Beyond Offload:
Cloudlets for Large-Scale Video Upload
[work in progress]

Pieter Simoens, Ph.D., Ghent University - CMU
Yu Xiao, Ph.D., Aalto University – CMU
Kiryong Ha, CMU
Zhuo Chen, CMU
Babu Pillai, Intel Research Lab
Mahadev Satyanarayanan, CMU

http://www.istc-cc.cmu.edu/

Today's use of tomorrow's Google Glass

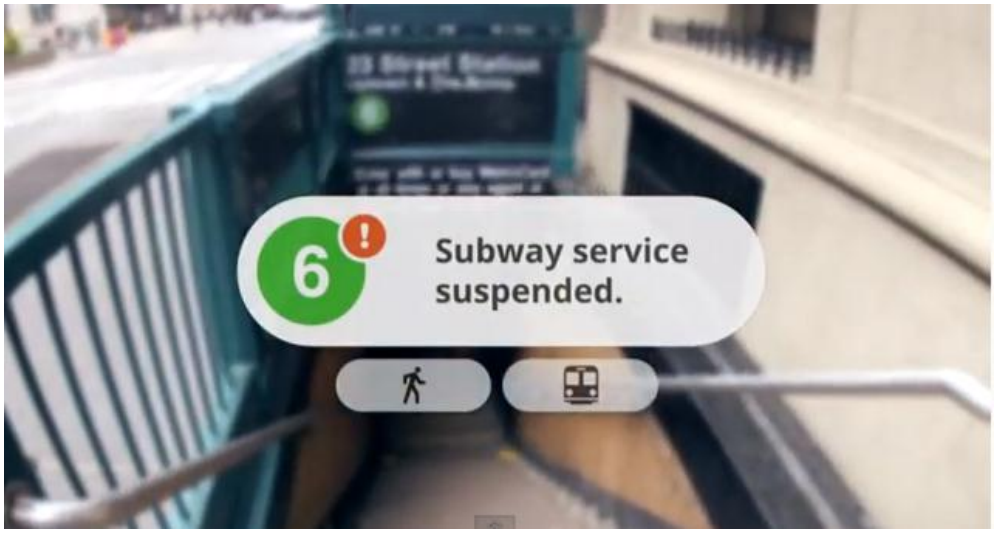

innovative user interface

Effortless capture

Huge number of uploads to be expected

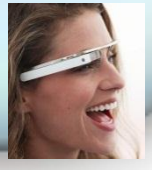

Diane Furstenberg, New York Fashion Week, Sept. 2012

Your video is/can be of value to others

© CNN

© Swellphone

Design of a scalable system for capturing, storing and content-based searching of crowd-sourced video segments

