

XIA Performance

Expressive \neq Expensive

Srini Seshan and Hui Zhang

Peter Steenkiste, Aditya Akella, Dave Andersen,
John Byers, David Eckhardt, Sara Kiesler,
Jon Peha, Adrian Perrig, Marvin Sirbu,

San Diego FIA PI meeting

CarnegieMellon

BOSTON
UNIVERSITY



XIA's Flat Addressing

Current Internet

IP address

128.2.10.162

XIA

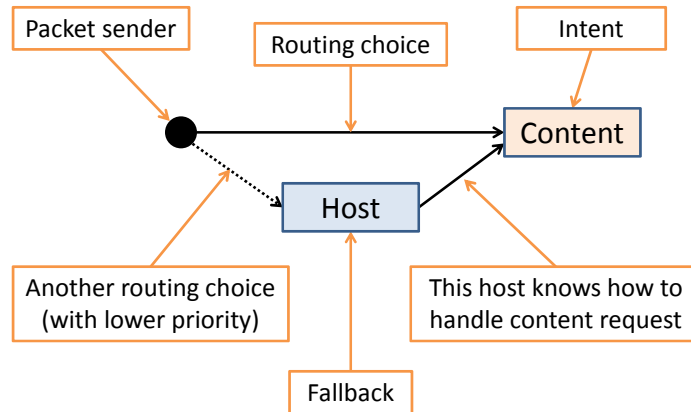
Principal type	Type-specific identifier
Host	0xF63C7A4...
Service	0x8A37037...
Content	0x47BF217...
Future	...

Hash of host's public key

Hash of service's public key

Hash of content

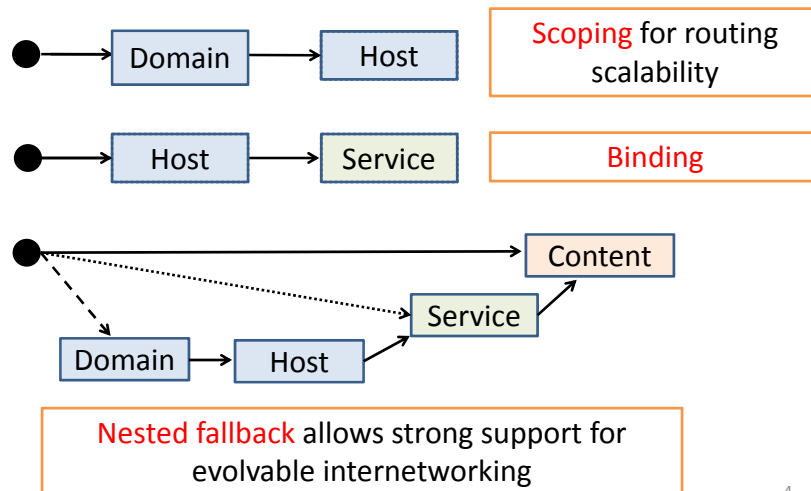
XIA's DAG-Based Addressing



A node can have **multiple outgoing edges**.
Outgoing edges have **priority** among them.

3


DAG Incorporates Key Networking Features



Nested fallback allows strong support for evolvable internetworking

4

Possible Performance Optimization “Knobs”



- Many choices: DAG, XID type, SID/CID routing, Scion vs NID, path selection, services, ..
- Examples: fault management, optimizing video distribution

