

BigBug

Practical Concurrency Analysis for SDN



Roman May
Ahmed El-Hassany
Laurent Vanbever
Martin Vechev

<http://sdnracer.ethz.ch>

ETH zürich

SDN **memory model**

Flow Tables are memory locations

Flow Table lookups are memory read operations

FlowMod messages are write operations

Events **ordering** in SDN

Ordered by **causality**

A flow table miss generates a Packet In event

Ordered by **barriers**

An OpenFlow Barrier request between two FlowMods messages

Partial events ordering is captured by **Happens-Before** model

What about **unordered** events?

SDN switches is free to process events at any order

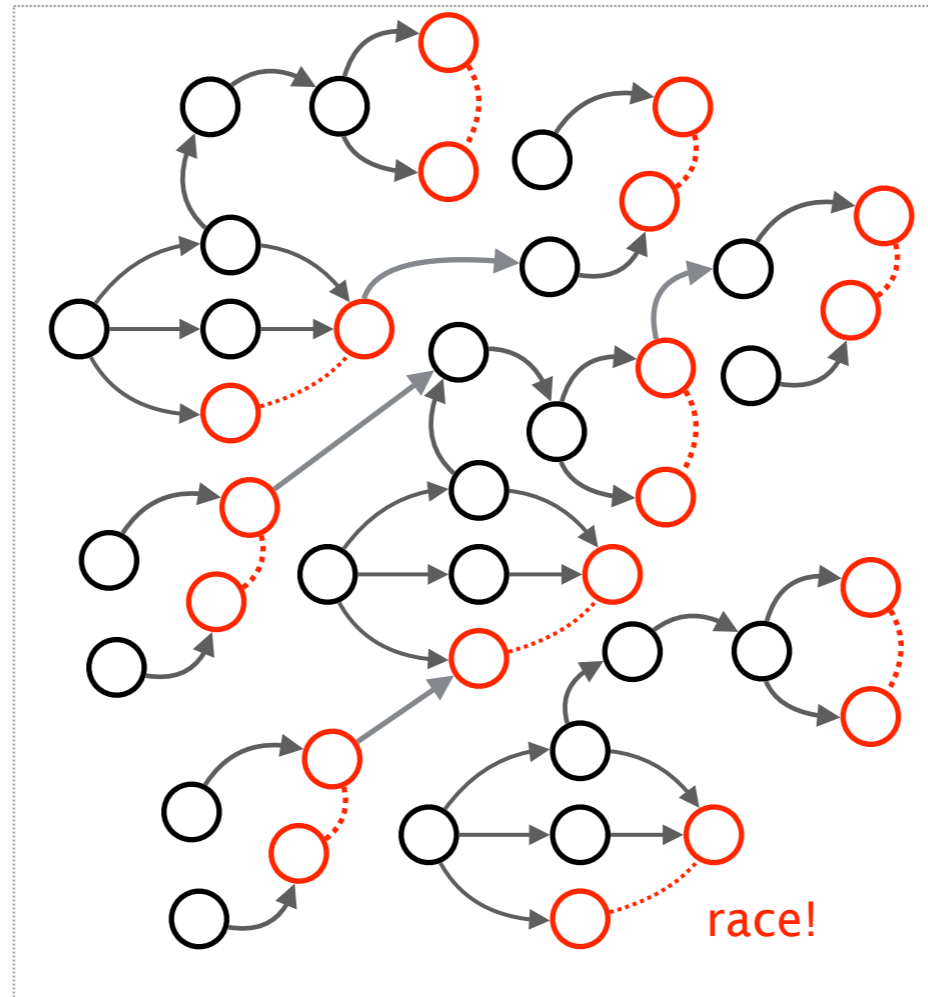
PacketOut can be processed before FlowMod

Unordered events may cause serious bugs in the network

To mention few: forwarding loops, policy violations, etc..

