

Simplifying Datacenter Network Debugging with PathDump



Praveen Tammana[†]




Rachit Agarwal[‡]




Myungjin Lee[†]

[†]University of Edinburgh, [‡]Cornell University




Network problems are inevitable

Loops



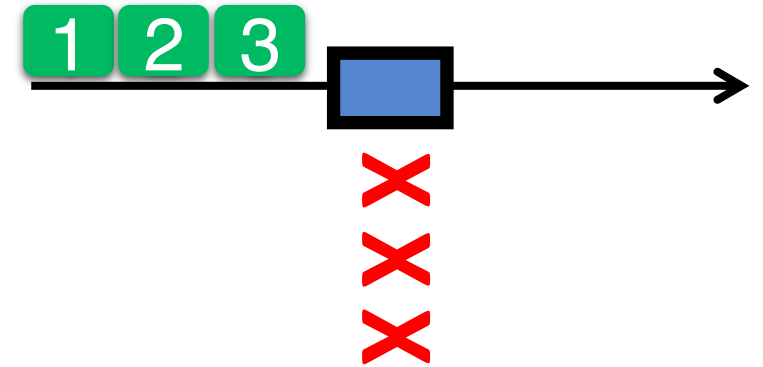
Failures, bugs

Silent random packet drops



Faulty interface

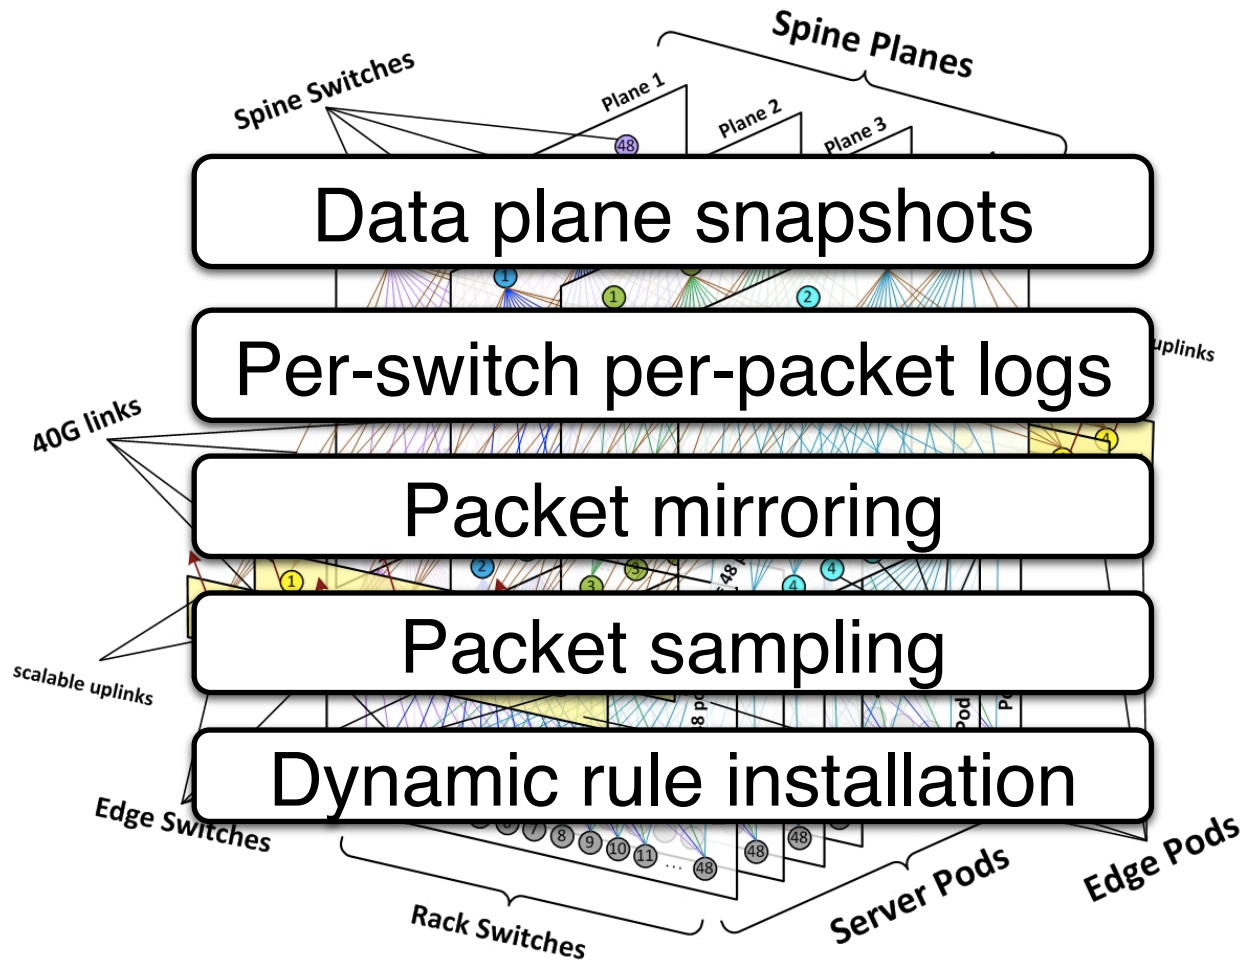
Black hole



Human errors

- Result: **Mismatch** between network behavior and operator intent
- Network debuggers
 - Existing designs: **in-network** techniques
 - Use programmability of network switches to capture debugging information