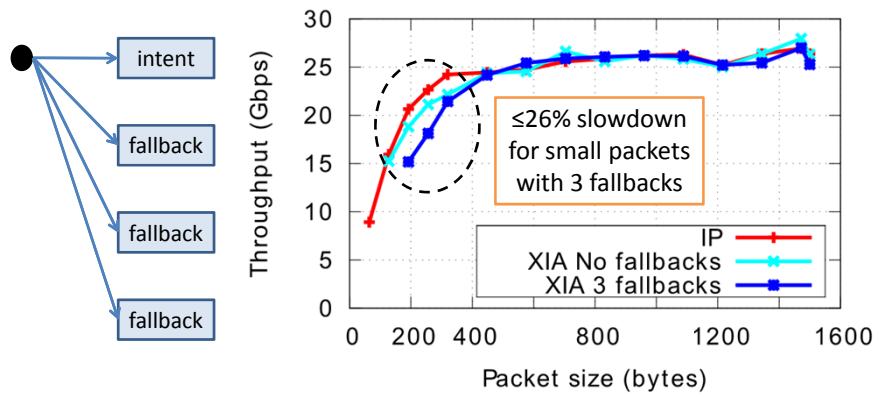
5

Outline

- XIA Performance Challenges/Opportunities
- **Packet Processing Performance (Data Plane)**
 - Processing DAGs
 - Large flat lookup tables
 - Congestion control
- Network-Wide Performance (Control Plane)
 - Application specific control planes
- Evaluation Metrics

6

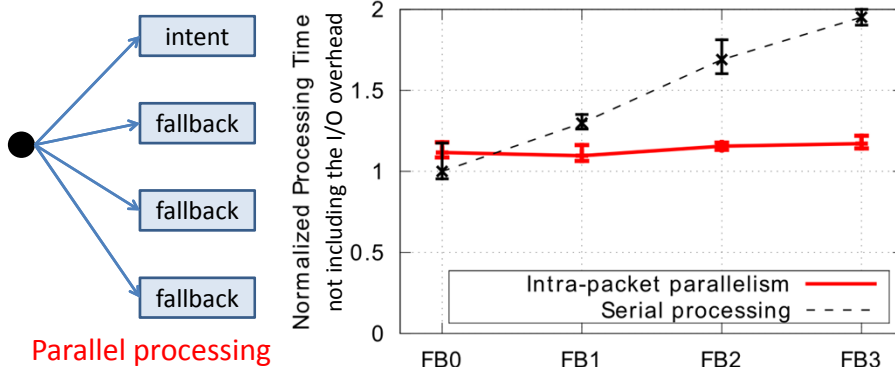
Can We Forward DAGs Rapidly? [NSDI 2012]



Click-based implementation on commodity hardware
351 K table entries based on a Route Views snapshot

7

Intra-Packet Parallelism for Bounded Processing Cost



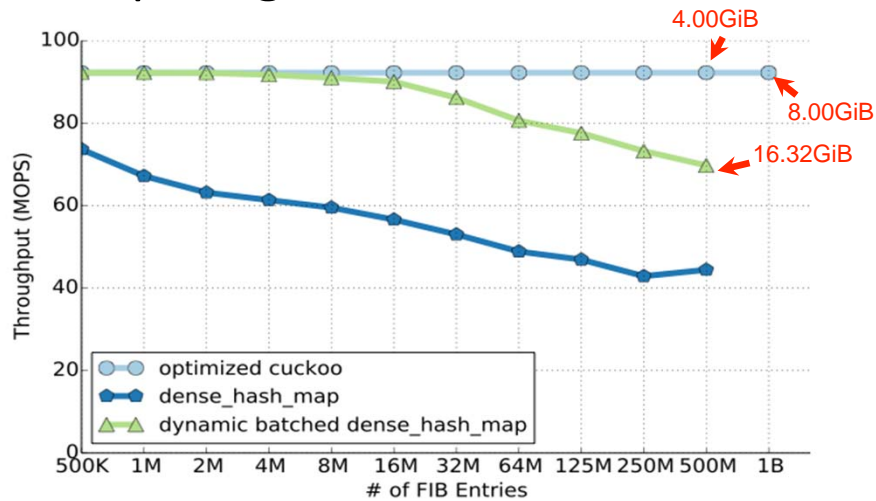
8

Large Flat Lookup Tables

- Can we build an x86-based software router that...
 - Handles 8x 10GbE ports at full line-rate
 - Handles arbitrarily large flat lookup fwd tables
 - Flow, host, and content routing as imagined uses; but
 - Also “build it, will come?” — raising expectations for what is possible from hardware!
- CuckooSwitch [CoNEXT 2013]

