

Network Topology

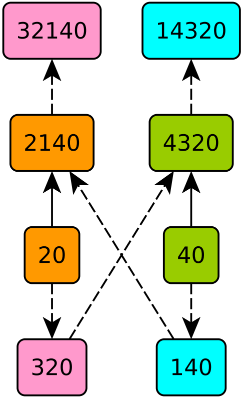

$\mathcal{PD}$

Strict SPP Example

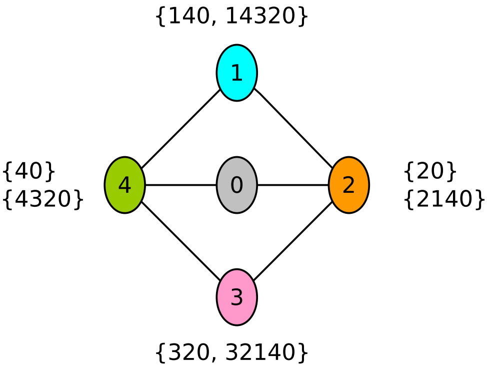

Network Topology

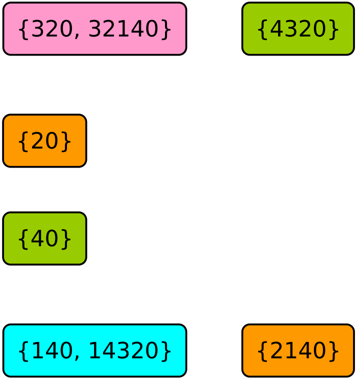

$\mathcal{MD}$

Strict SPP Example

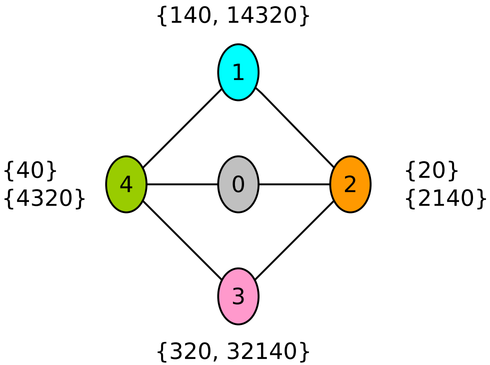

Network Topology

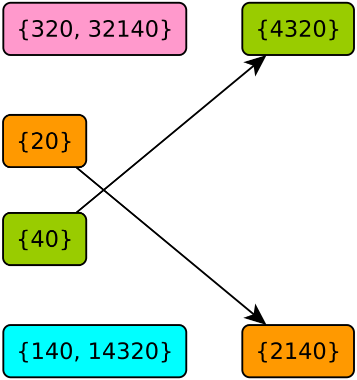

$\mathcal{MD}$

Strict SPP Example

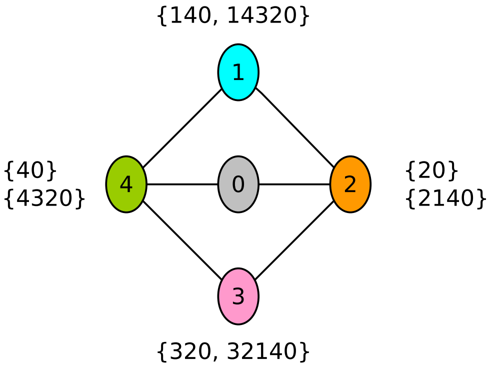

Network Topology

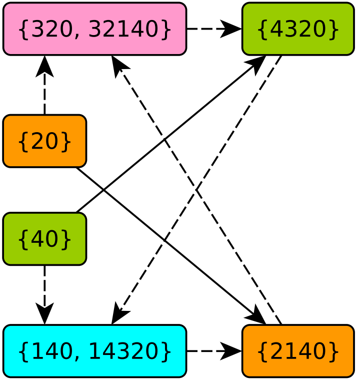

$\mathcal{MD}$

$\mathcal{MD}$ has Cycle, $\mathcal{PD}$ is Acyclic

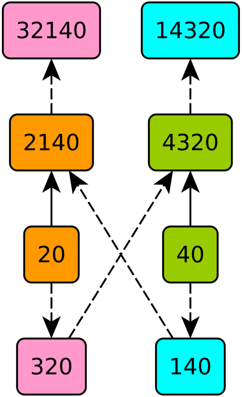

$\mathcal{PD}$

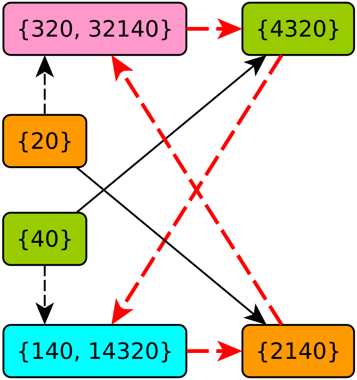

$\mathcal{MD}$

The Only Refinement That Has a Cycle

Both $\mathcal{PD}$ and $\mathcal{MD}$ will be:

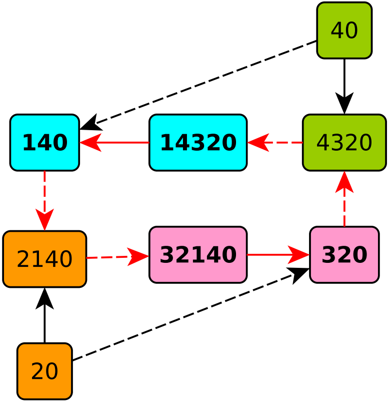

Refinement: specification where every router has its paths strictly ordered

A Methodology for Safety (I)

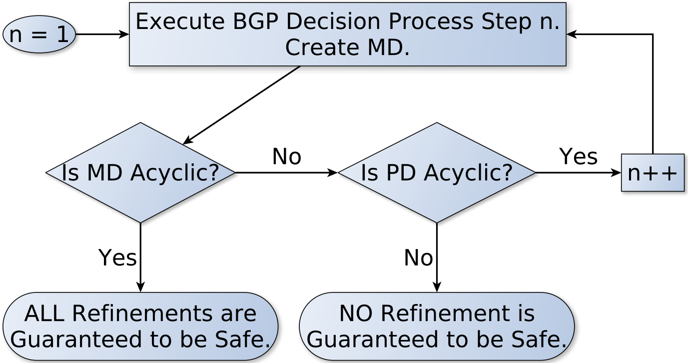

MD: Multipath Digraph

PD: Paths Digraph

A Methodology for Safety (II)

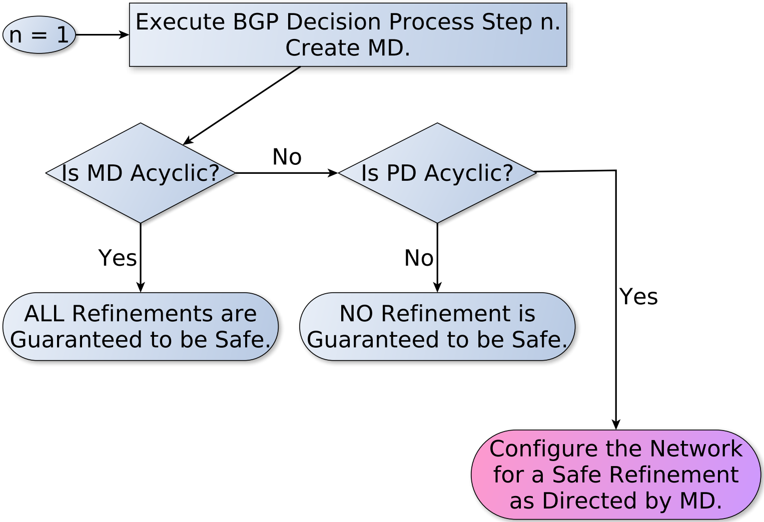

Example

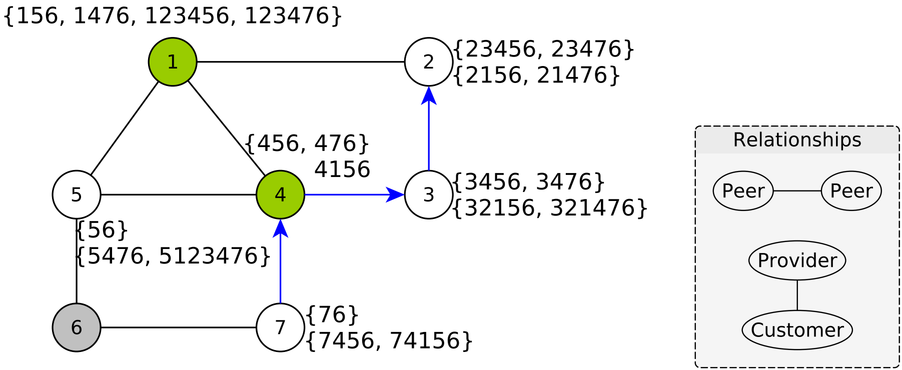

Specification of ASes with Node 6 as Destination

Nodes 1, 4: prefer peer routes equally to customer

Sessions 7 $\rightarrow$ 4, 4 $\rightarrow$ 3, 3 $\rightarrow$ 2: announce peer routes (plus customer)