concurrency violations

SDN concurrency **violations**




100,000s of events

1,000s of reported **concurrency violations**

SDNRacer: PLDI'16 & SOSR'15

SDN concurrency **violations**

**External
Concurrency
Analyzer**




**concurrency
violations**

10—1000s

reports

SDN concurrency **violations**


**External
Concurrency
Analyzer**



**concurrency
violations**

10—1000s

reports




**representative
per-violation
graphs**

1 to 6

reports


BigBug Pipeline

BigBug

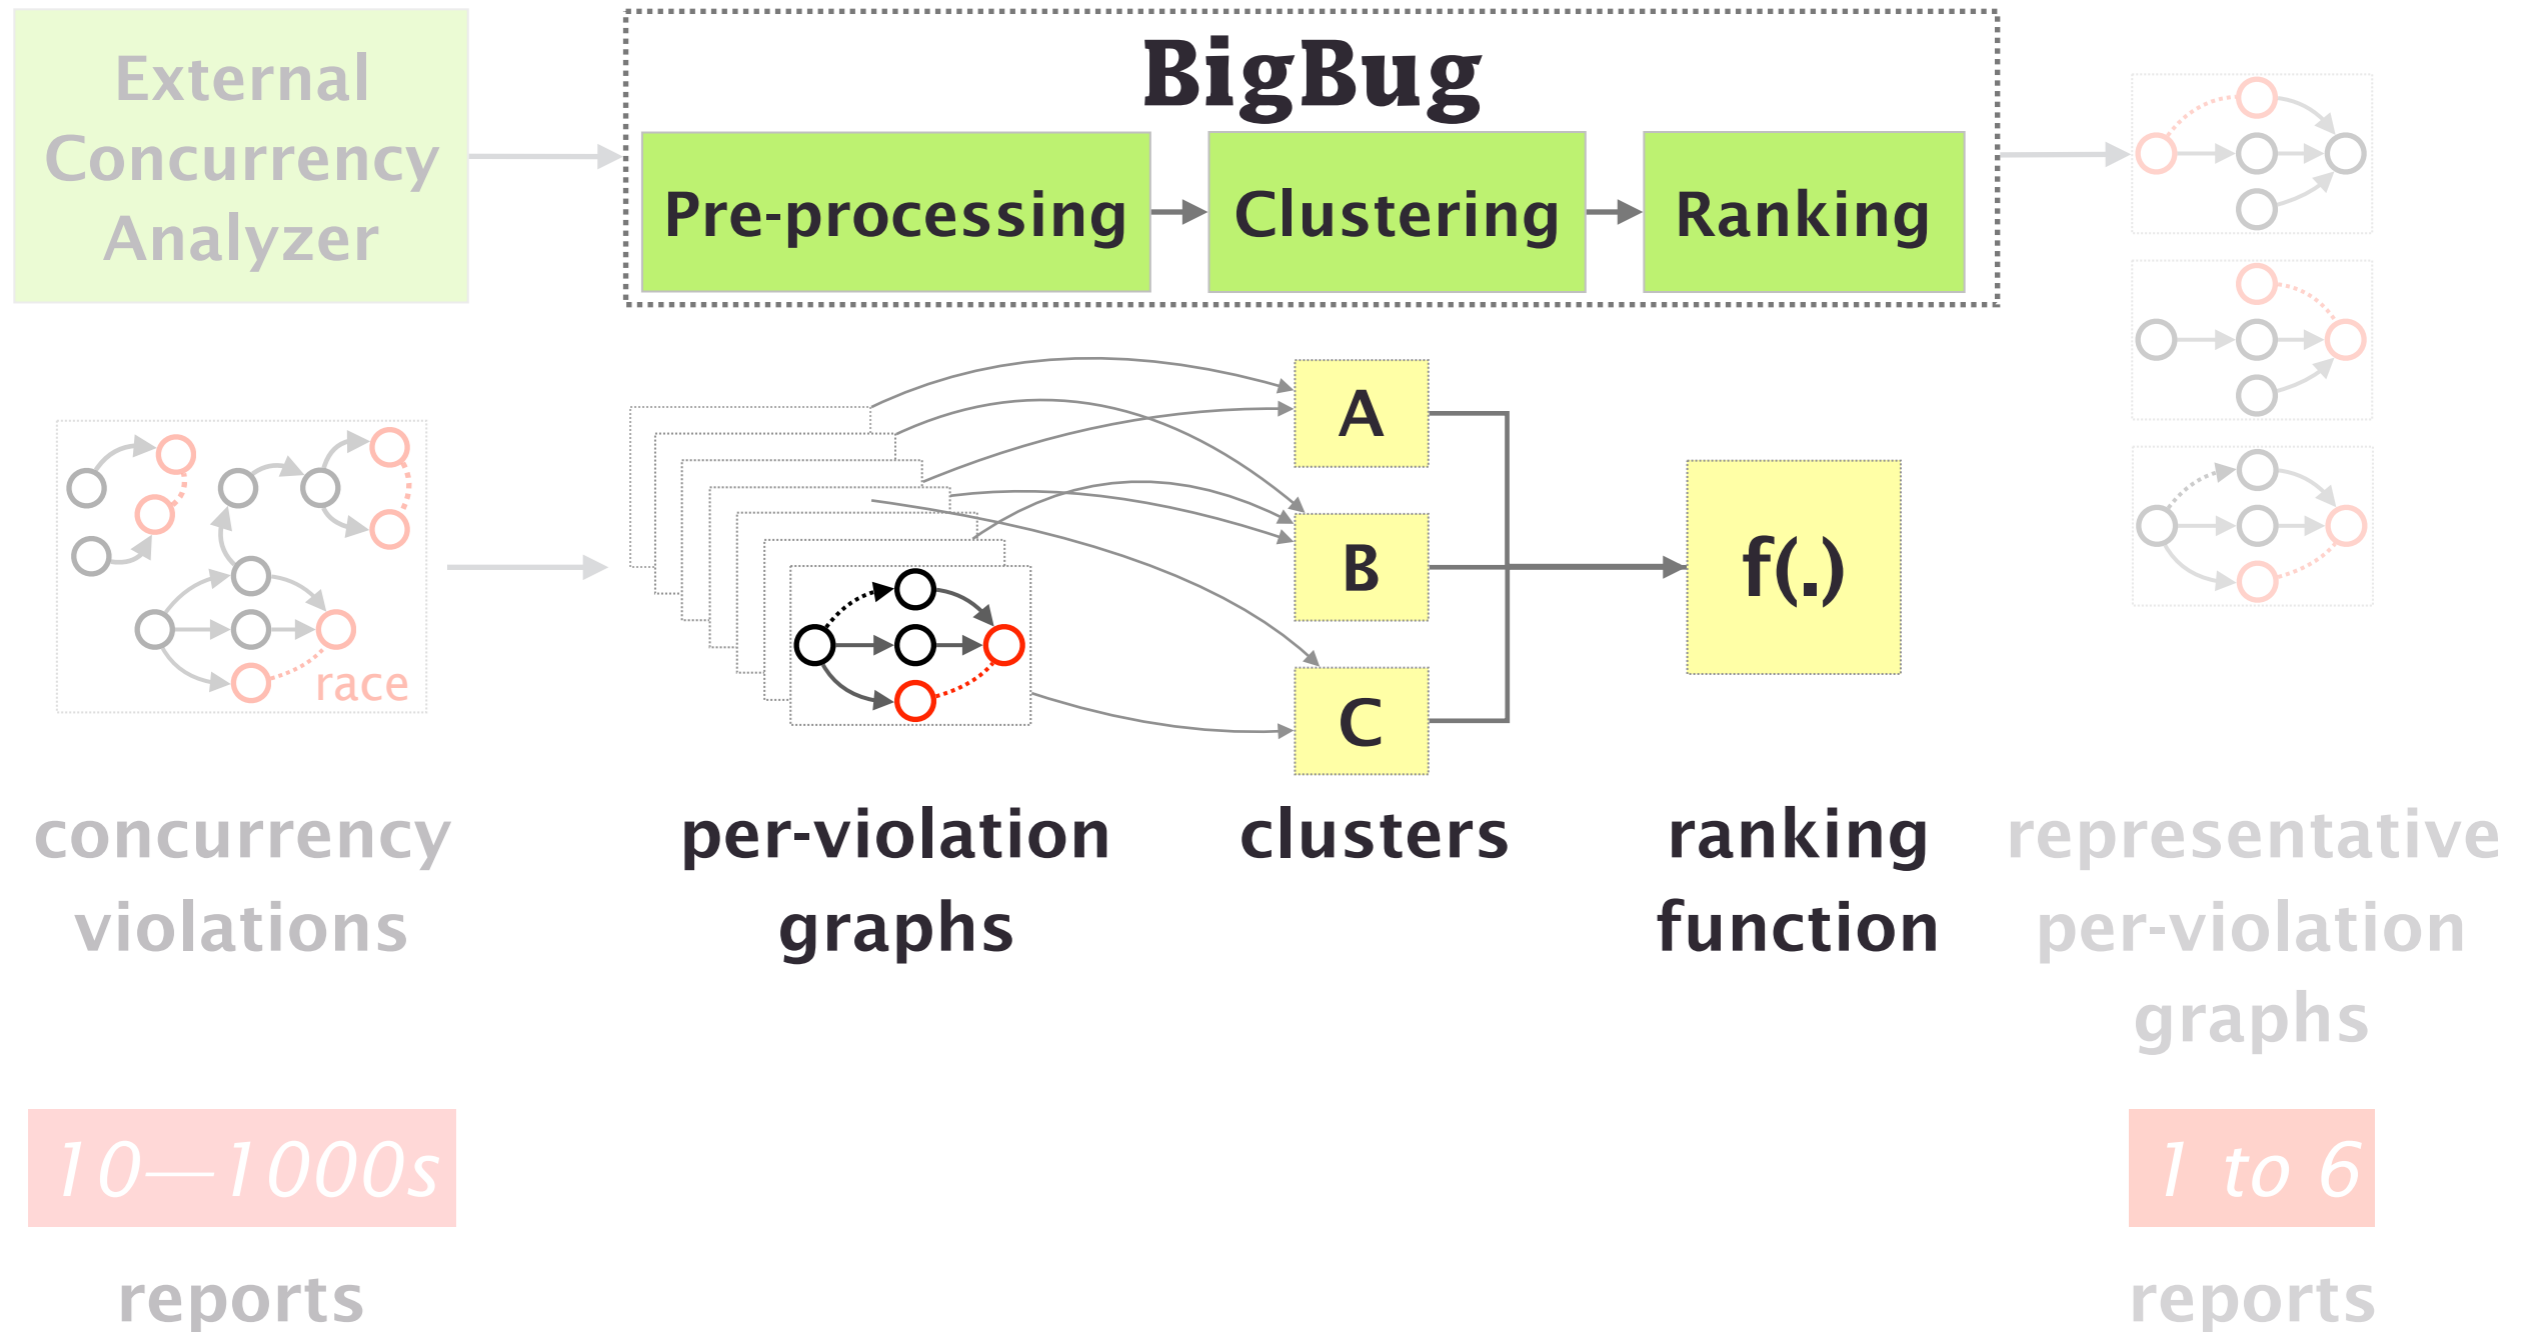


BigBug Pipeline

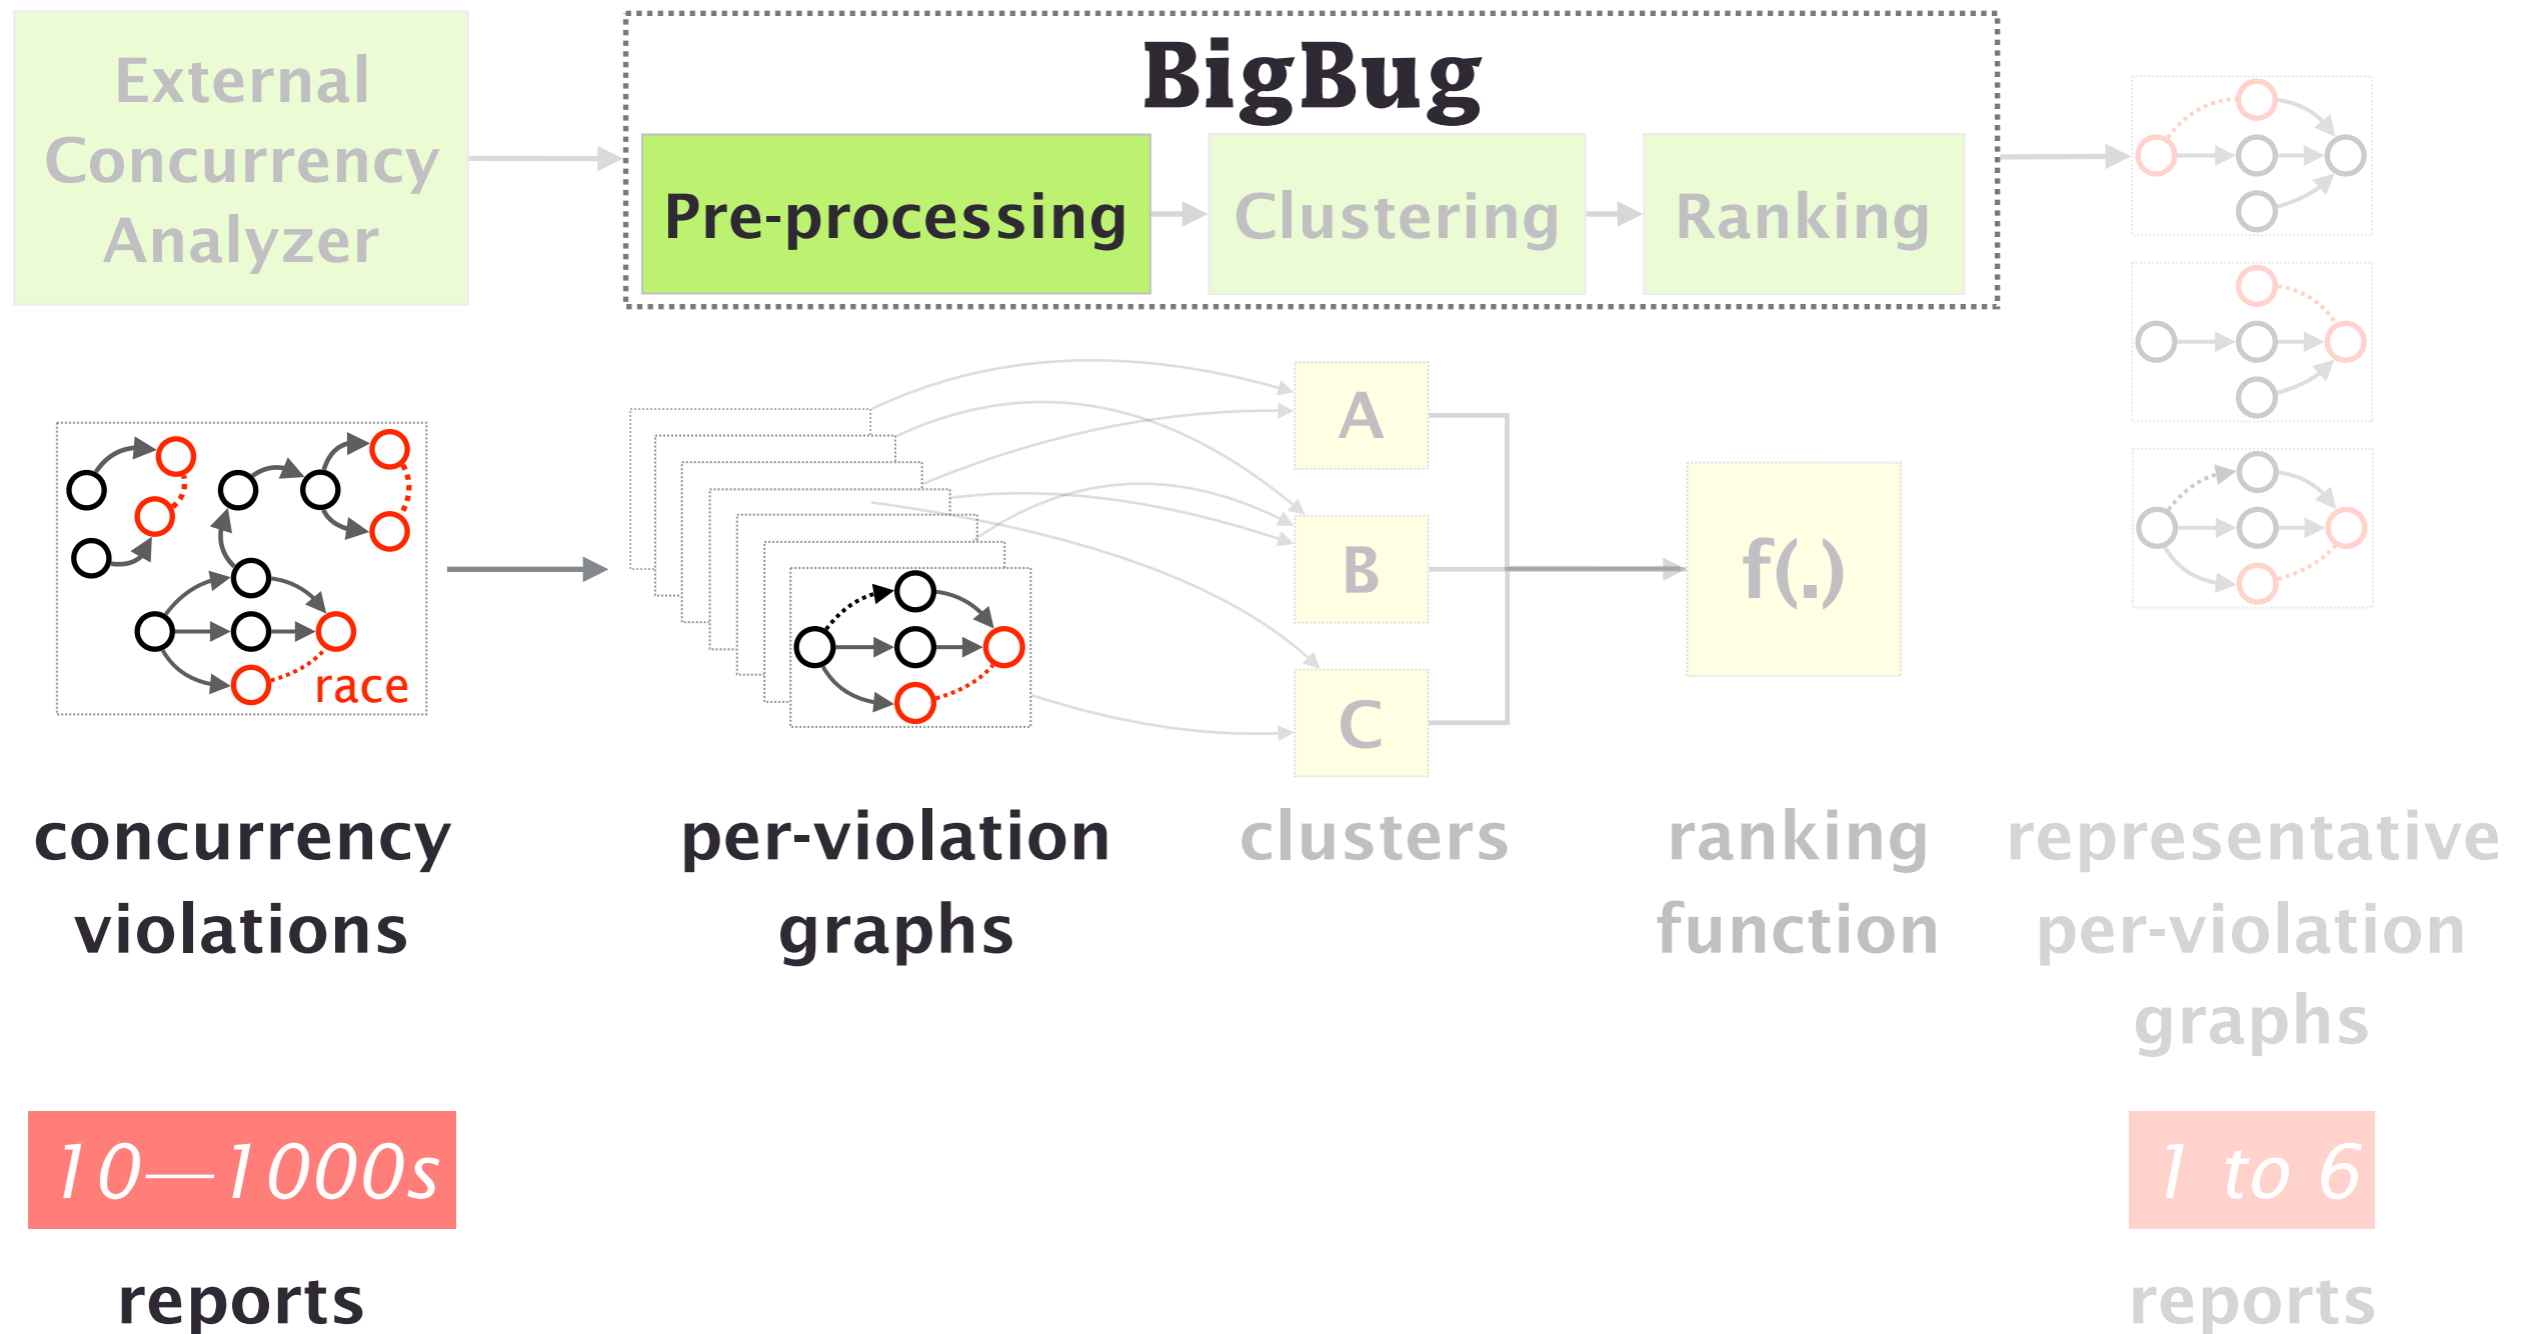
BigBug



BigBug Pipeline



BigBug Pipeline



BigBug


Pre-processing



Clustering




Ranking



100,000s of events


1,000s of reported concurrency violations



100,000s of events

1,000s of reported concurrency violations


Extract a **subgraph** per each **violation**.



1,000s of reported concurrency violations




Extract a **subgraph** per each **violation**.