Research questions

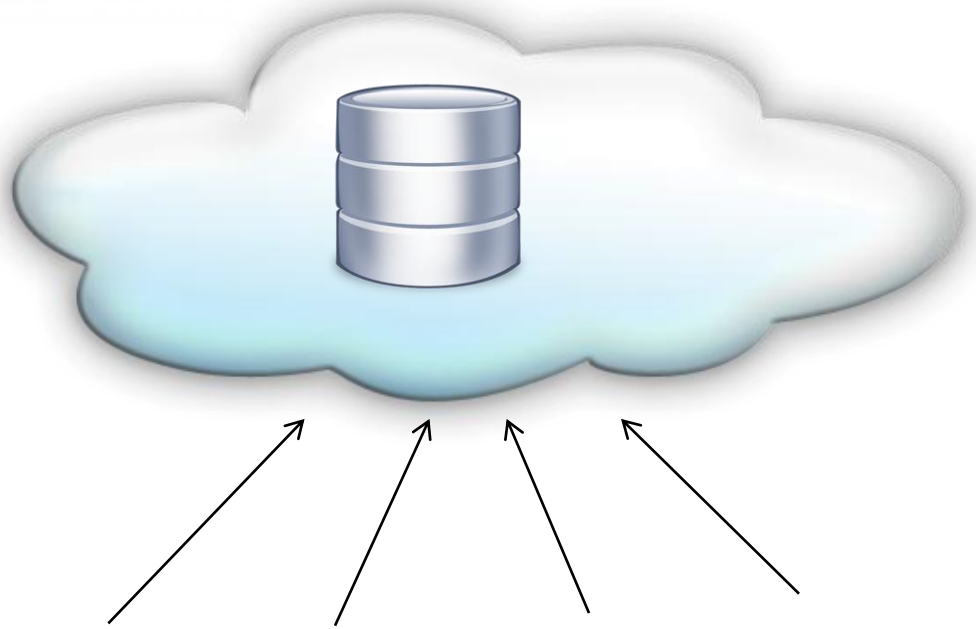

How to make the system scalable?

Why would I share my personal video?

Challenge 1: massive upload bandwidth

upload

1 hour of video uploaded each second
equivalent of **3600 users** streaming!

480p (SD): 2.5 Mbps per stream → 9 Gbps (SD)
1080p (HD): 8.5 Mbps per stream → 30.6 Gbps (HD)

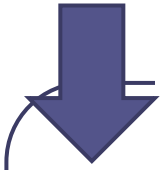

download

4 billion hours watched per month
equivalent **5.5 million users** watching

480p (SD): 1.25 Mbps per stream → 6.8 Tbps (SD)
1080p (HD): 3.75 Mbps per stream → 20.6 Tbps (HD)

Challenge 1: massive upload bandwidth

750 Gbps (SD)
2550 Gbps (HD)

Verizon Press Release
Dec 2011

verizon Solutions & Products Industries Resources

Home > About Us > News Center

About Us

- Executive Team
- Global Network
- IP Innovation
- Green Initiatives
- News Center
  - Global Archive
  - Media Contacts
  - Network Facts
  - Podcasts
- Global Events
- Career Center

News Center

Verizon First Global Service Provider to Deploy 100G on U.S. Long-Haul Network

Ciena Equipment Supports Advanced Technology With Coherent Optical Transport Solution

September 12, 2011

NEW YORK - In another industry first, Verizon has deployed 100G (gigabits per second) technology on an ultra-long-haul optical system on a portion of the company's U.S. backbone network.

distributed approach at the network edge needed
“Content Distribution Network” in reverse

Challenge 2: Efforts of and Incentives for sharing

The value of crowd-sourcing grows with the volume and diversity of entries in the video catalogue

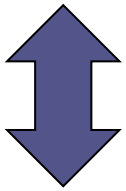

People capture personal highlights and manually select scenes to share

incentive **to capture**

make catalogue more **diverse** with daily and common situations

incentive **to share**

make scene selection as **easy** as the capture process

Challenge 2a: Incentive to *capture*

Make daily, common scenes financially attractive

Many parties might be willing to pay for access to the catalogue

- videos reflect personal taste
  - advertisers, tourist bureaus …
- hindsight view on time and place
  - insurance companies

pay-per-view

GigaSight service provider

royalties

consumers

producers

Challenge 2b: Incentive to *share*

Trusted entity needed to handle original video
Privacy guarantee must be effortless for the user