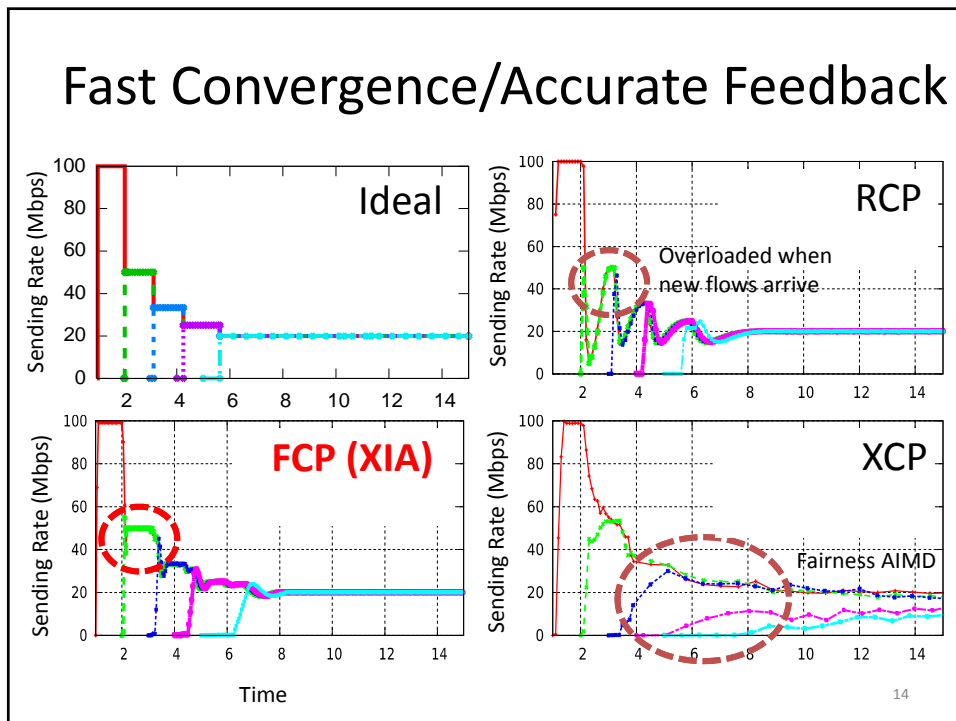
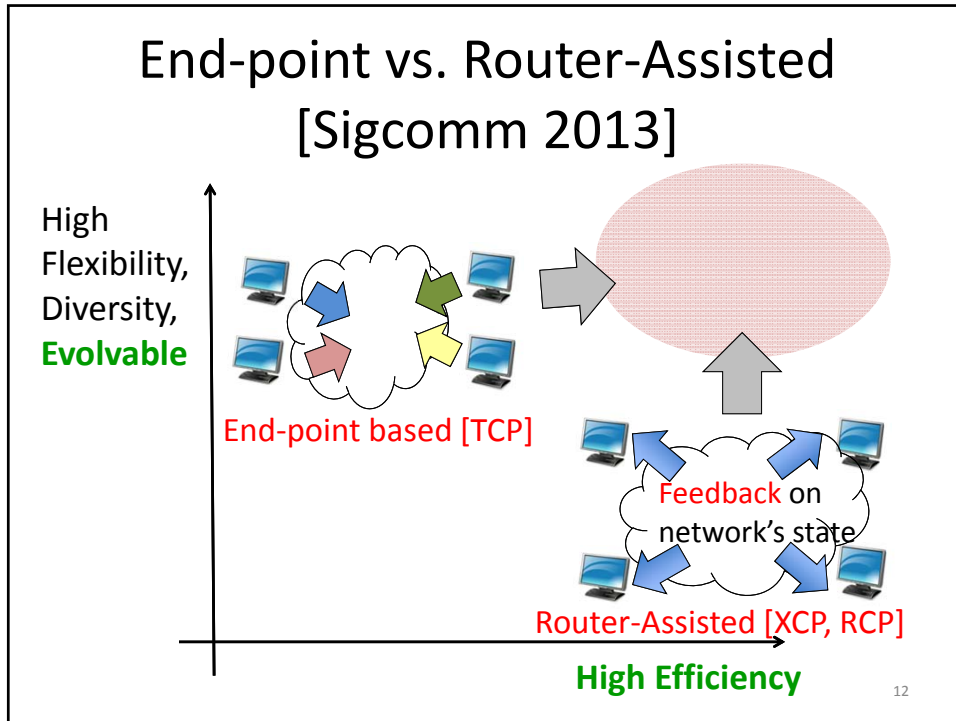
9