$\mathcal{MD}$ Has No Cycle

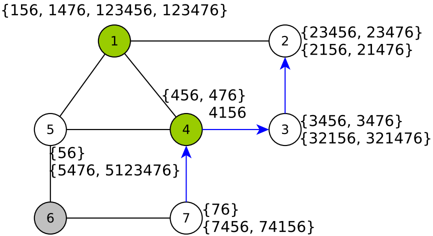

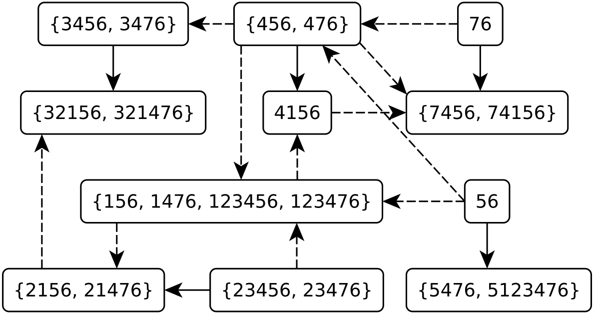

$\mathcal{MD}$ with Partial Information

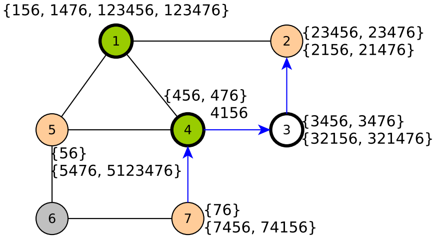

Group $\mathcal{K}$: Nodes 1, 3, 4
Known configurations

Group $\mathcal{U}$: Nodes 2, 5, 7
Unknown configurations

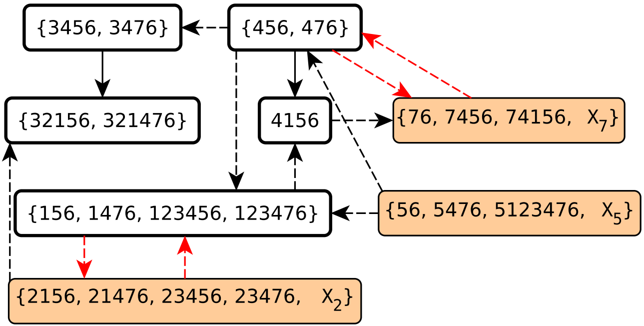

Conclusion

- ISPs can implement a richer set of BGP policies without sacrificing safety and determining themselves the level of autonomy.
- The complexity of the SPP safety analysis can be reduced by partially executing the BGP decision process without losing accuracy.
- Operators receive feedback even when paths are only partially known.
- We plan to implement a tool that evaluates the proposed approach. See Poster Session.

Questions?

Thank you

Debbie Perouli
depe@purdue.edu

Relationship of Cycles in $\mathcal{MD}$ and $\mathcal{PD}$

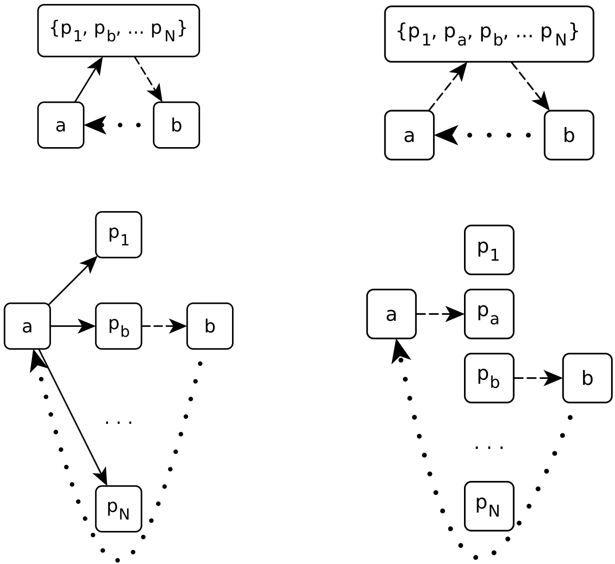

Double Backup Wedgie (I)

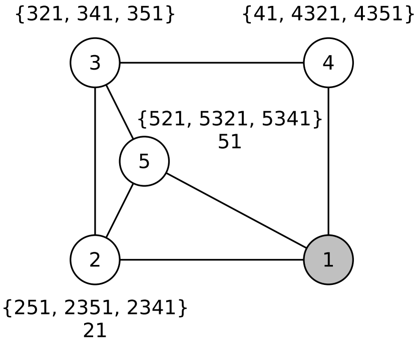

Specification

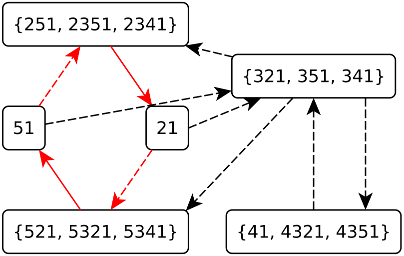

$\mathcal{MD}$

Double Backup Wedgie (II)

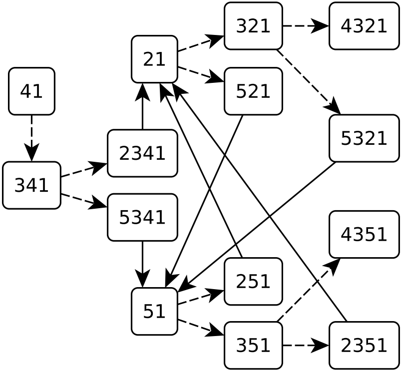

$\mathcal{PD}$

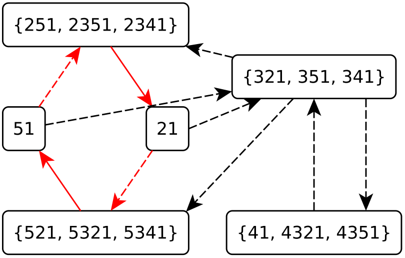

$\mathcal{MD}$