User trade-off: personal and context-sensitive

Denaturing: removing private scenes from captured video

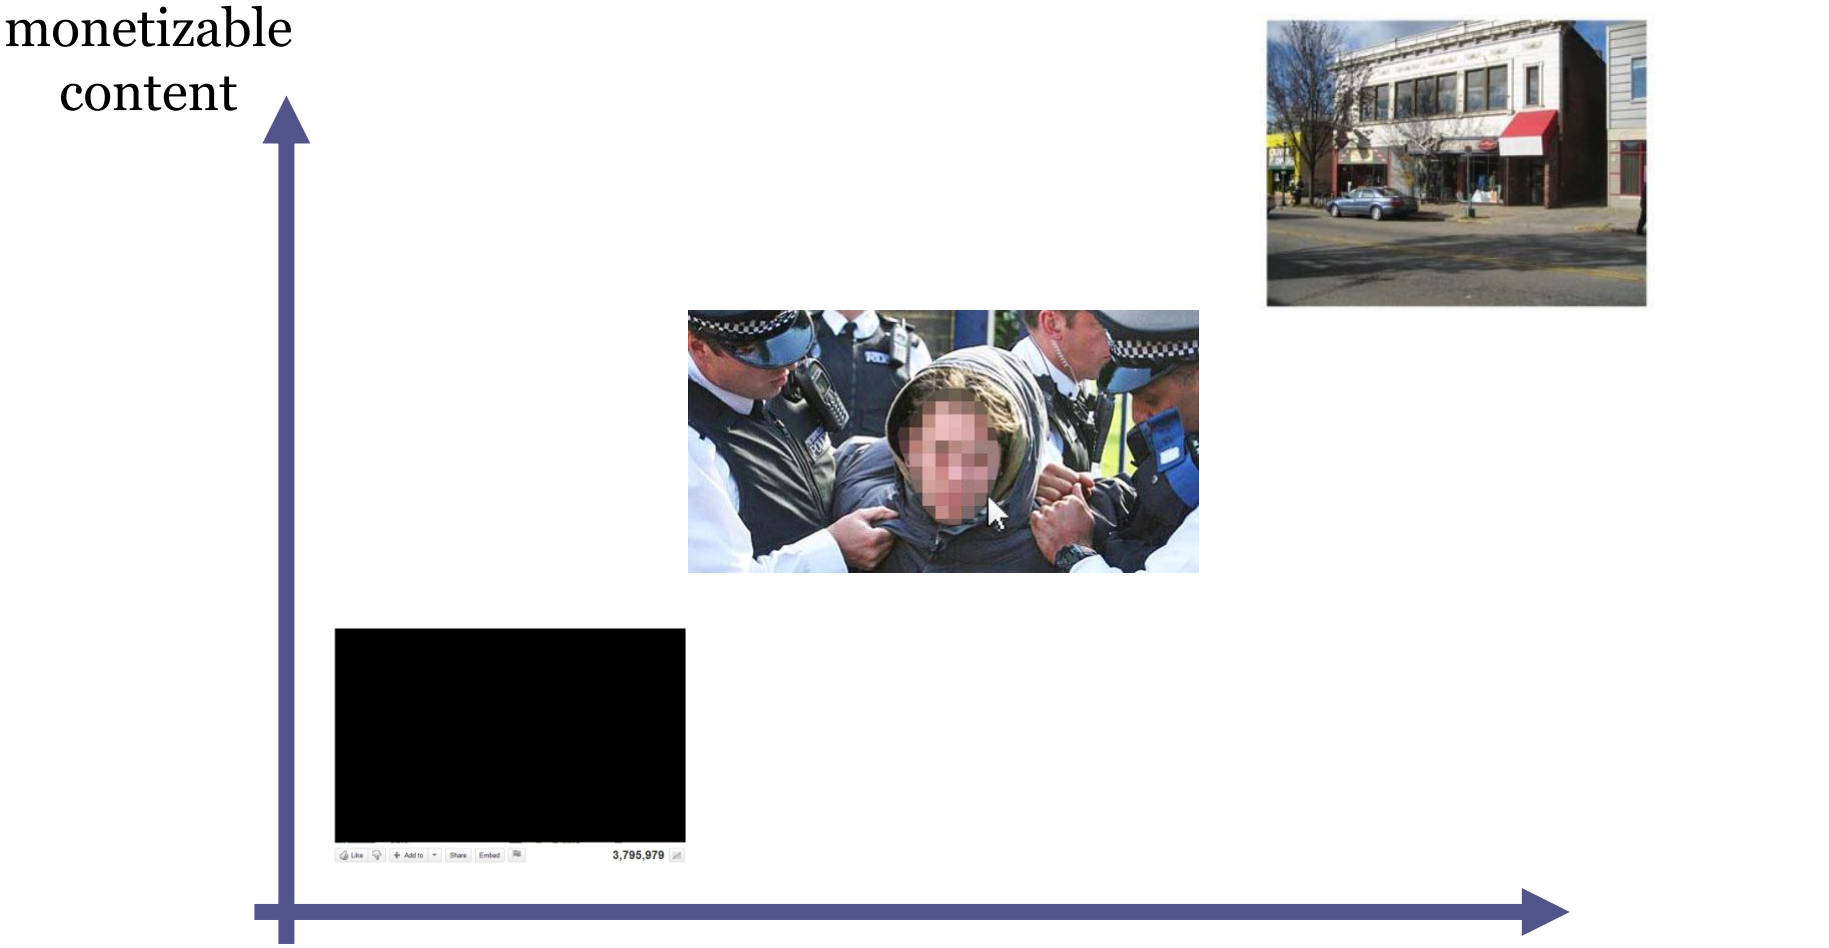

Cloudlet-based architecture

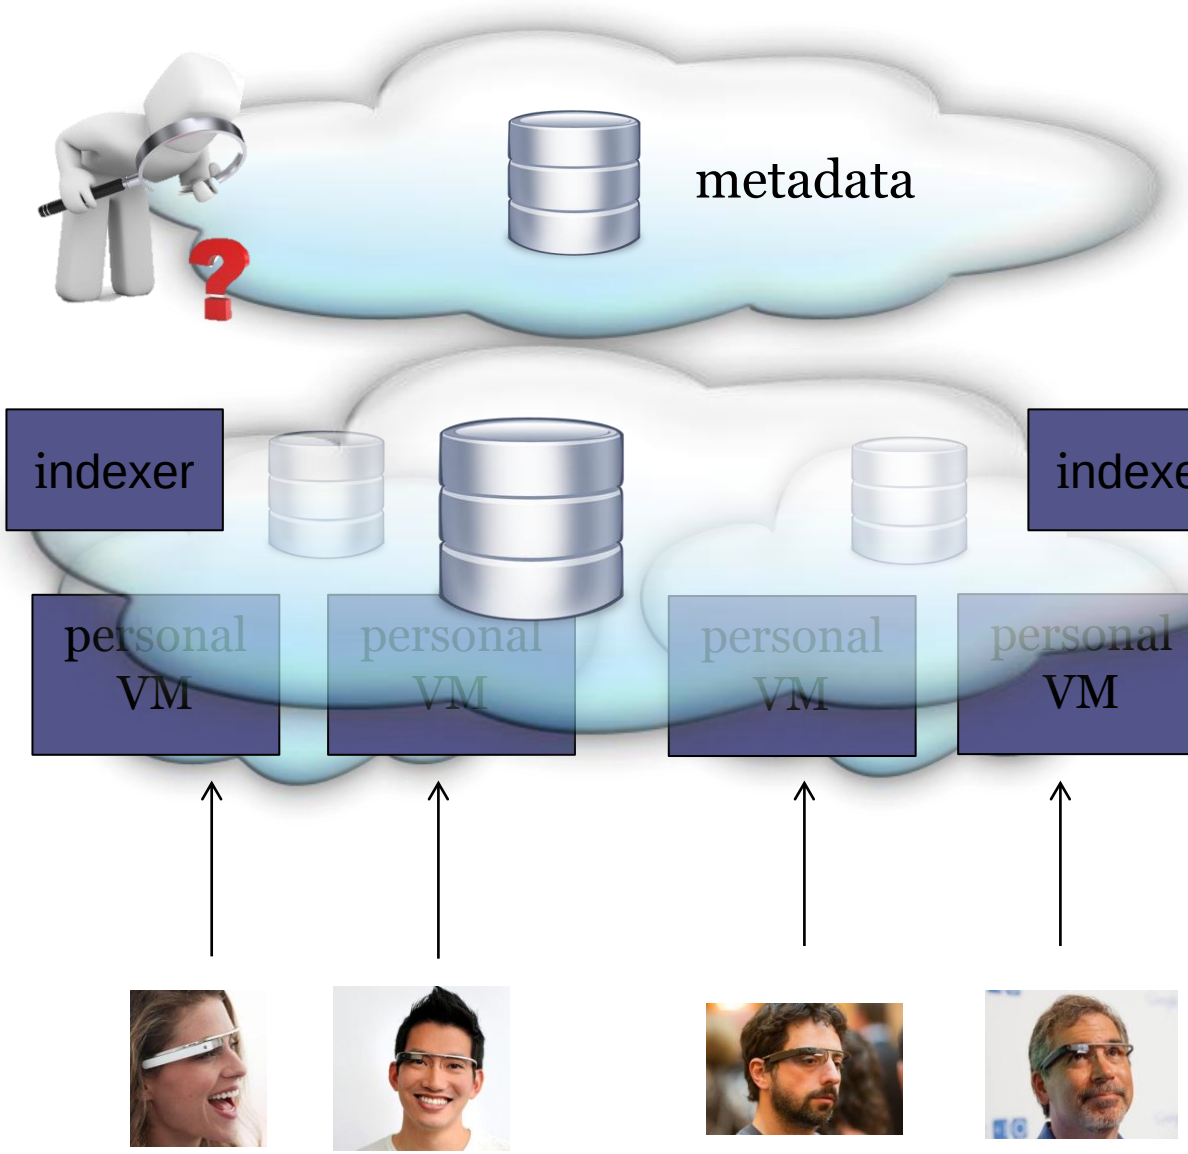

content-based search