Comparing with Other Hash Tables



XIA packet processing can scale.

10

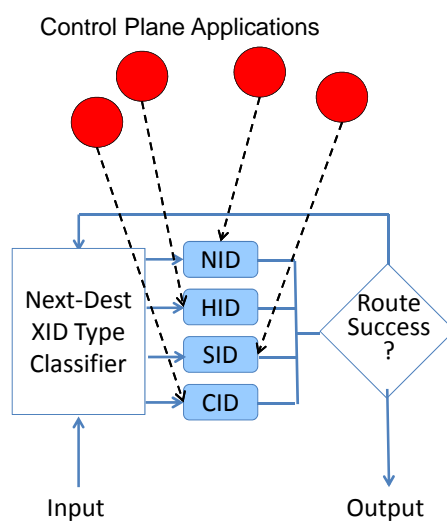


Outline

- XIA Performance Challenges/Opportunities
- Packet Processing Performance (Data Plane)
 - Processing DAGs
 - Large flat lookup tables
 - Congestion control
- Network-Wide Performance (Control Plane)
 - Application specific control planes
- Evaluation Metrics

15

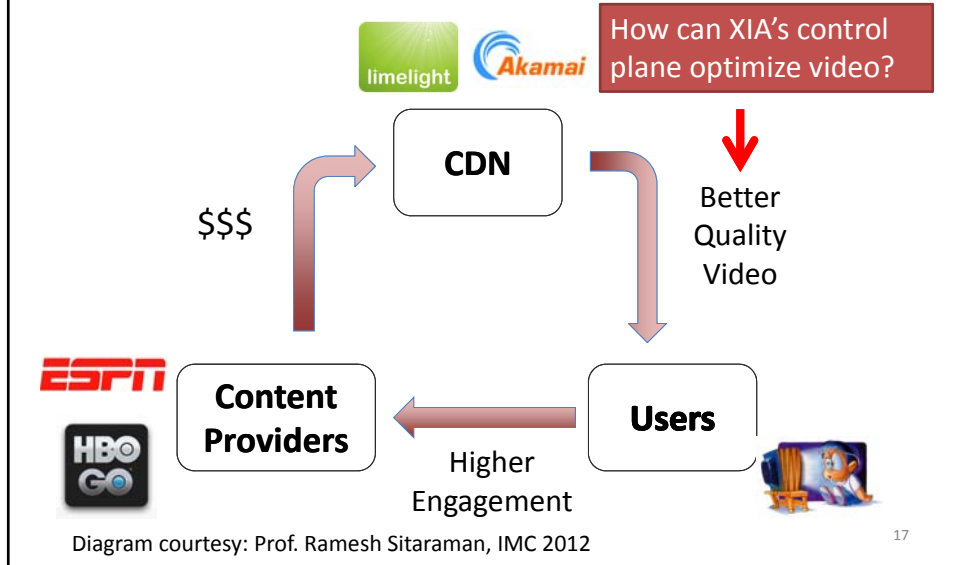
XIA Packet Processing Pipeline



- Principal-independent processing defines how to interpret the DAG
 - Core architecture
- Principal-dependent processing realizes forwarding semantics for each XID type
 - Logically: one forwarding table per XID type
 - Reality: anything goes, e.g., no forwarding table
- Control plane sets up forwarding for each principal type