Complex networks and debuggers

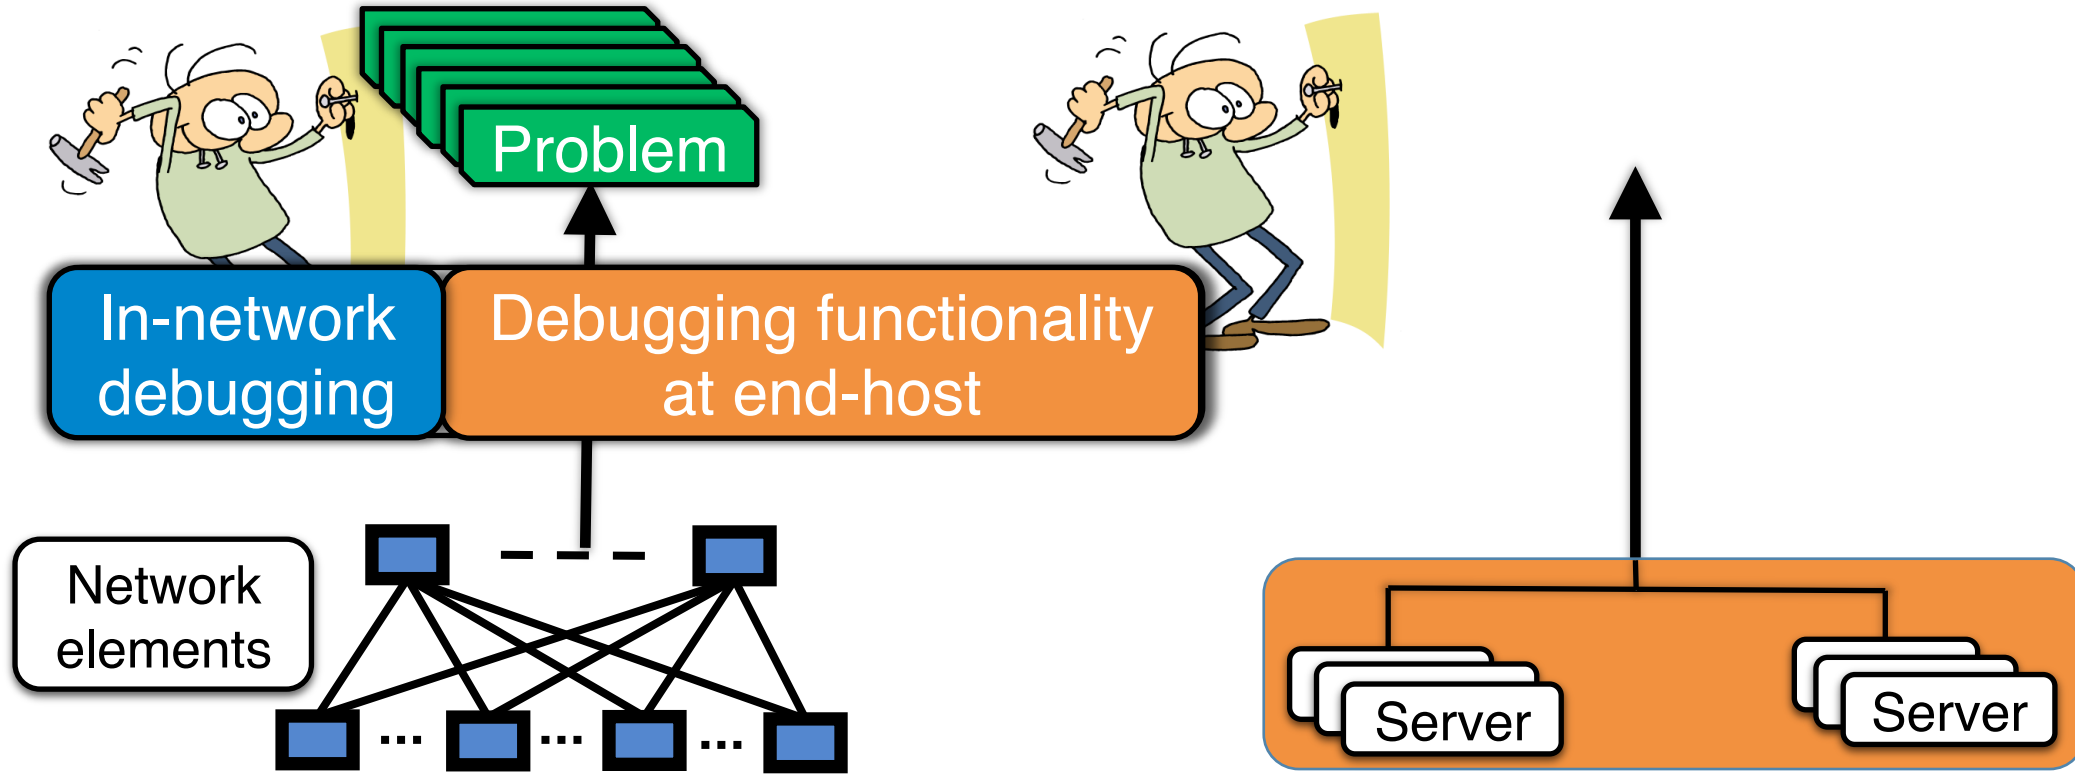


--source: TechRepublic.com

Complex networks

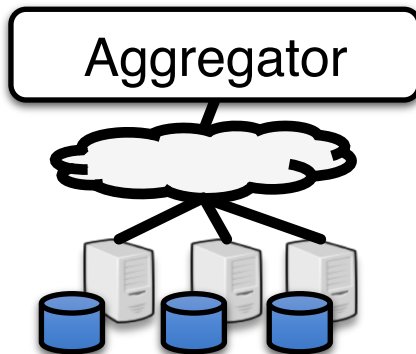


**Network debuggers
even more complex**

PathDump: (Simple) In-network + End-hosts



- Use end-hosts for most debugging problems
- In-network functionality for a small number of debugging problems

PathDump in a nutshell



- Before forwarding a packet, checks a condition
- If met, embeds its ID into packet header
- No data plane snapshots
- Captures each and every packet header
- No per-switch per-packet logs
- Stores and updates flow-level statistics
- No packet sampling
- Exposes API for debugging purposes
- No packet mirroring
- No dynamic rule installation
- Enables slicing-and-dicing of statistics across flows (potentially stored at various end-hosts)