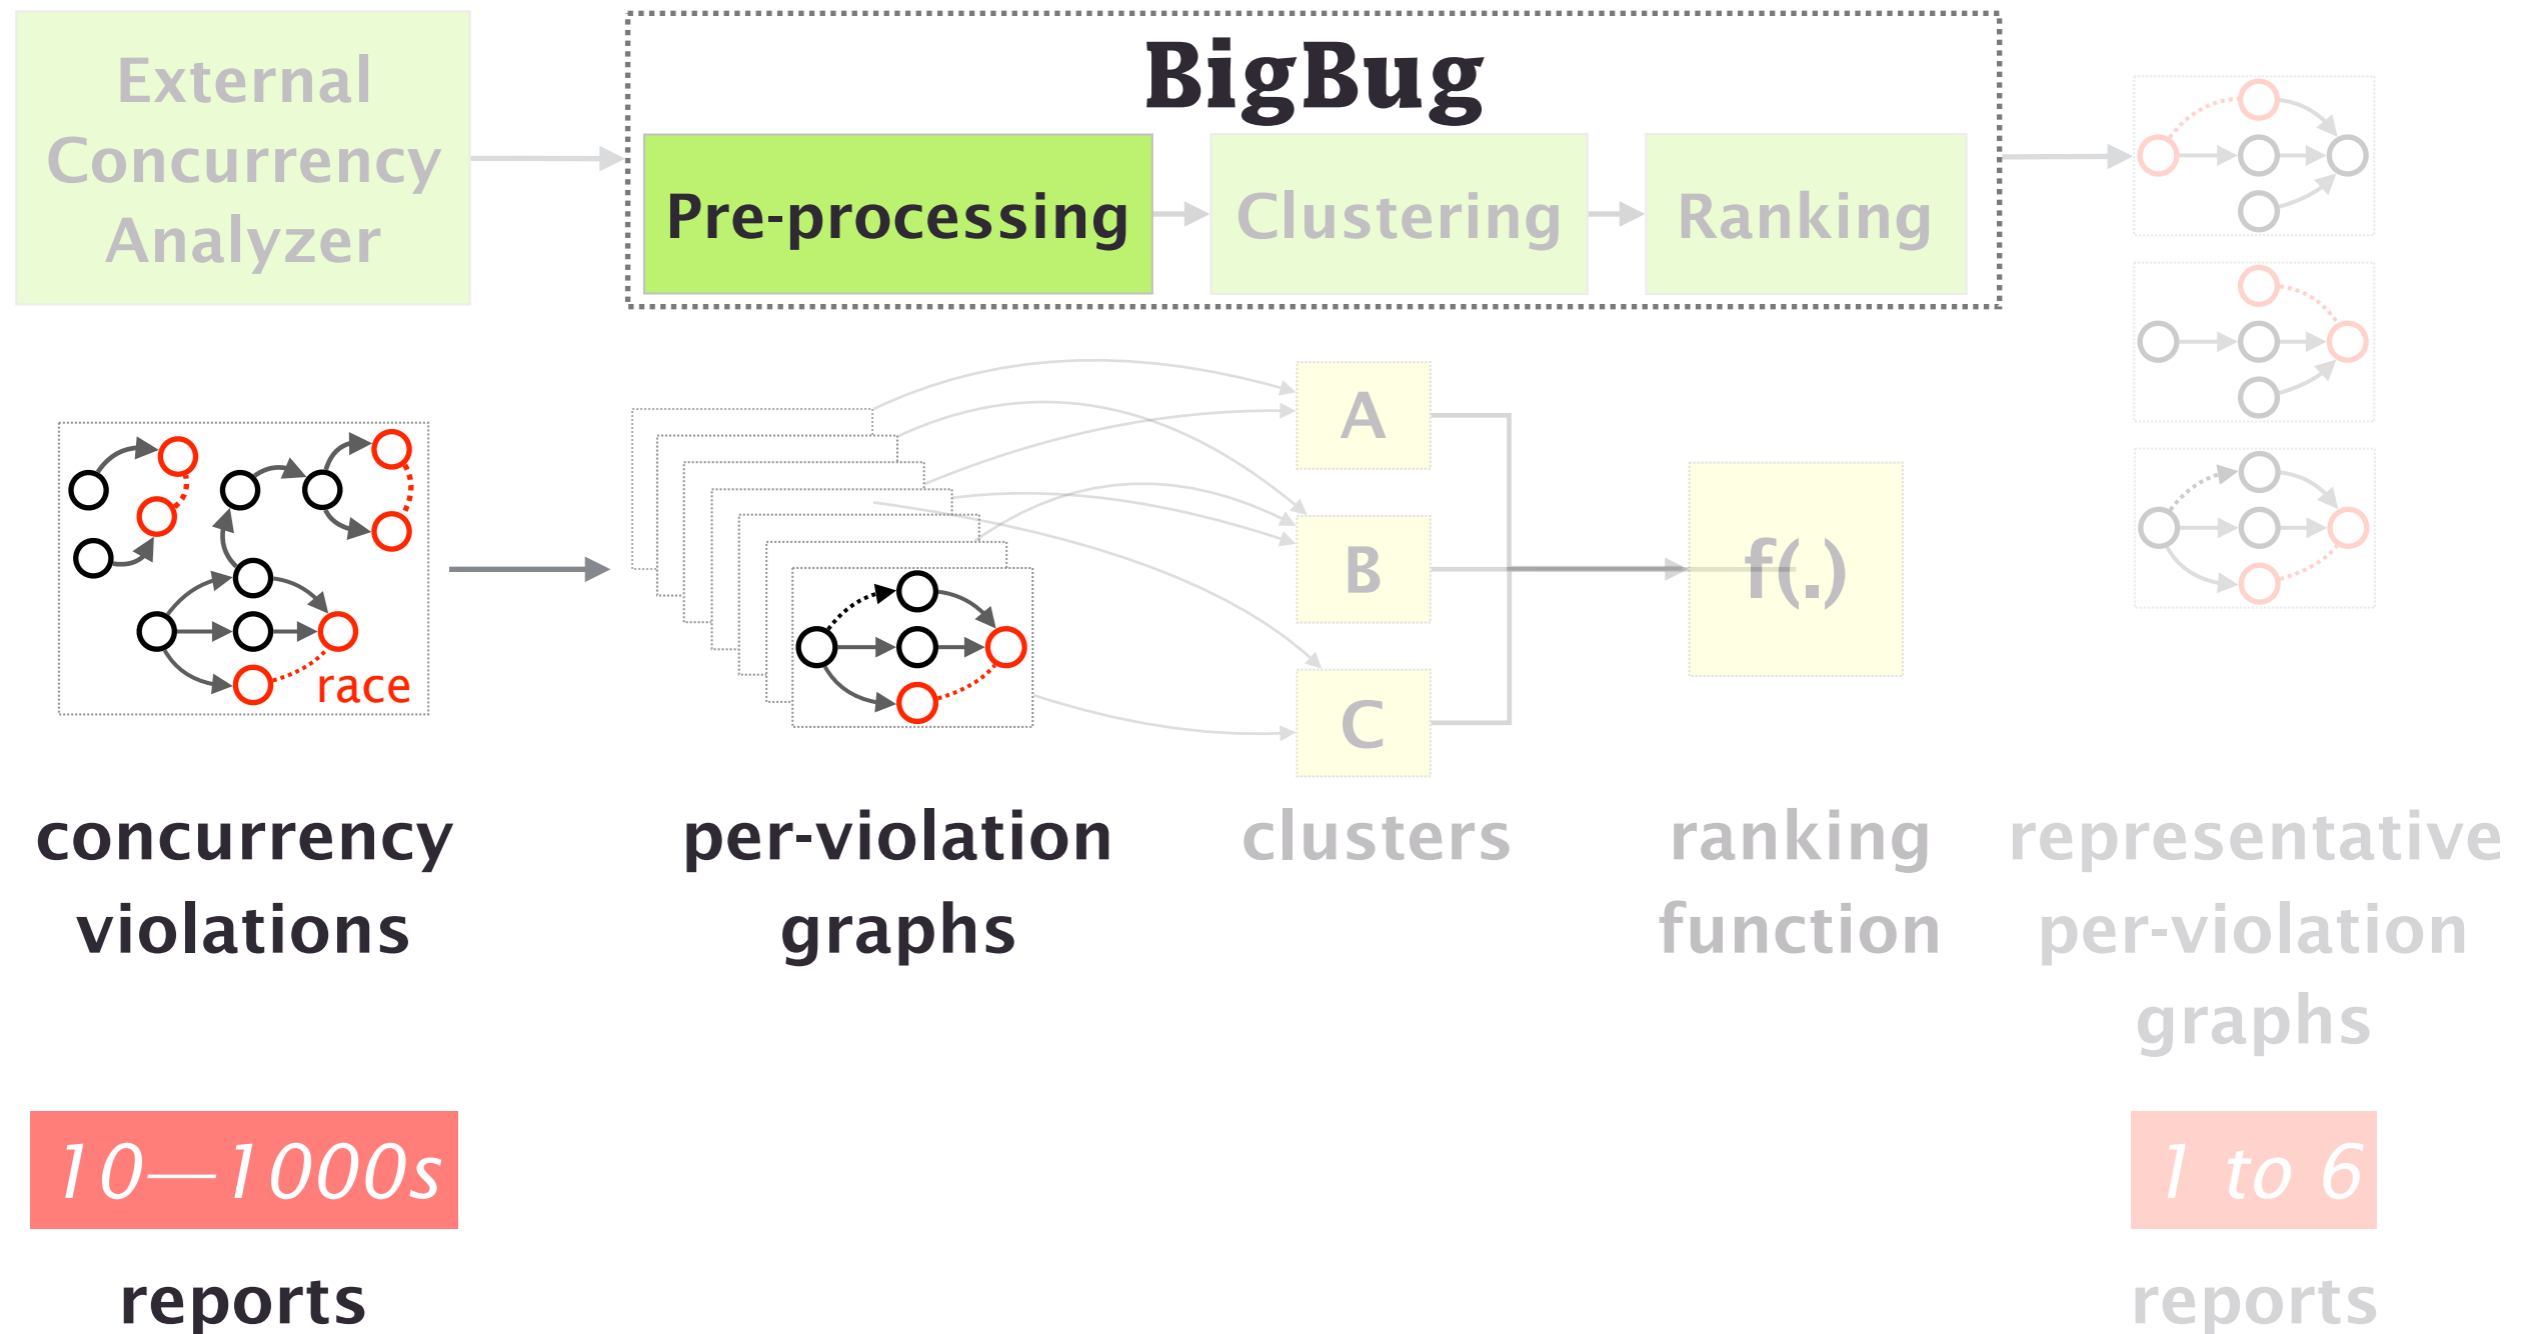
100,000s of events

1,000s of reported concurrency violations

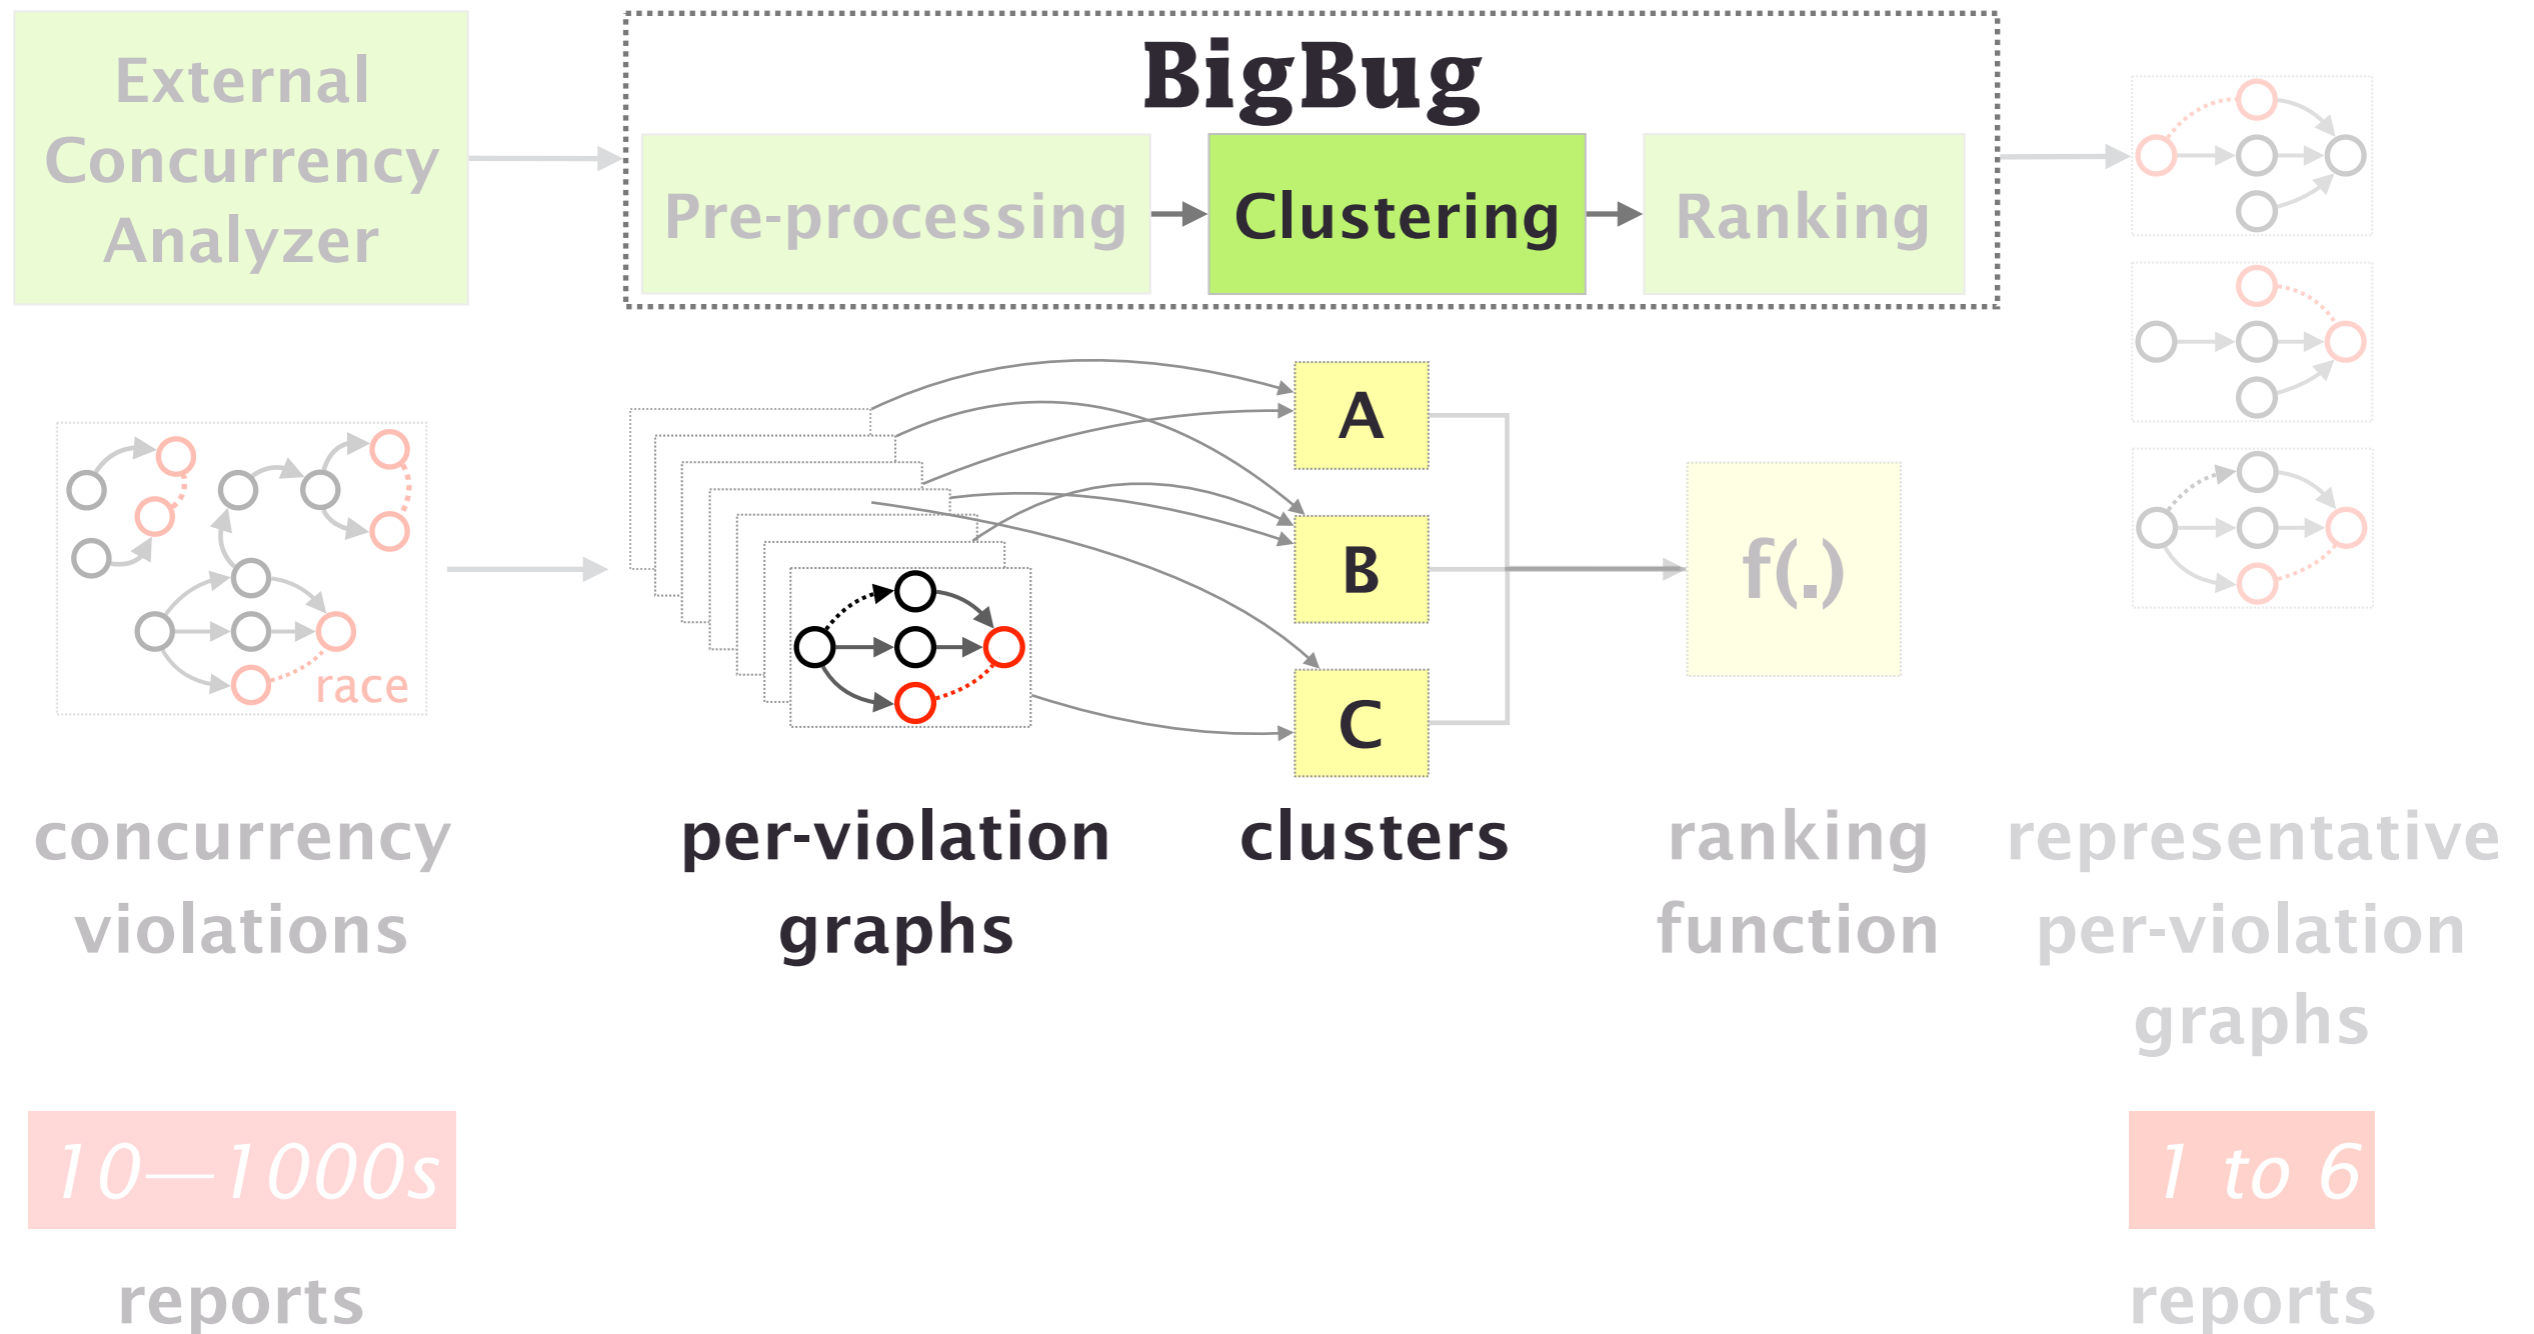


1,000s of violations subgraph

BigBug Pipeline




BigBug Pipeline





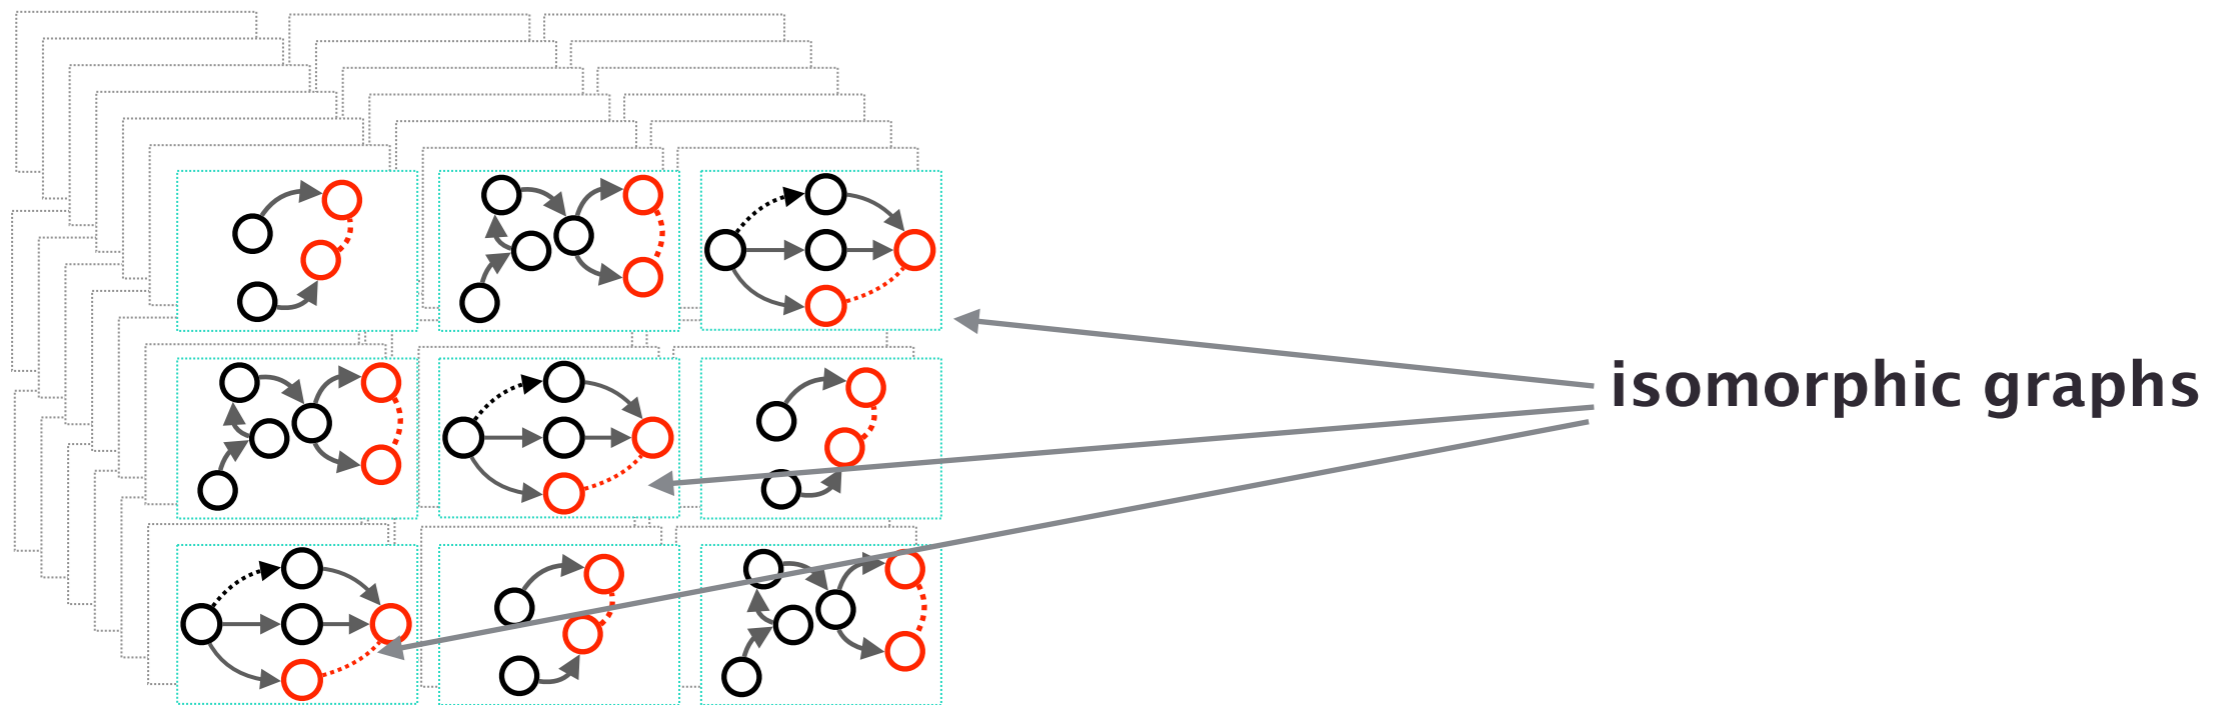
Isomorphic clustering

Domain specific features based clustering




1,000s of violations subgraph

Identical events sequences may trigger the same controller bug.




1,000s of violations subgraph

Group identical subgraphs into clusters




Group identical subgraphs into clusters

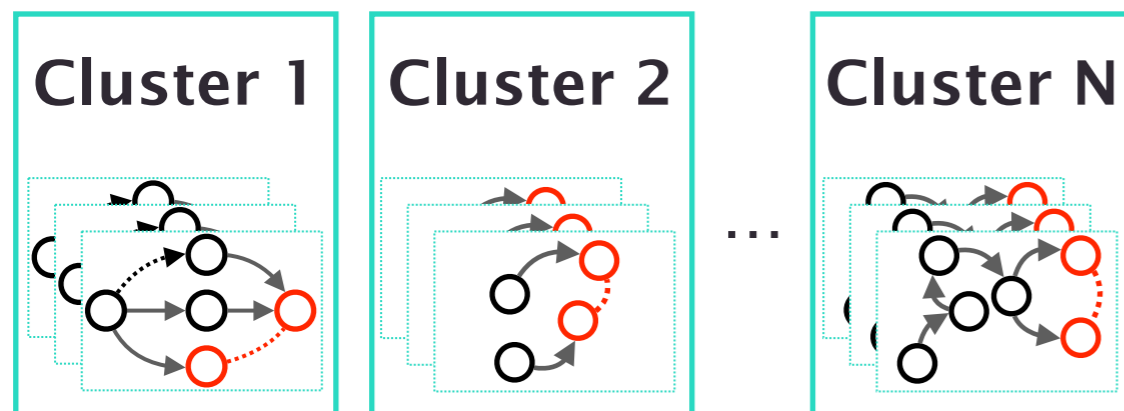


1,000s of violations subgraph

100s of isomorphic clusters




100s of isomorphic clusters




Not all bugs are triggered by *identical events*.

100s of isomorphic clusters



How can we **measure** the **relevance** of two violations reports to the a **source bug**?


100s of isomorphic clusters



100s of isomorphic clusters


How can we **measure** the **relevance** of two violations reports to the a **source bug**?

Domain specific features




7 domain specific features are used to compute a **distance matrix** between the different isomorphic clusters.