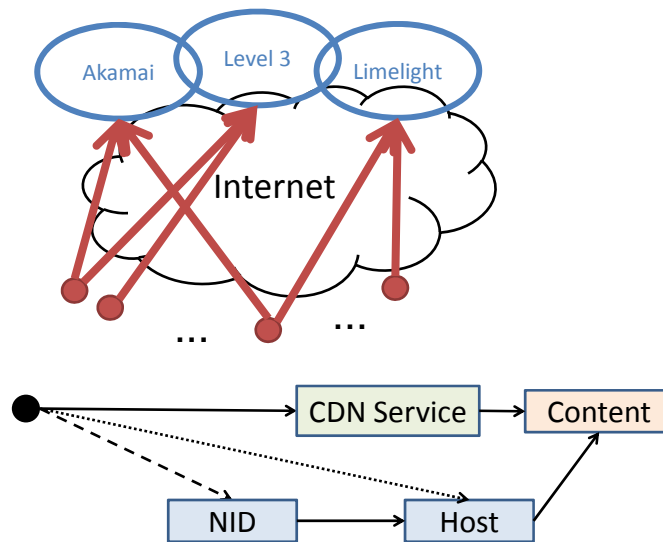
16

Control Plane: Video Case Study



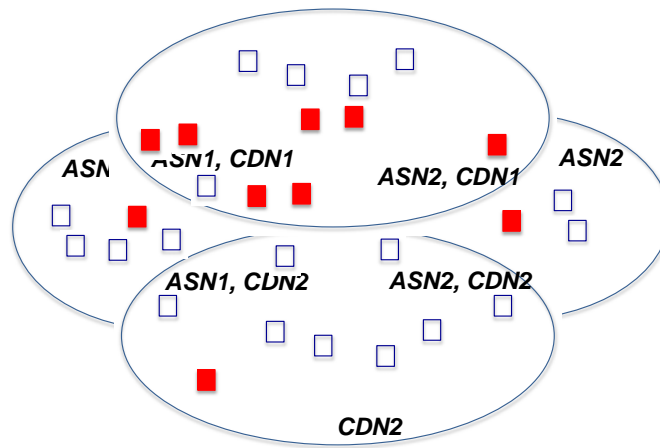
Internet Fault Management: The Opportunity of Video Layer Inference

- Video delivery involves many entities
 - Content providers
 - CDNs
 - ISPs
- Performance issues can come from any of them




Performance Fault Isolation: Critical Clusters [CoNEXT 2013]

■ Problem session □ Good quality session



20

Live Content Delivery on a CDN



$m_1: \{S\} \rightarrow \{B\}$
 $m_2: \{S\} \rightarrow \{A\}$

(a) Example scenario

$m_1: T_1$
 $m_2: \text{No bw.}$

(b) No central control

$m_1: T_2$
 $m_2: T_3$

(c) With central control

- Wide-area traffic-engineering critical for good video delivery performance
- Video is different from other services (or content)
 - Long-lived sessions, high-bandwidth constraints, adaptive behavior, etc.

21

Possible Directions

- Naming → we can give different clients different DAGs to control their routing
- Routing → we can use controls over CID routing to optimize video without impacting other traffic
- XID types → we can give video its own XID type

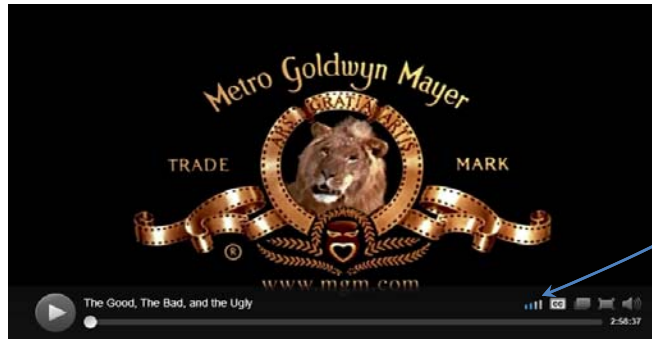

22

Outline

- XIA Performance Challenges/Opportunities
- Packet Processing Optimization (Data Plane)
 - Processing DAGs
 - Large flat lookup tables
 - Congestion control
- Network-Wide Optimization (Control Plane)
 - Application specific control planes
- Evaluation Metrics

23

How Do We Evaluate Performance?