PathDump architecture

1. Switch embeds unique ID (e.g., link ID)

- Packet header space limitation

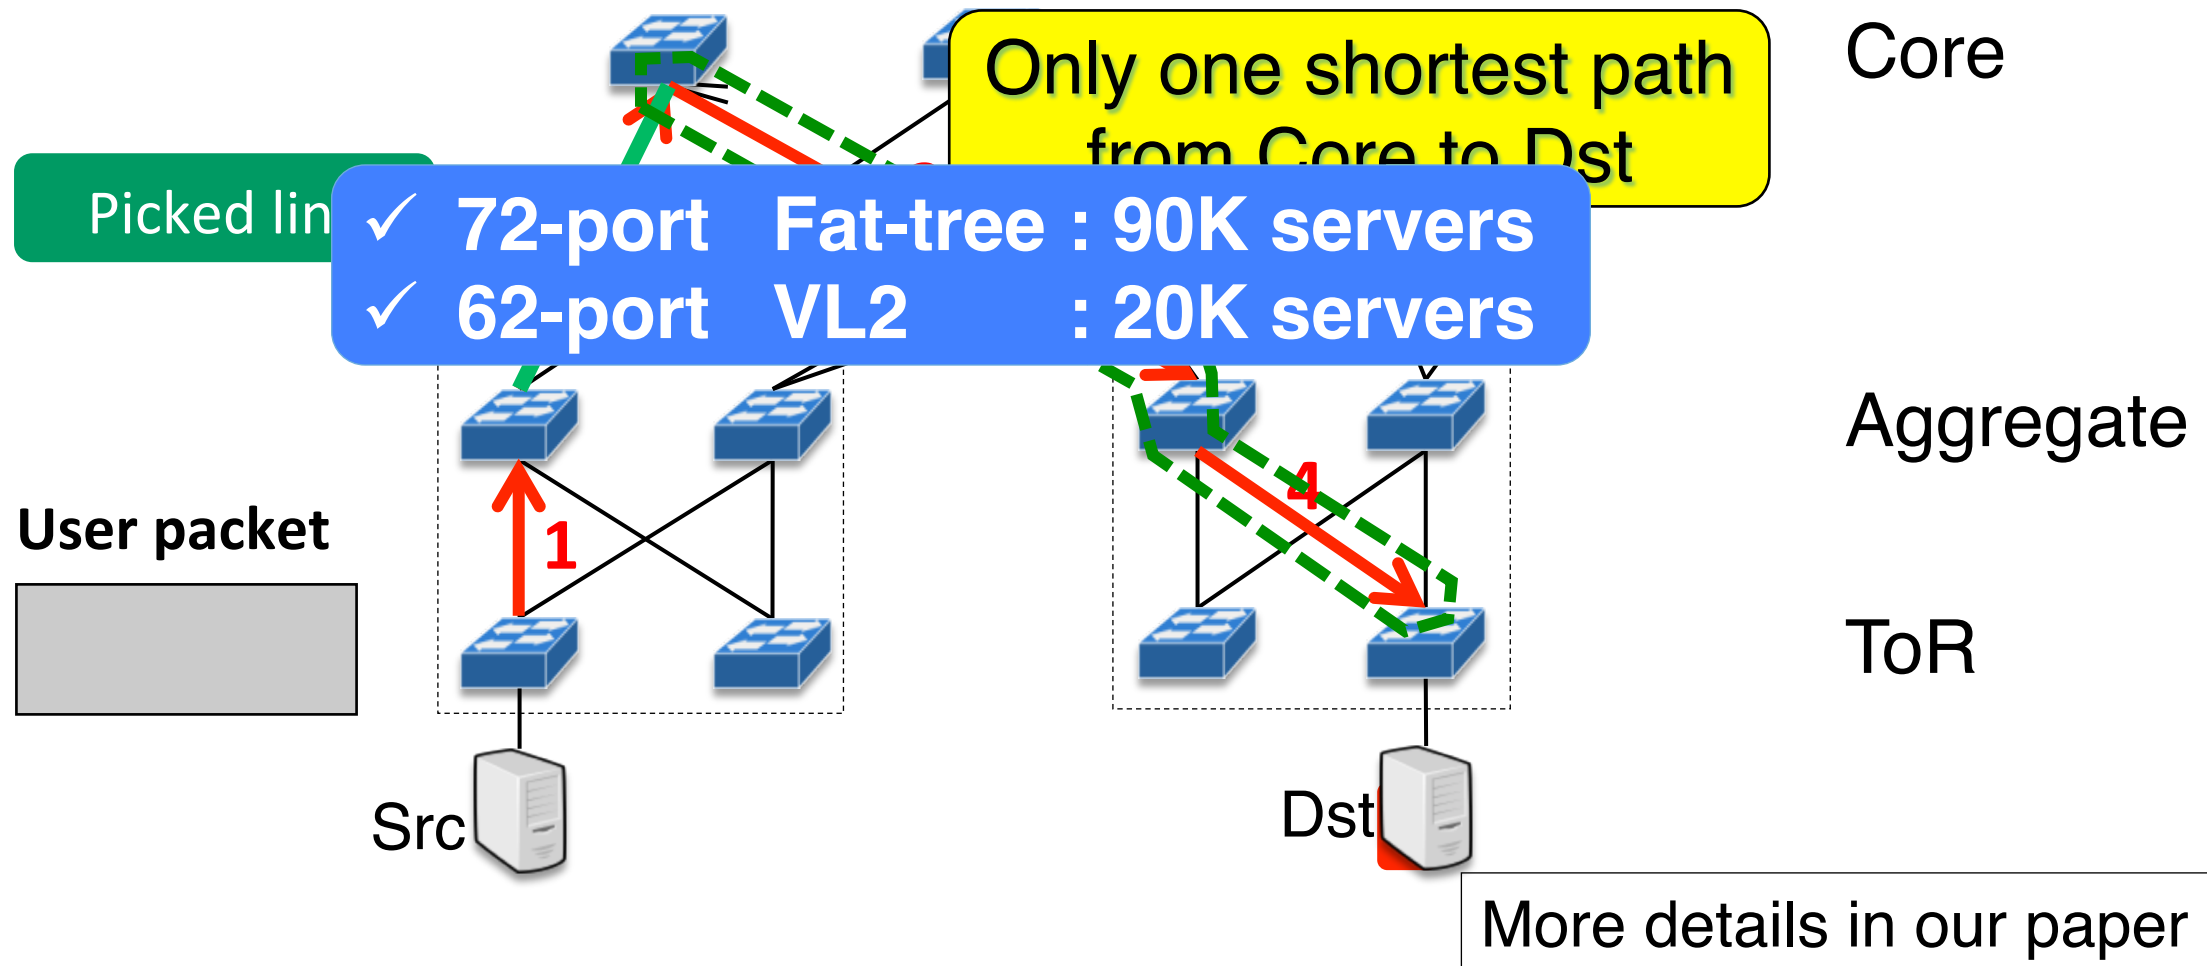
Em