100s of isomorphic clusters

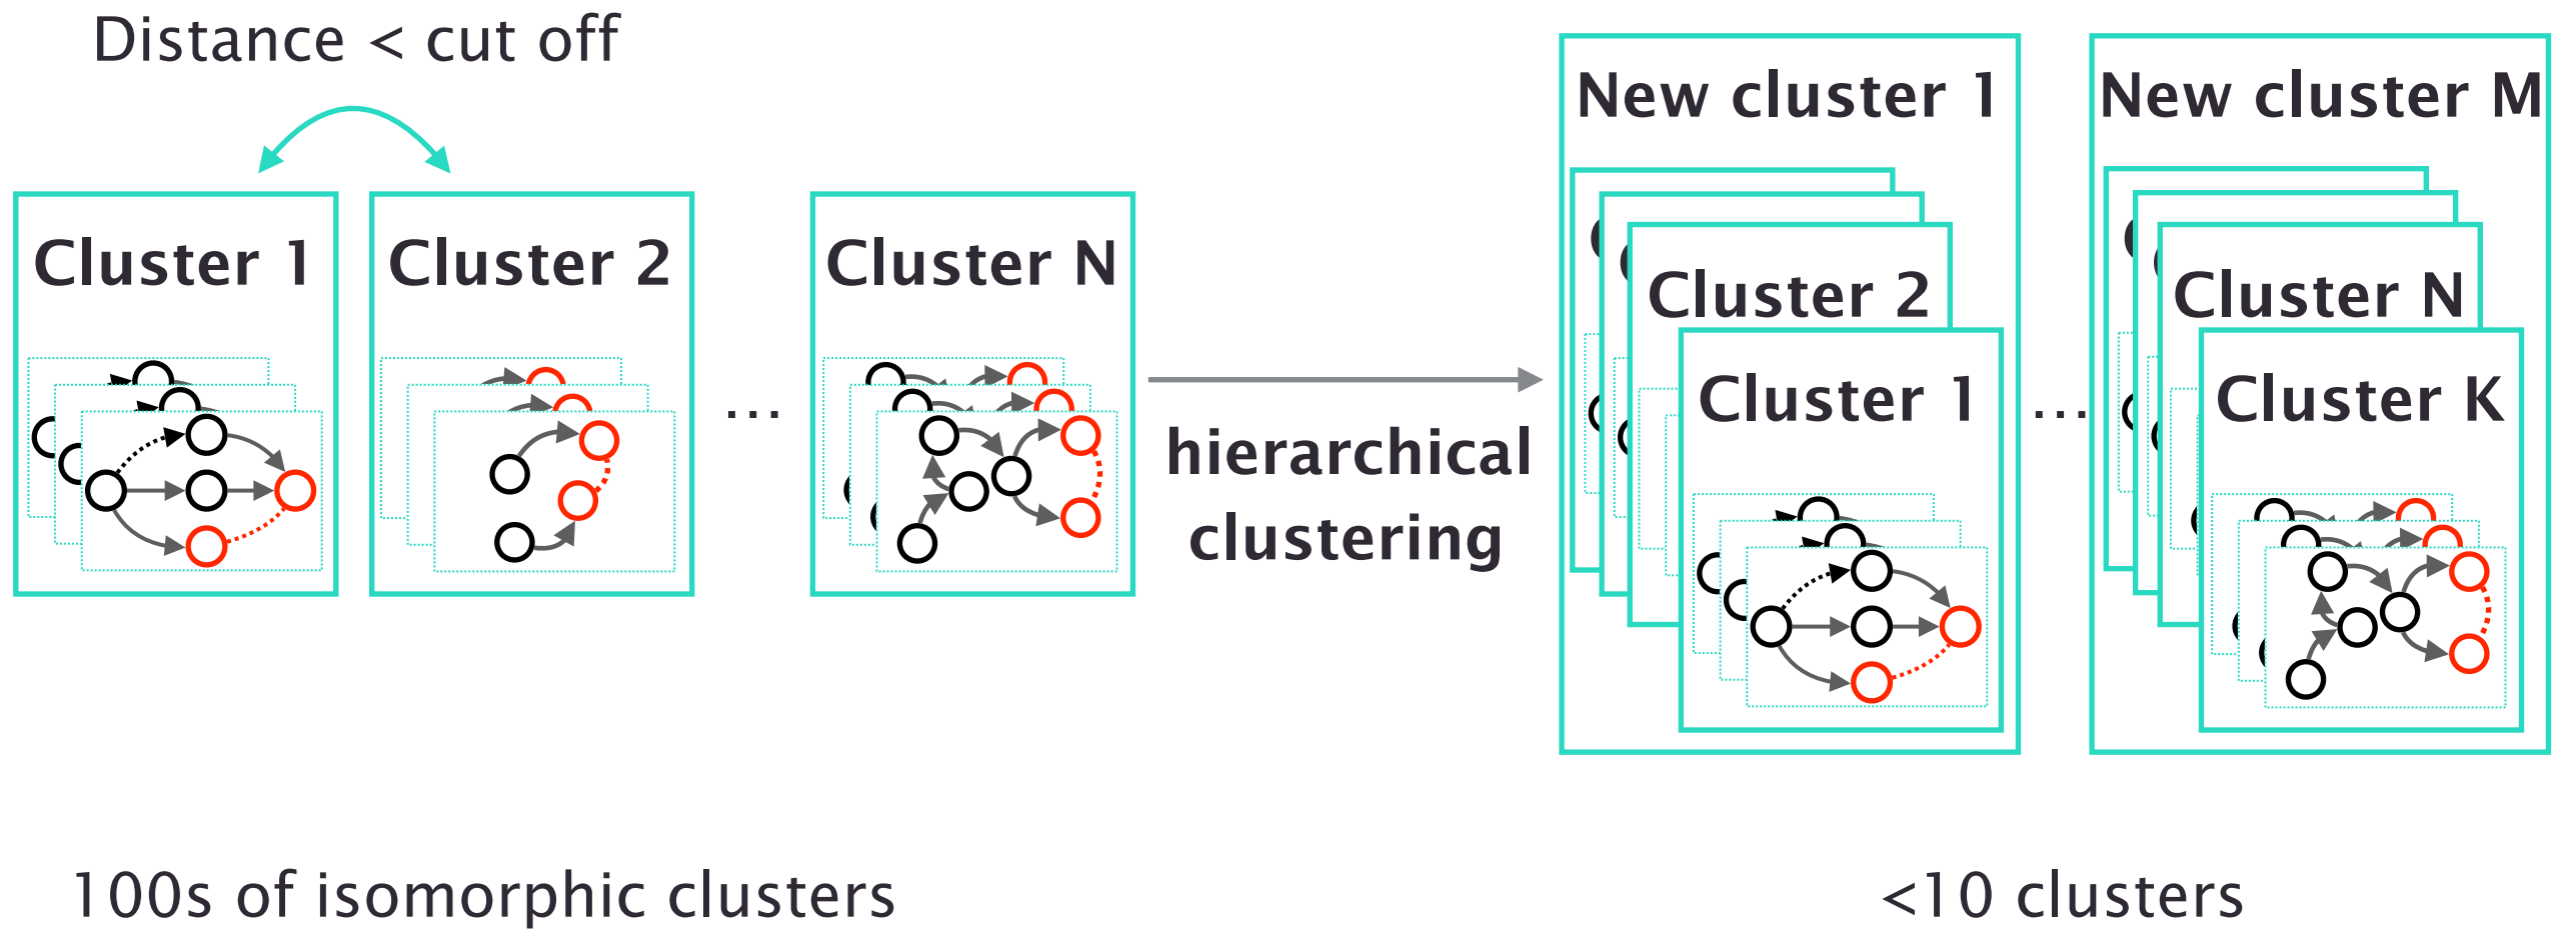


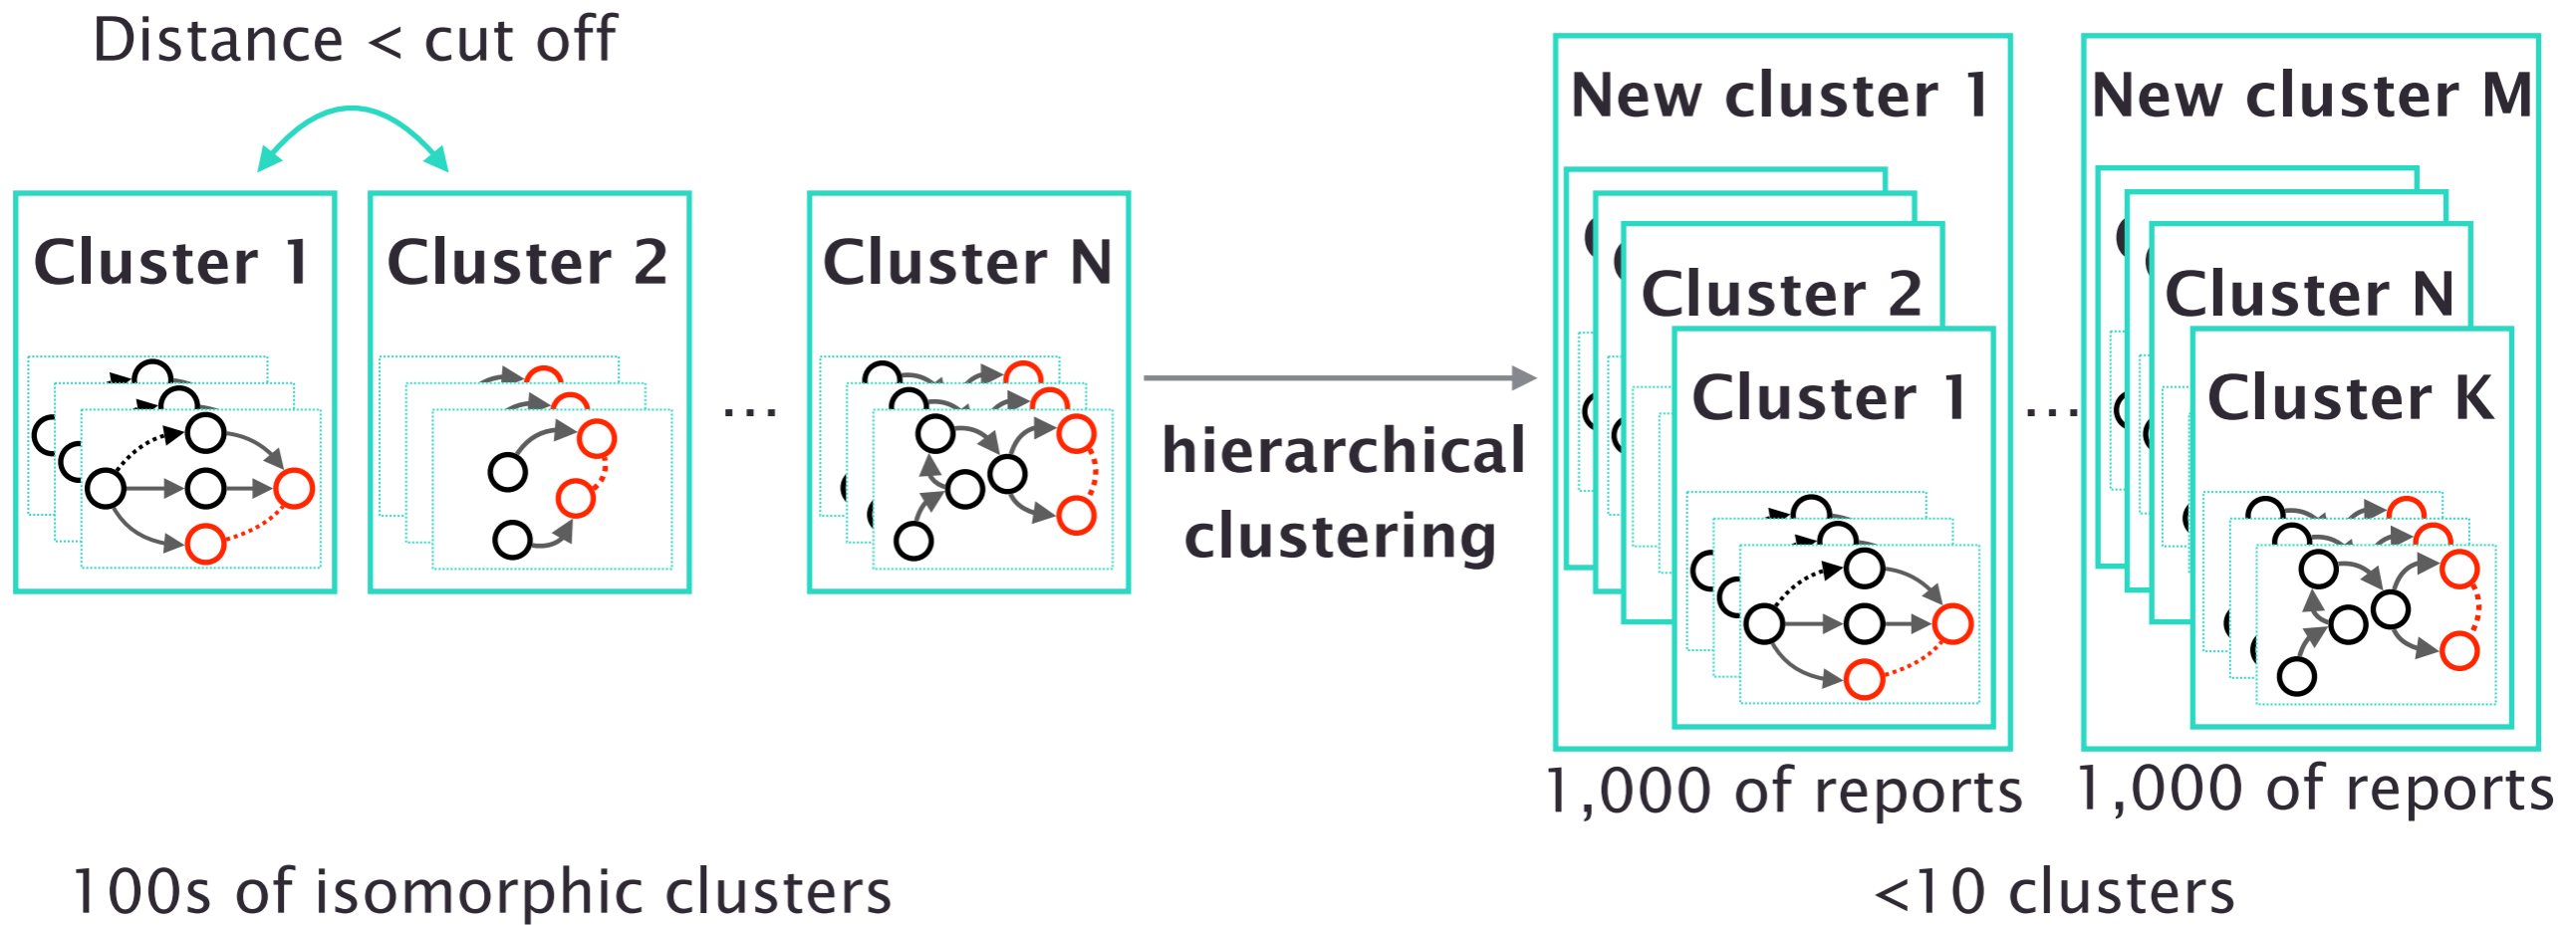
100s of isomorphic clusters

Hierarchical clustering algorithm is used to group **close** clusters together.