24

How Do We Evaluate Performance?

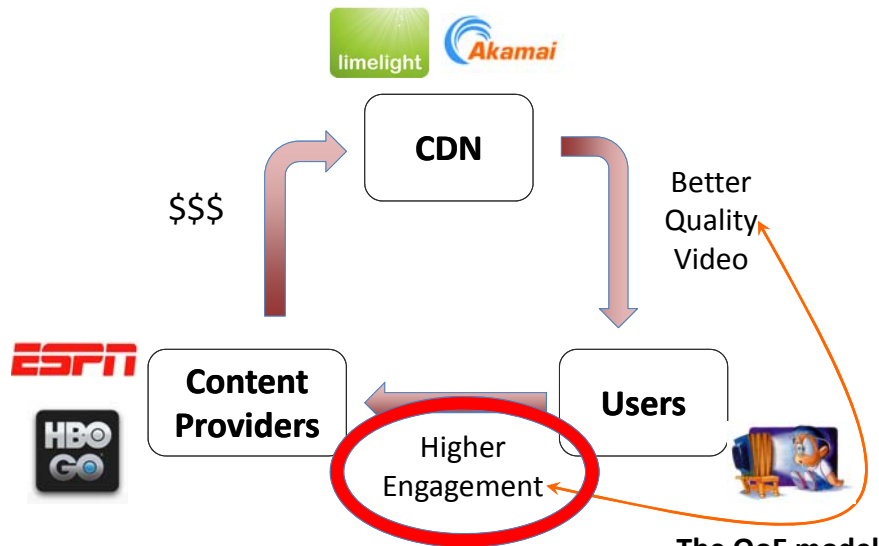
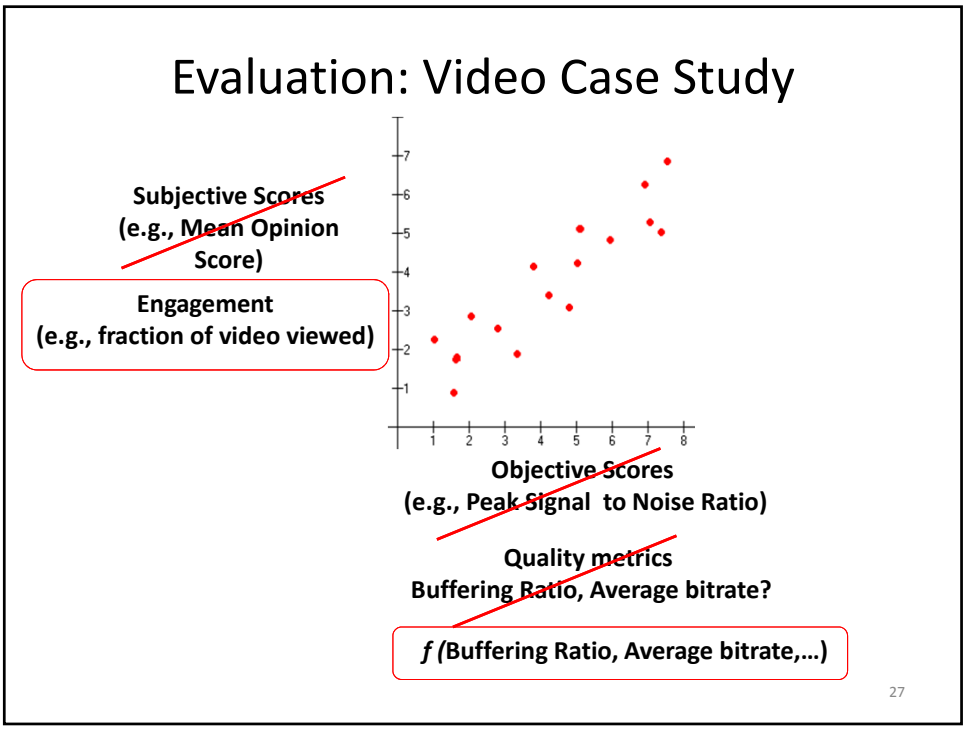
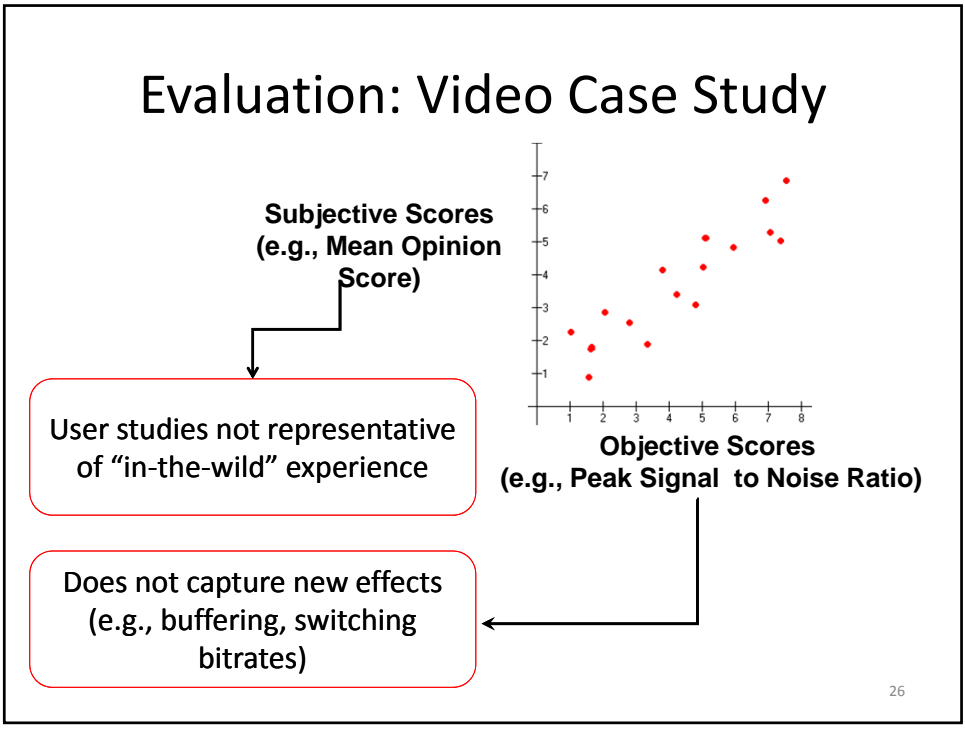
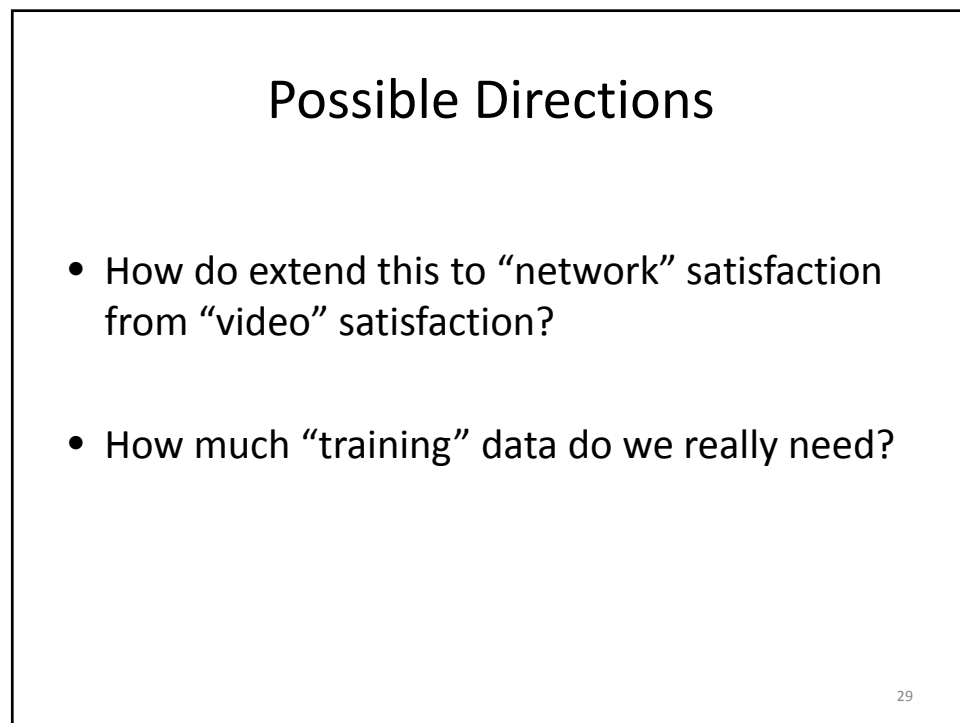