- Cherrypick [SOSR'15] for current deployments




PathDump architecture

1. Switch embeds unique ID (e.g., link ID)



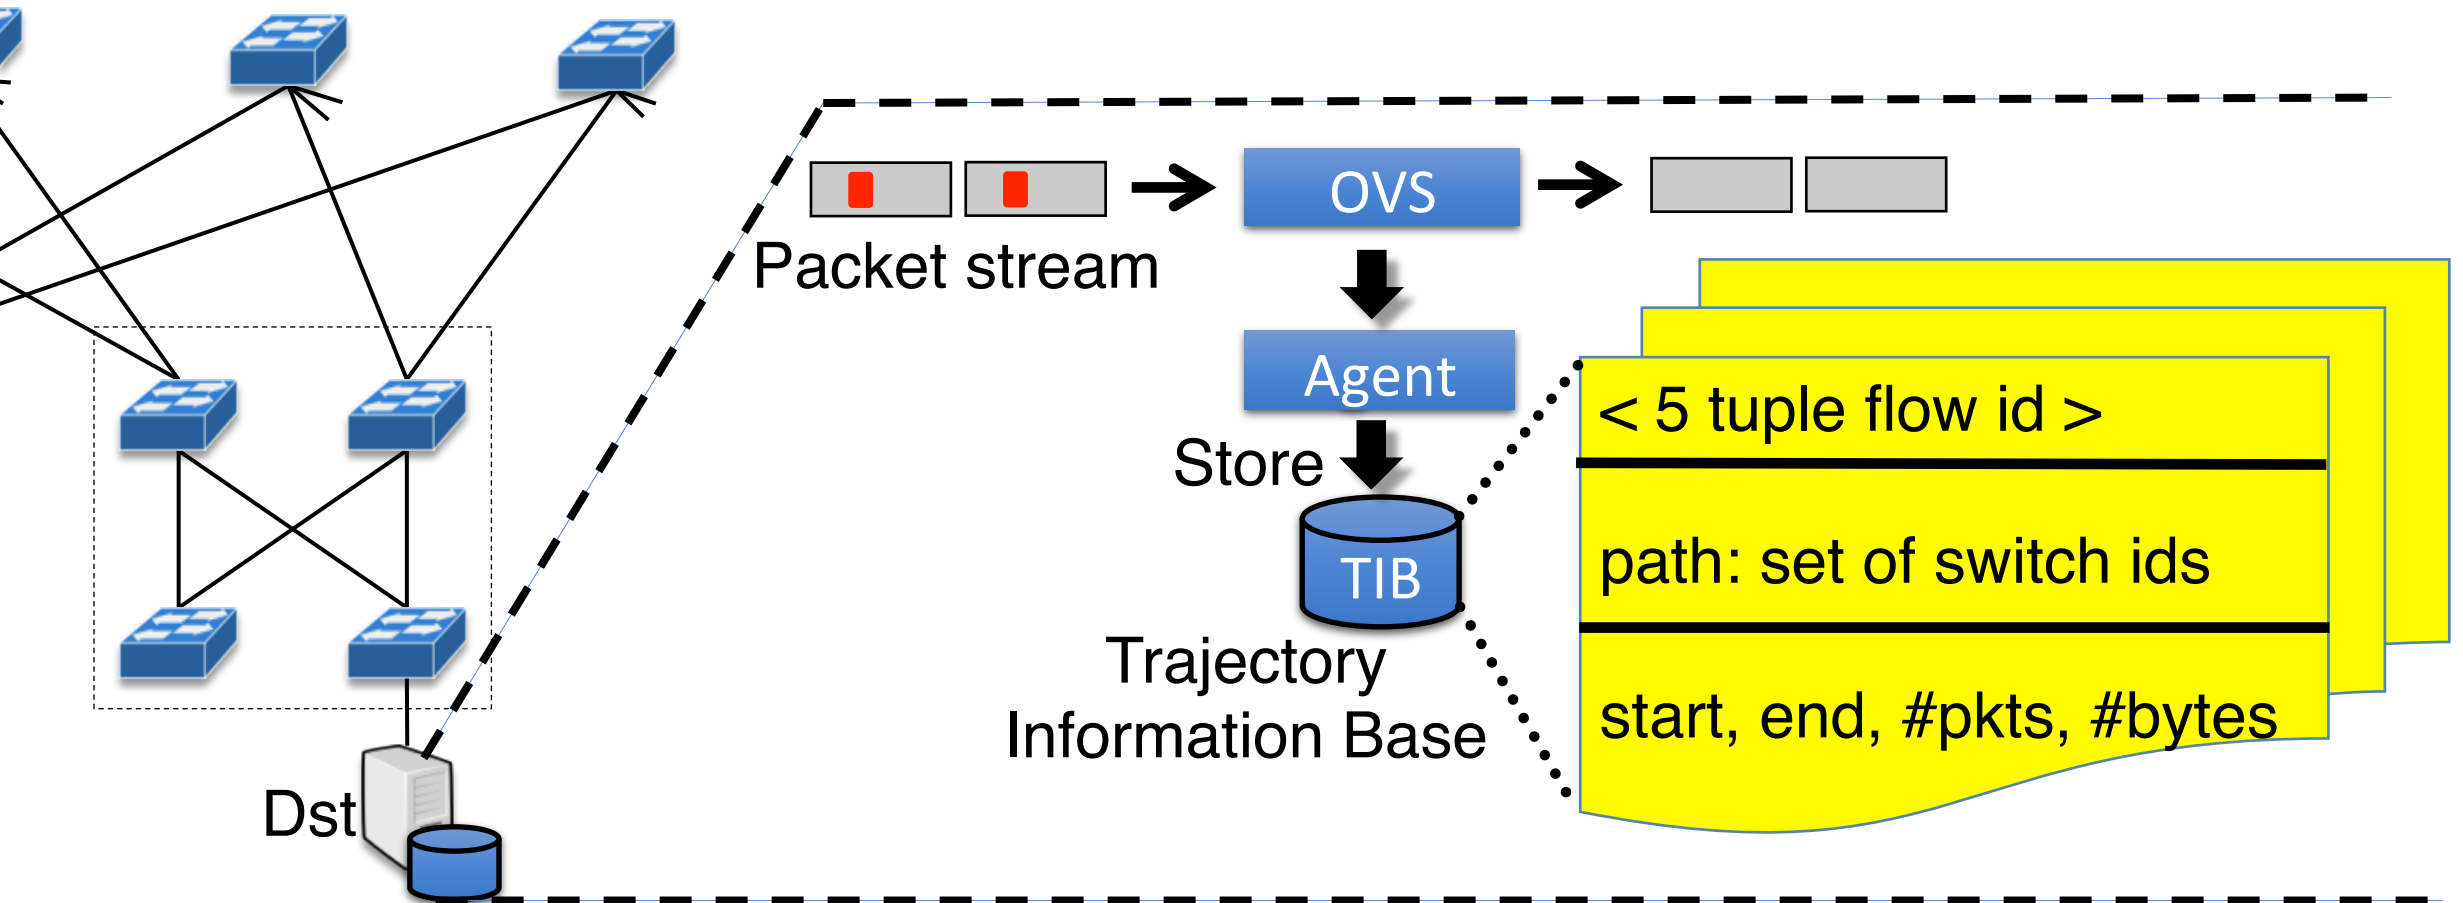
PathDump architecture