100s of isomorphic clusters

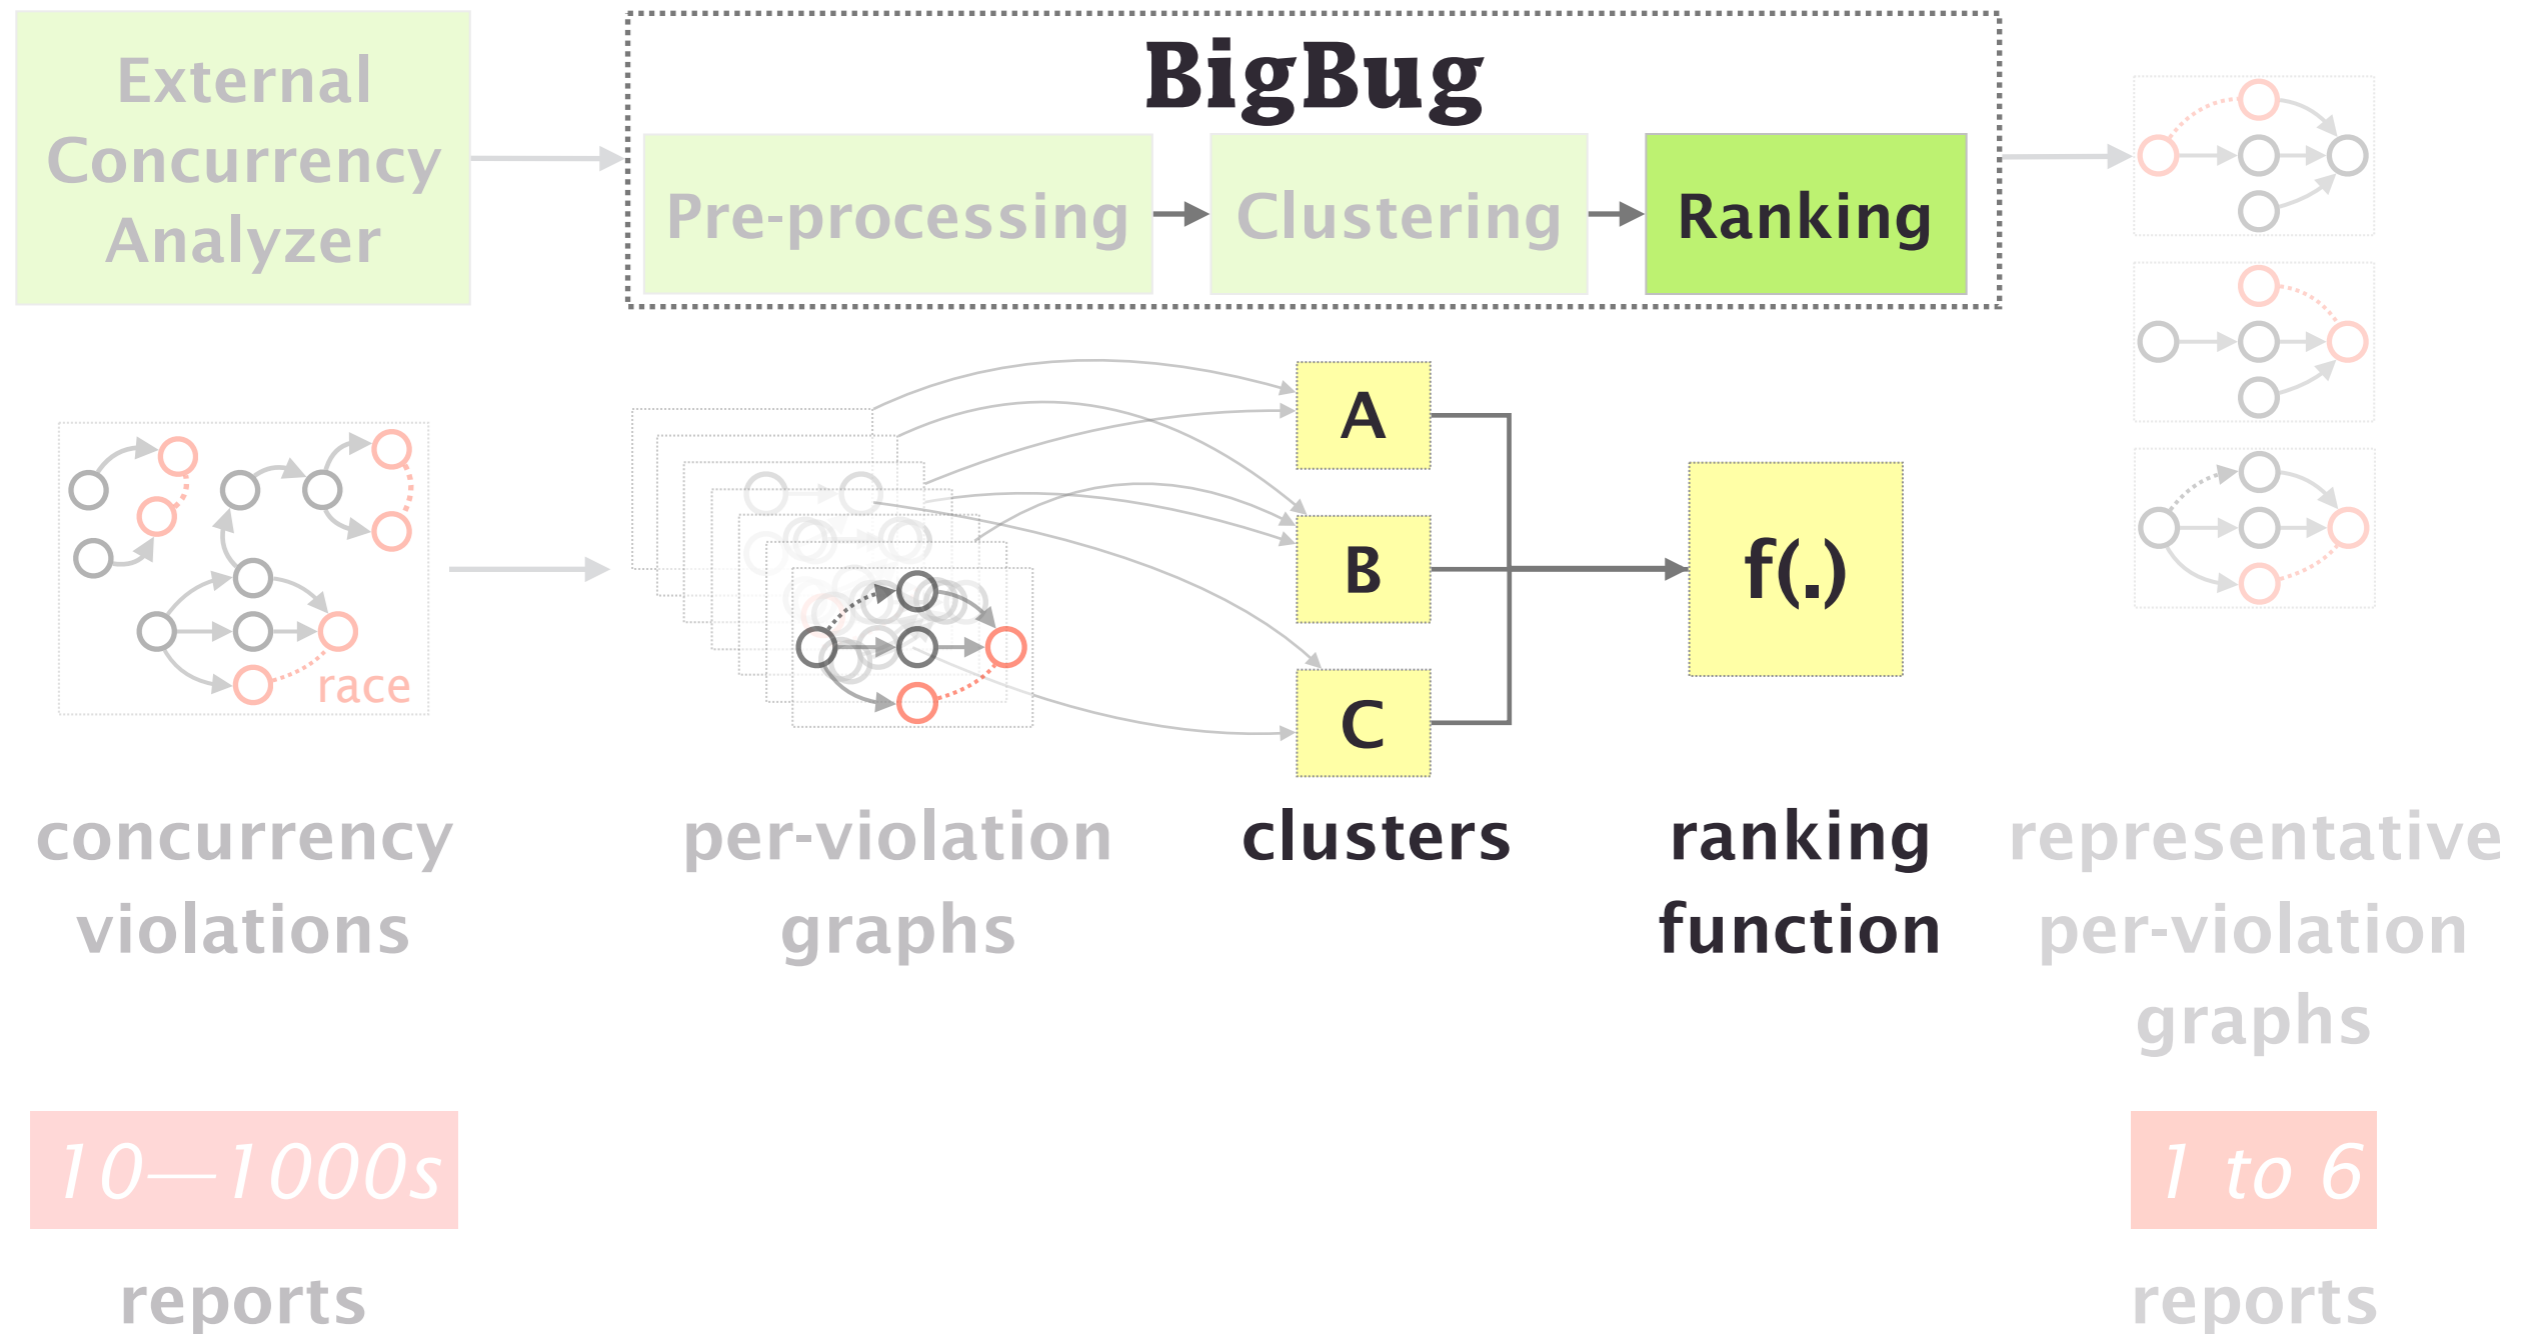





BigBug Pipeline

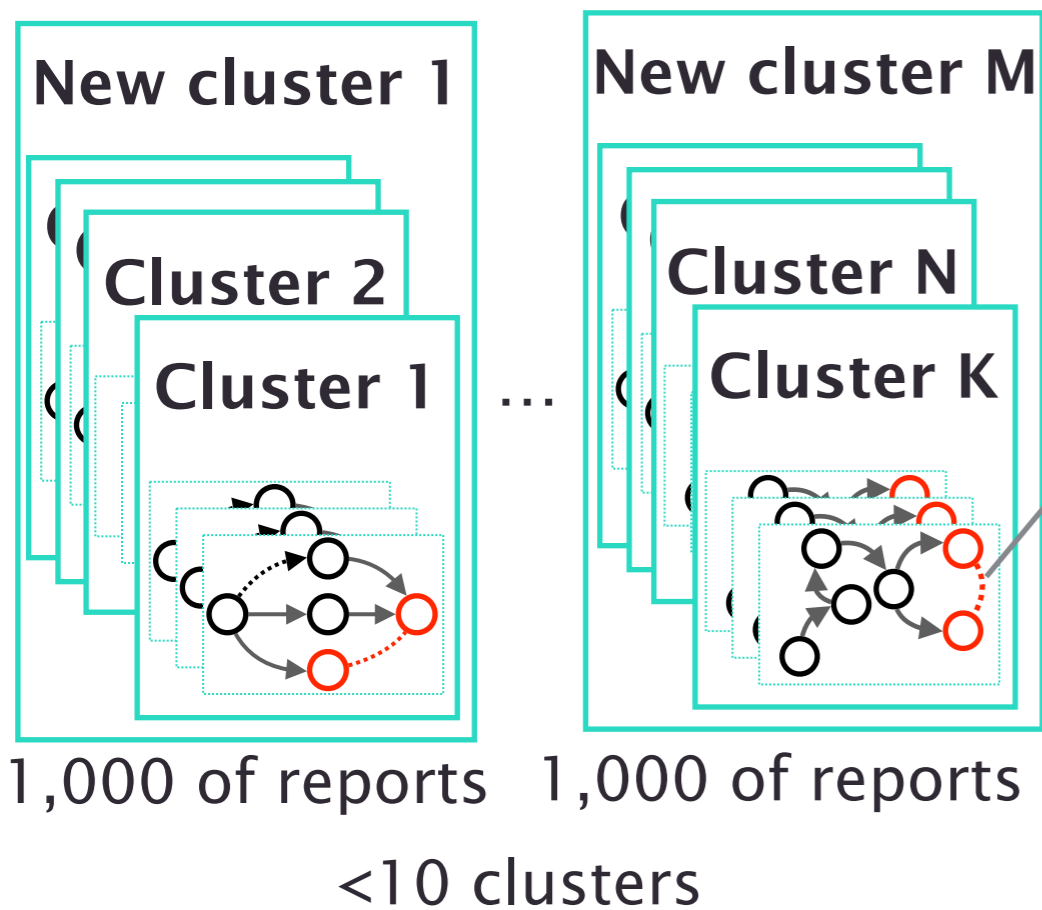


BigBug Pipeline







Select the **most representative violation**