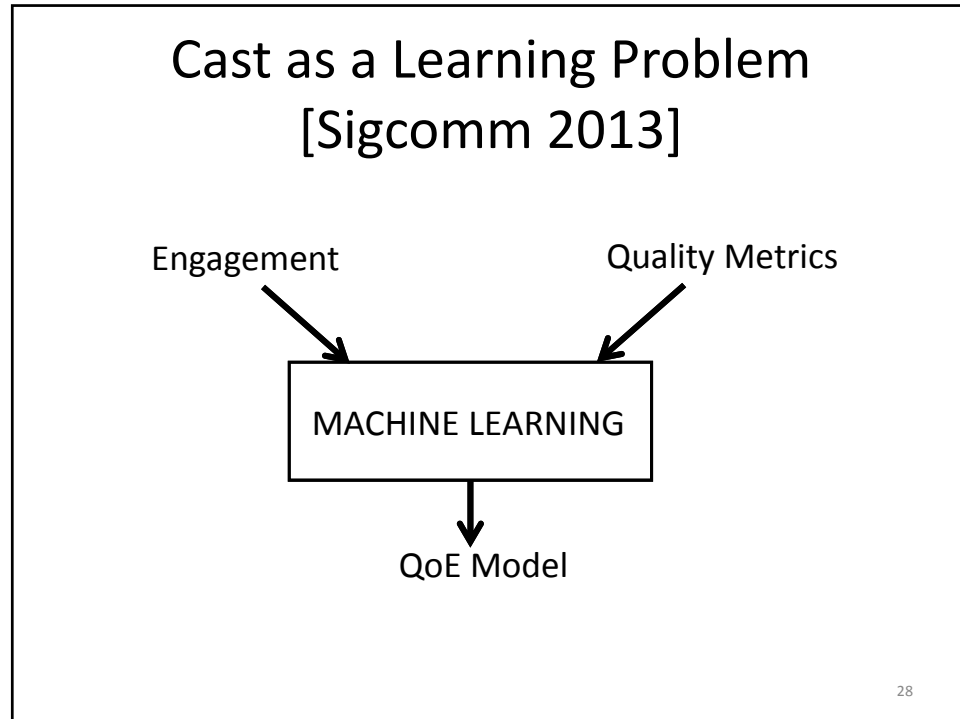


Diagram courtesy: Prof. Ramesh Sitaraman, IMC 2012

25





Outline

- XIA Performance Challenges/Opportunities
- Packet Processing Optimization (Data Plane)
 - Processing DAGs
 - Large flat lookup tables
 - Congestion control
- Network-Wide Optimization (Control Plane)
 - Application specific control planes
- Evaluation Metrics

30