2. End-host captures packet path and updates flow-level statistics



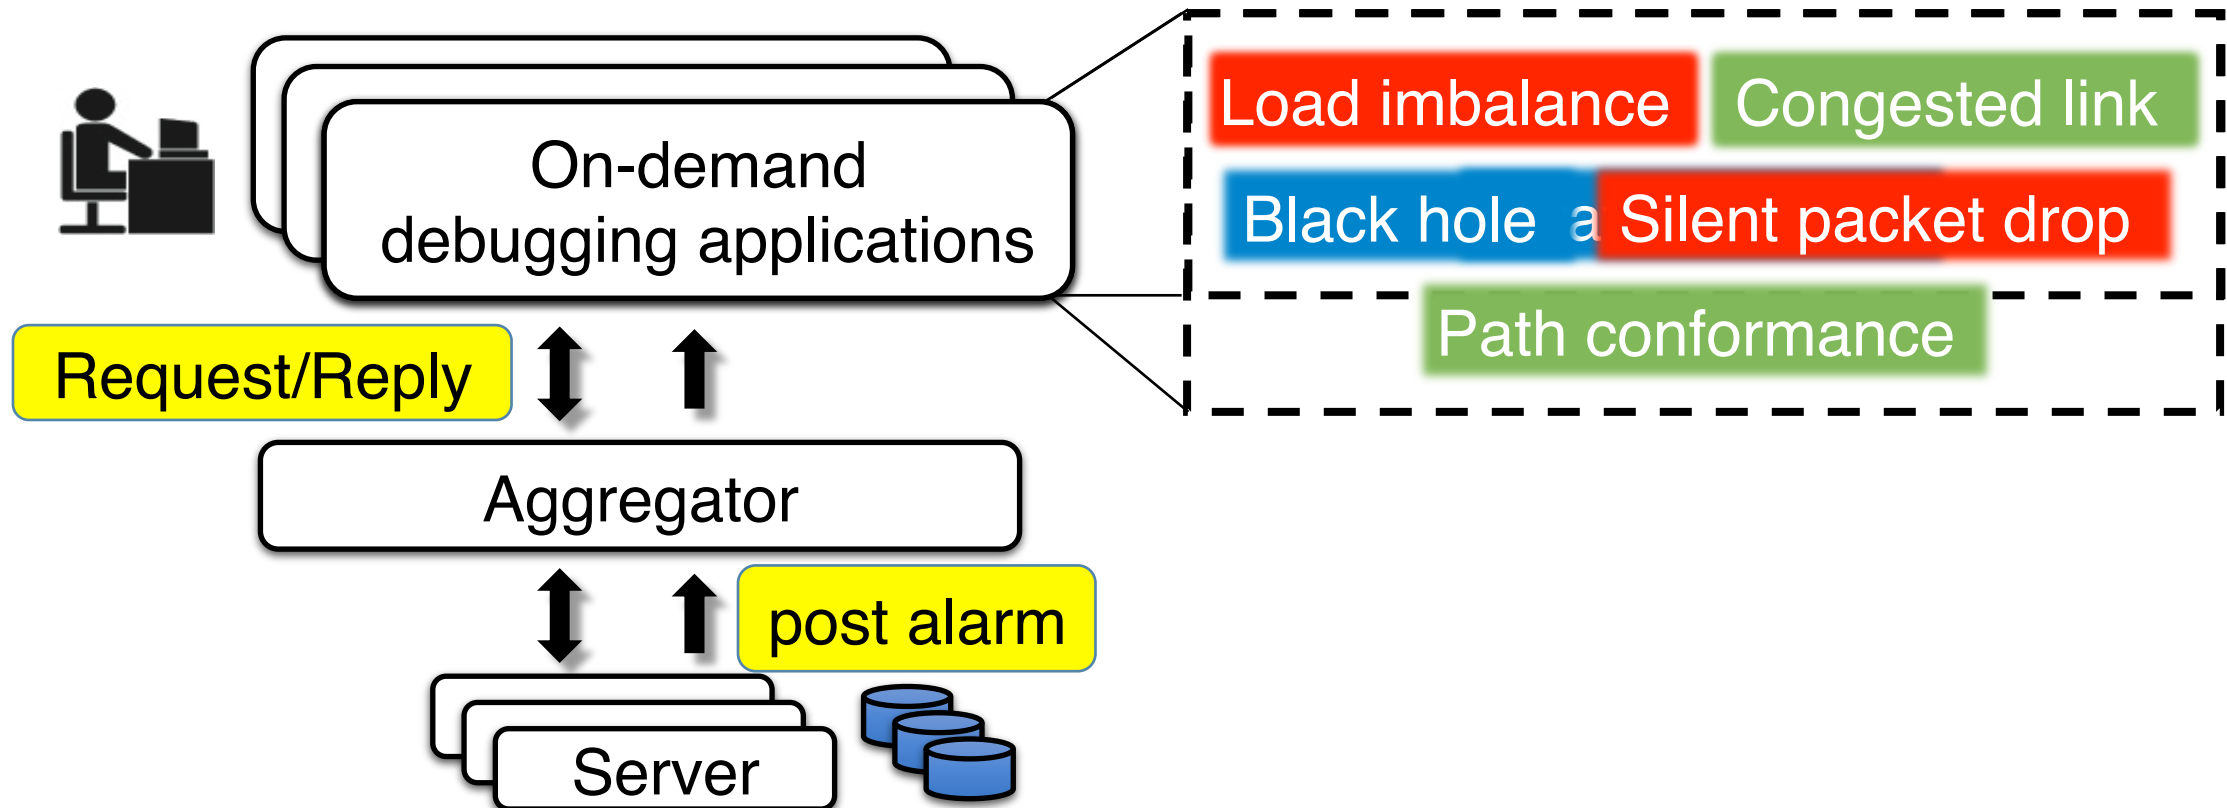
PathDump architecture

2. End-host captures packet path and updates flow-level statistics