1. Exhibits largest number of domain specific features represented in the cluster.



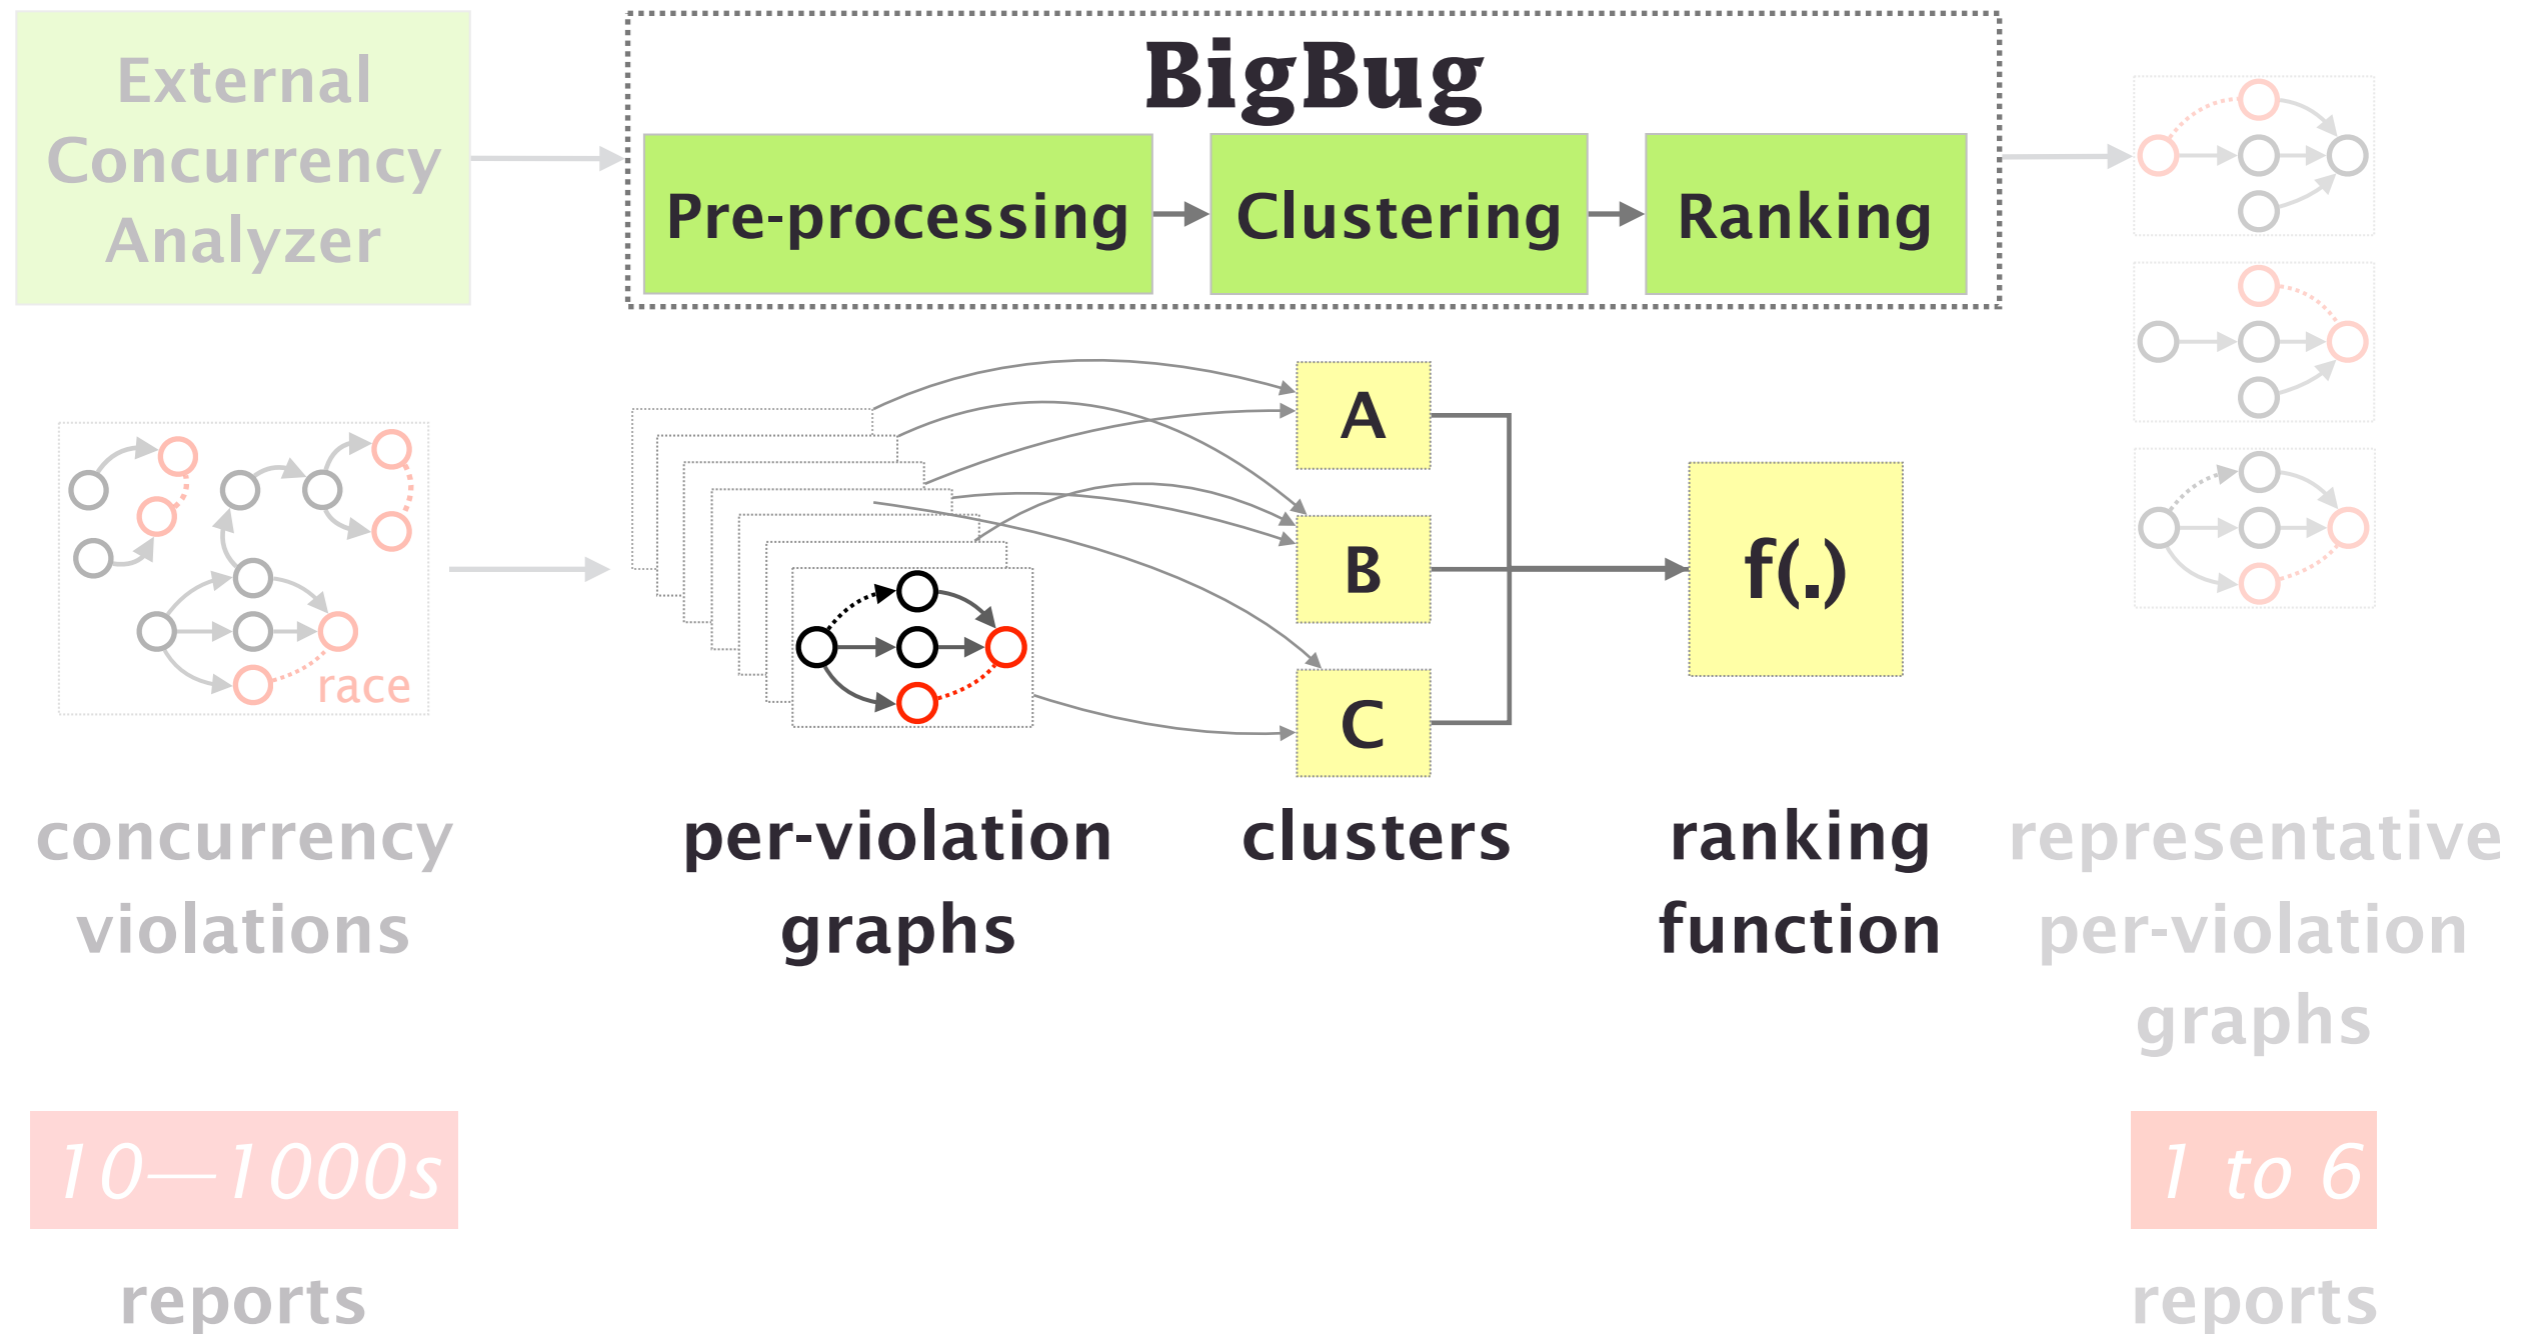
1. Exhibits largest number of domain specific features represented in the cluster.
2. Has the smallest size.



