

UC San Diego

Technical Reports

Title

Scalable Measurement: Finding some Elephants in a Swarm of Ants

Permalink

<https://escholarship.org/uc/item/99t93038>

Authors

Estan, Cristian
Varghese, George

Publication Date

2001-02-12

Peer reviewed

Scalable Measurement: Finding some Elephants in a Swarm of Ants

Cristian Estan and George Varghese

February 12, 2001

Abstract

While the number of concurrent flows passing through backbone routers is large (more than 250,000), measurements show a small number of flows (say 10%) represent a large fraction of the traffic. Thus a desirable goal for a measurement algorithm at a backbone router is to identify (say) the top 10% of the flows using not much more than 10% of the memory needed to keep track of all flows. A similar goal would be to determine any flows that use more than (say) 1% of the link bandwidth (to catch hot spots) using only a small amount of high speed memory. Keeping state for each flow (to tell whether a flow is “large”) is ruled out. Our paper describes an algorithm for identifying large flows using small memory and small per packet processing bounds. Our algorithm scales well even with increases in bandwidth and number of flows.