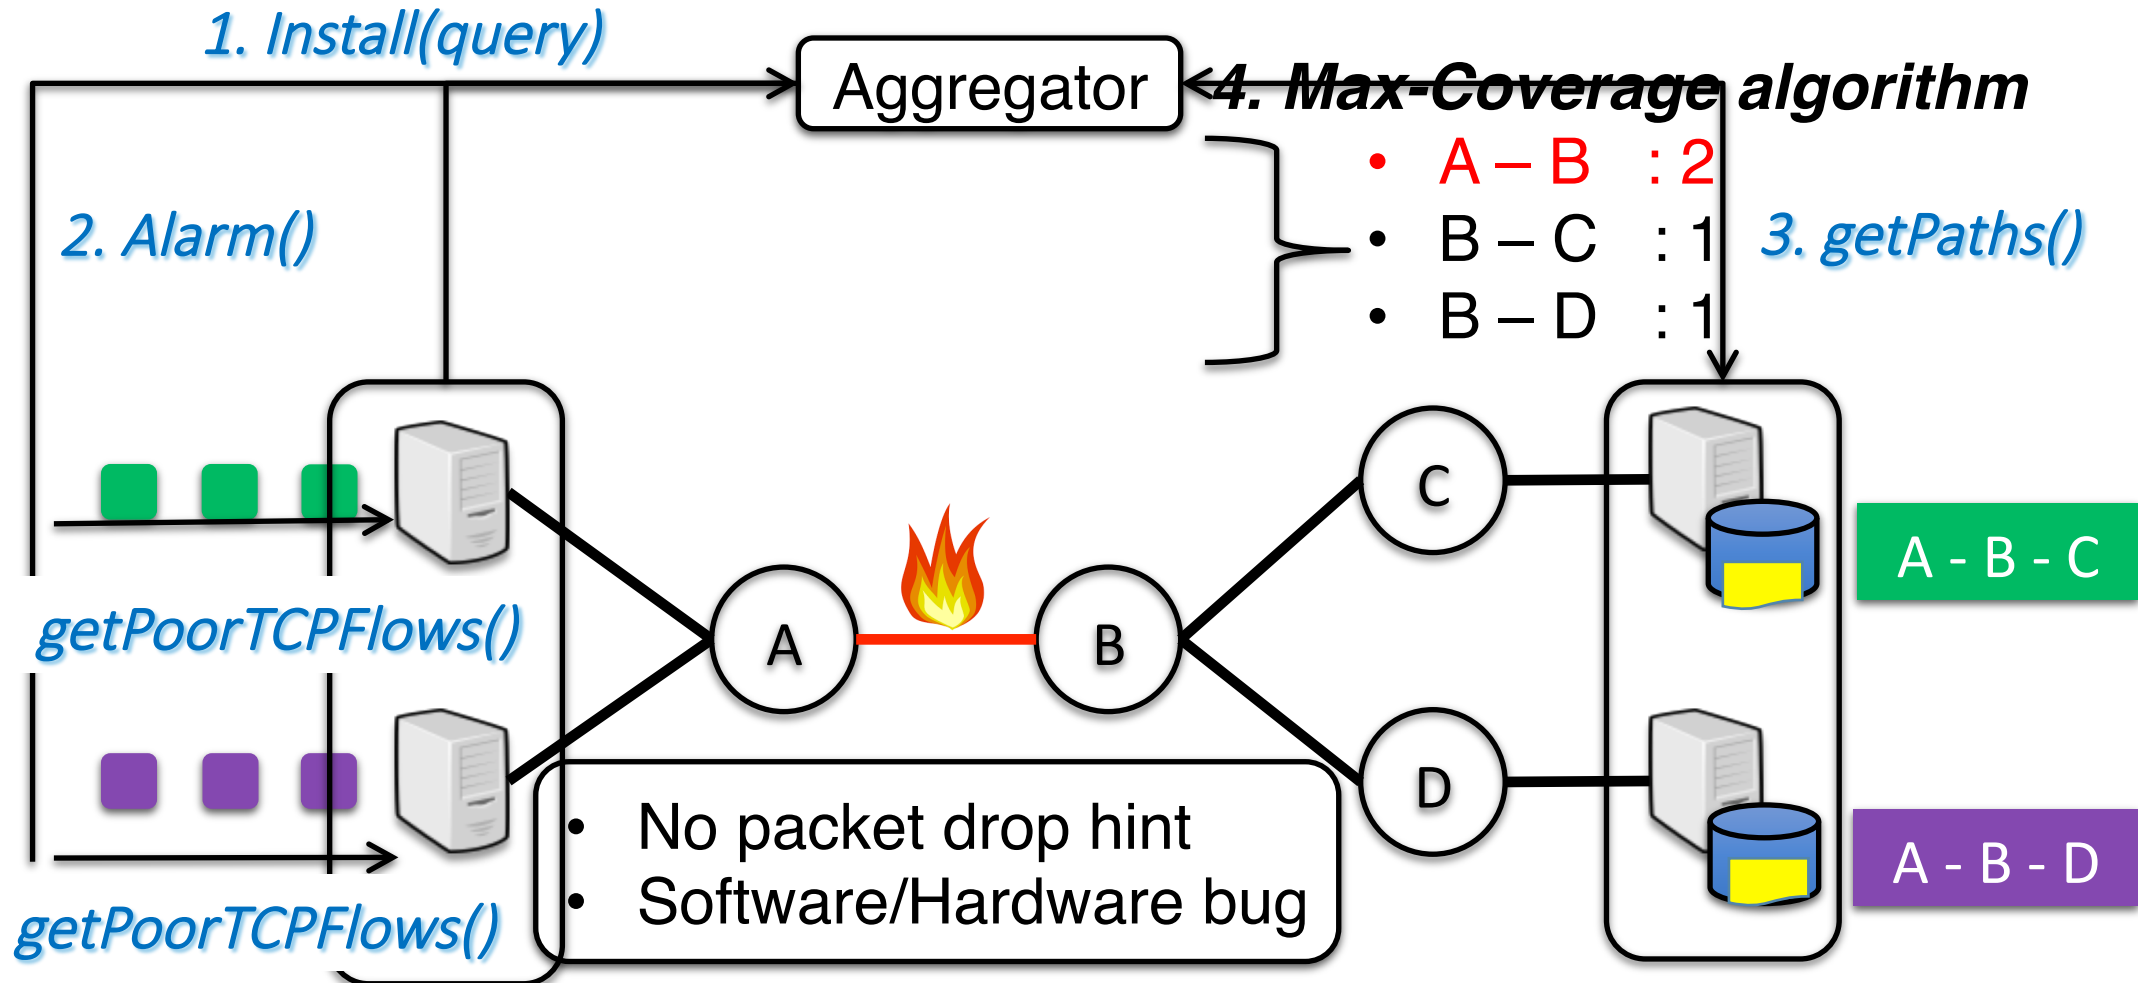


PathDump architecture

3. Aggregator runs debugging applications On-demand vs. Event-driven



Example: Silent random packet drop diagnosis



Other debugging applications

- Load imbalance diagnosis
- Real-time routing loop detection
- Blackhole diagnosis
- TCP performance anomaly diagnosis
 - TCP incast and outcast
- Traffic measurement
 - Traffic matrix, heavy-hitter detection, etc.

More details in our paper

<https://github.com/PathDump>