→ ranking



BigBug Pipeline




One switch star topology

Floodlight load balancer application

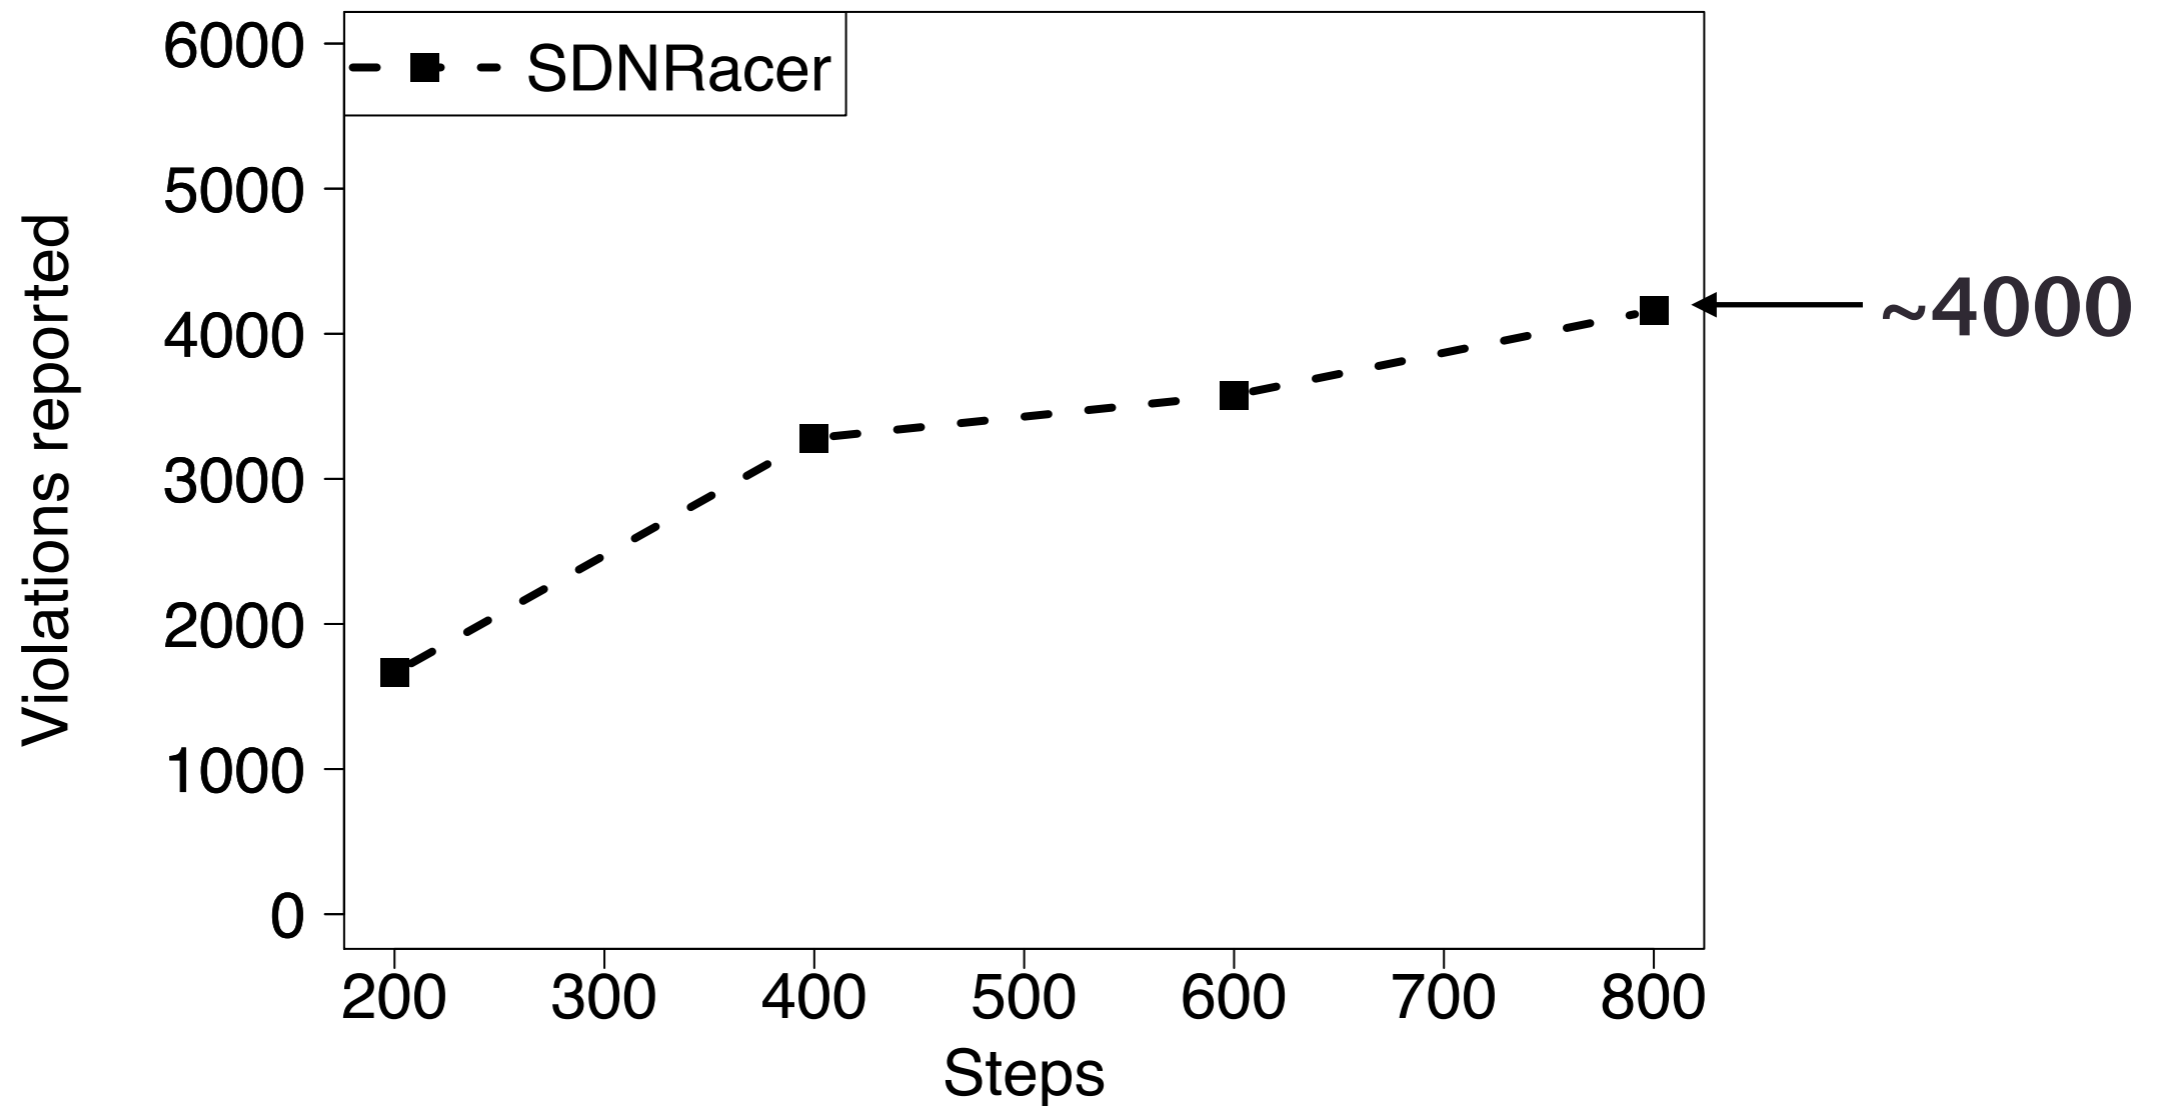
Various simulation lengths

Violations in Floodlight Loadbalancer




The median of 15 repetitions, on Star topology

Violations in Floodlight Loadbalancer



The median of 15 repetitions, on Star topology

Violations in Floodlight Loadbalancer



The median of 15 repetitions, on Star topology

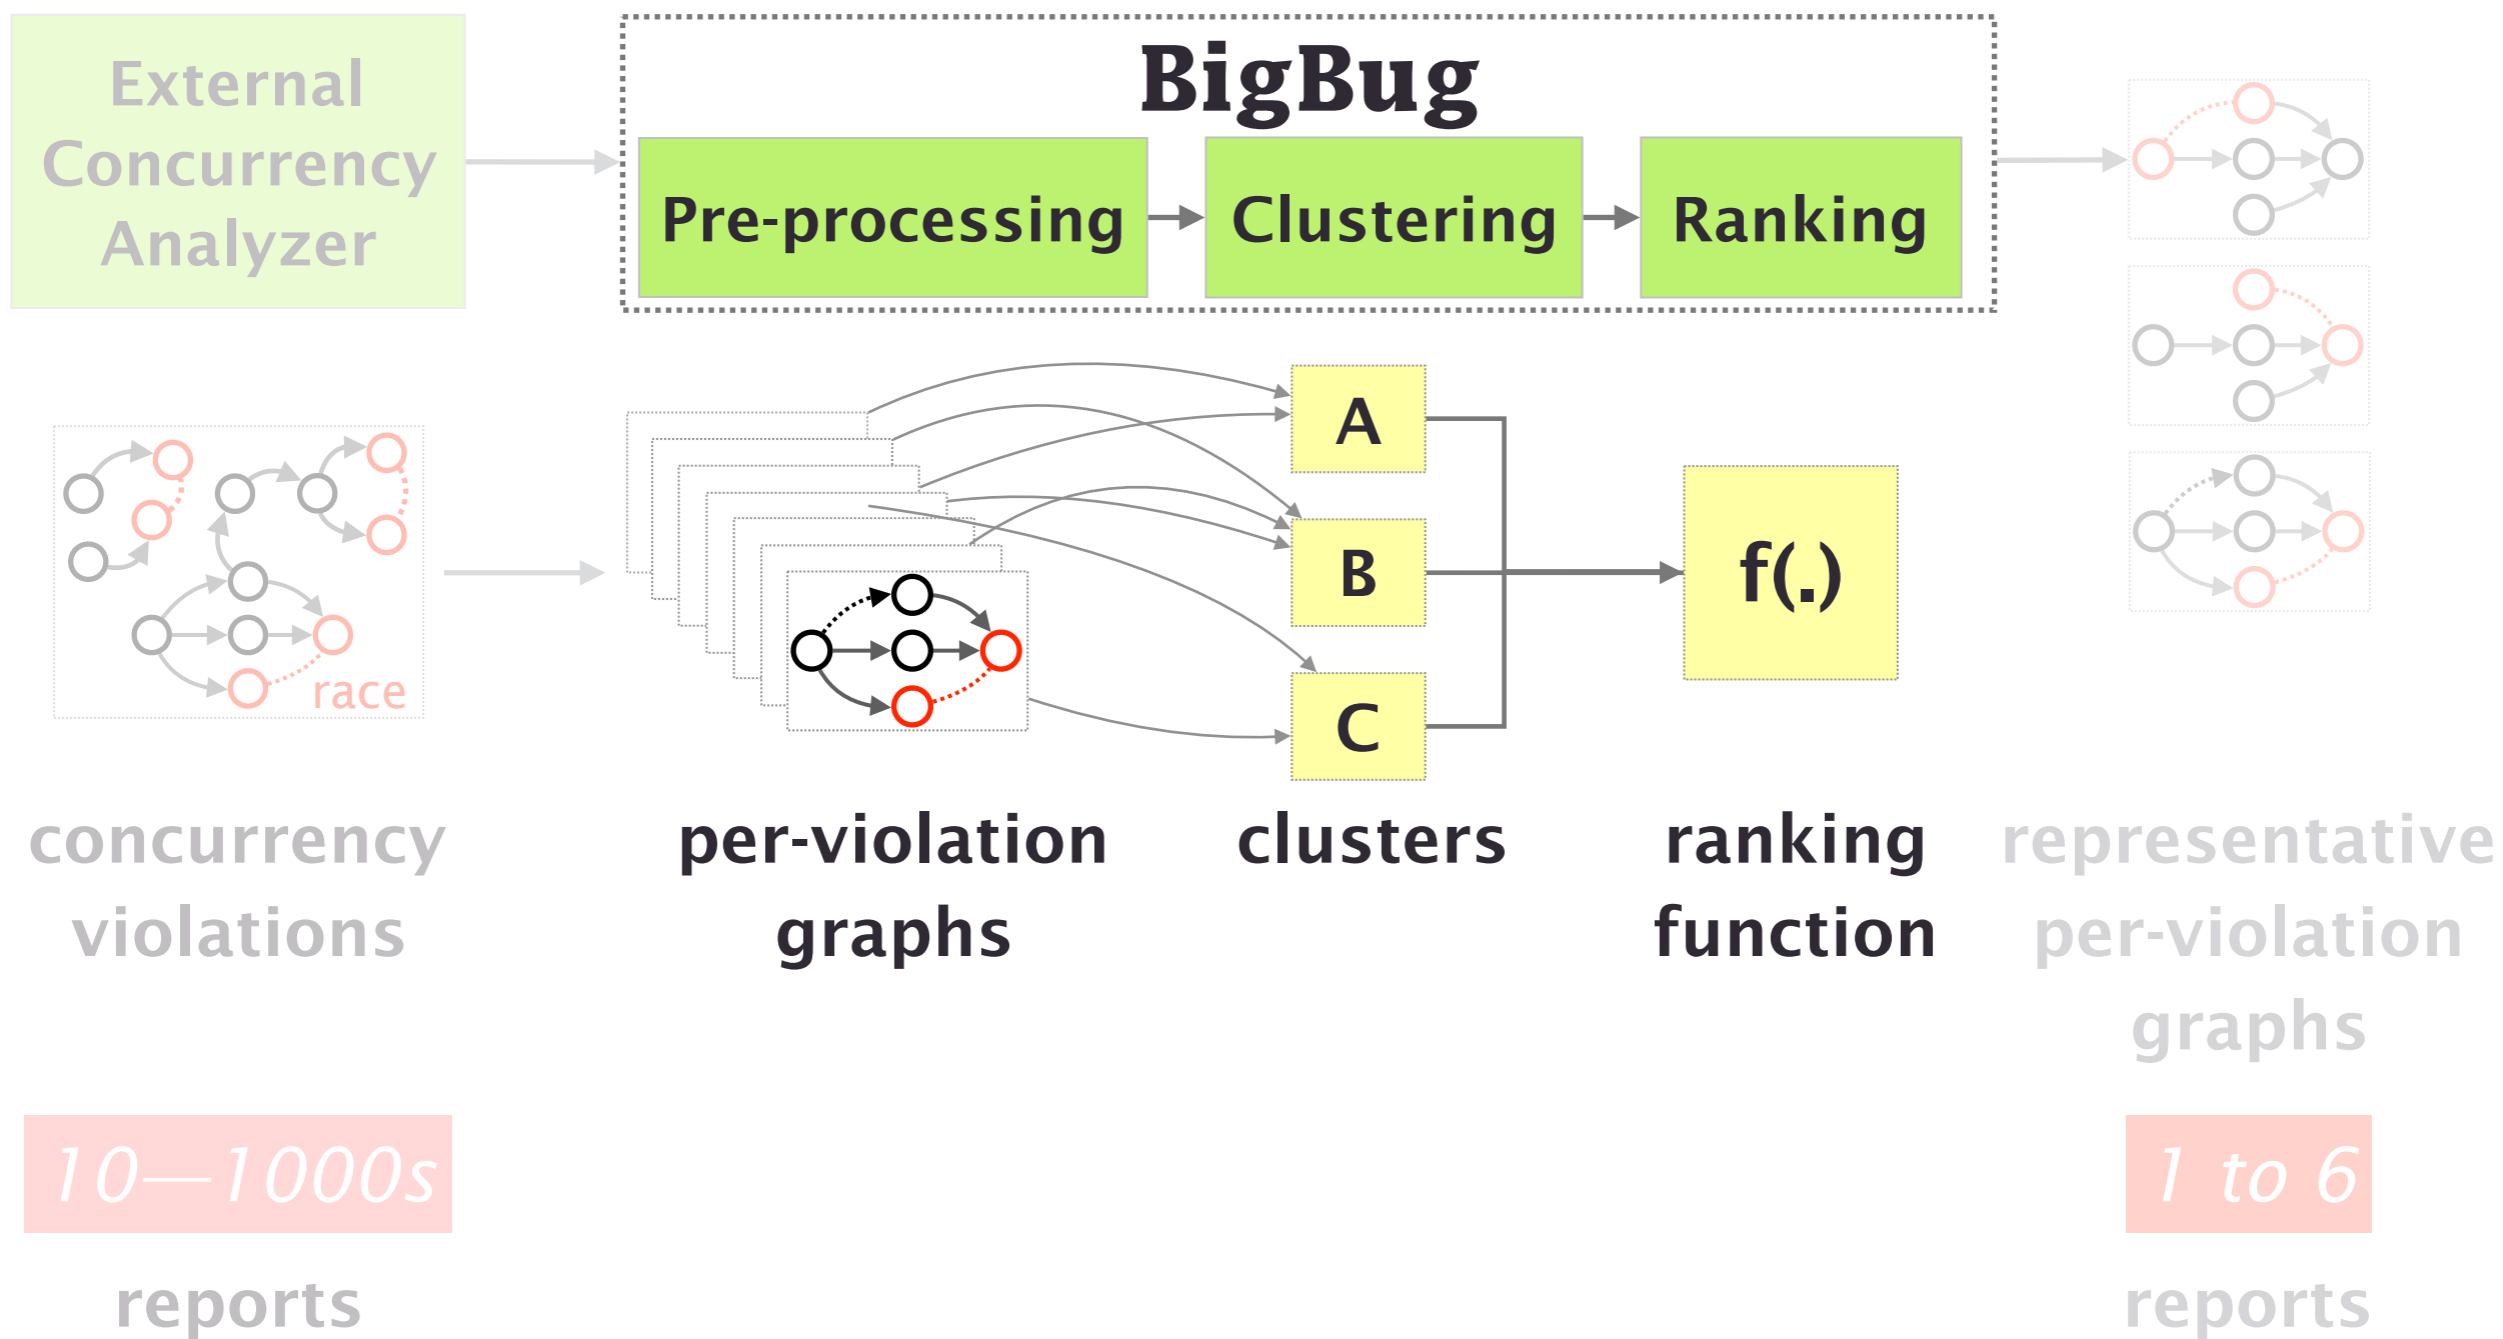
Running 6 applications

Binary Tree Topology

91 to 1910 concurrency violations reported

BigBug reports **less than 6** violations

across all our experiments



<http://sdnracer.ethz.ch>