scoped by time and location

storage and indexing
at the network edge

effortless denaturing in
personal VM

easy configuration
of privacy settings

Upload and denaturing

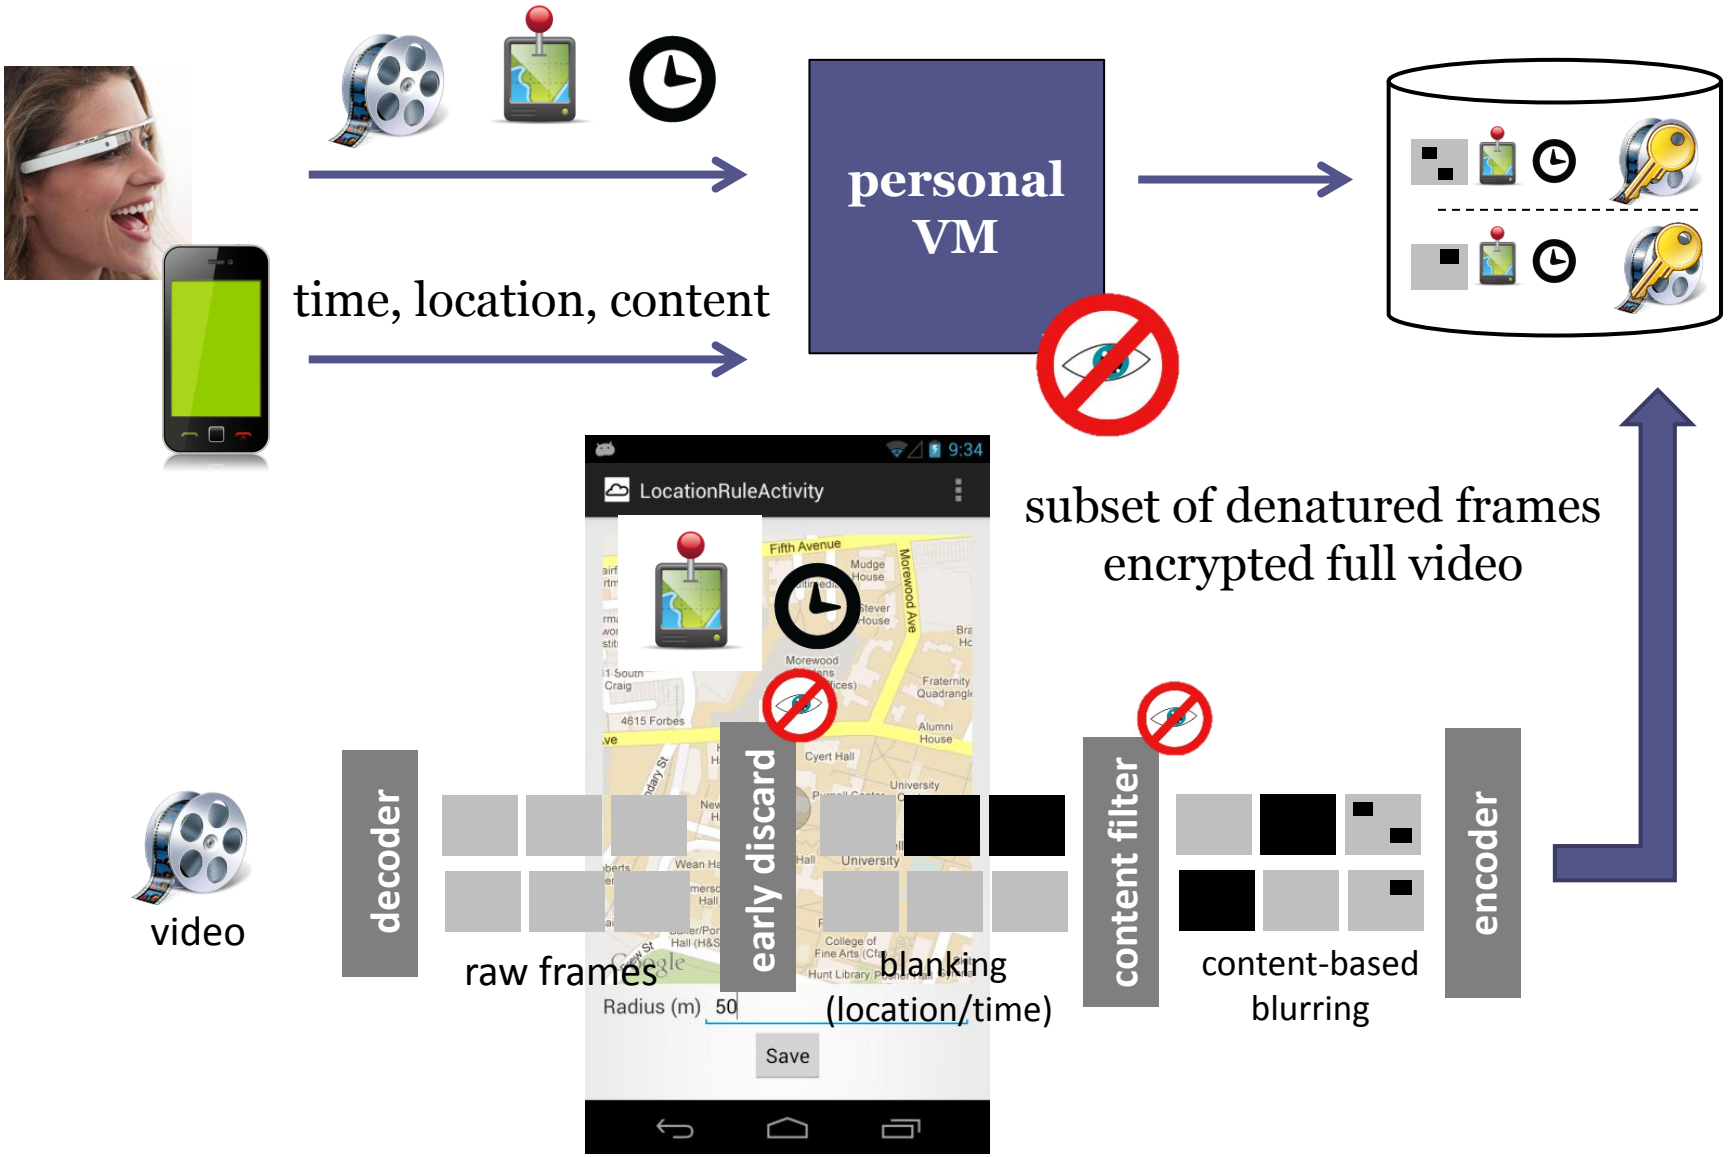

Denaturing

Indexing

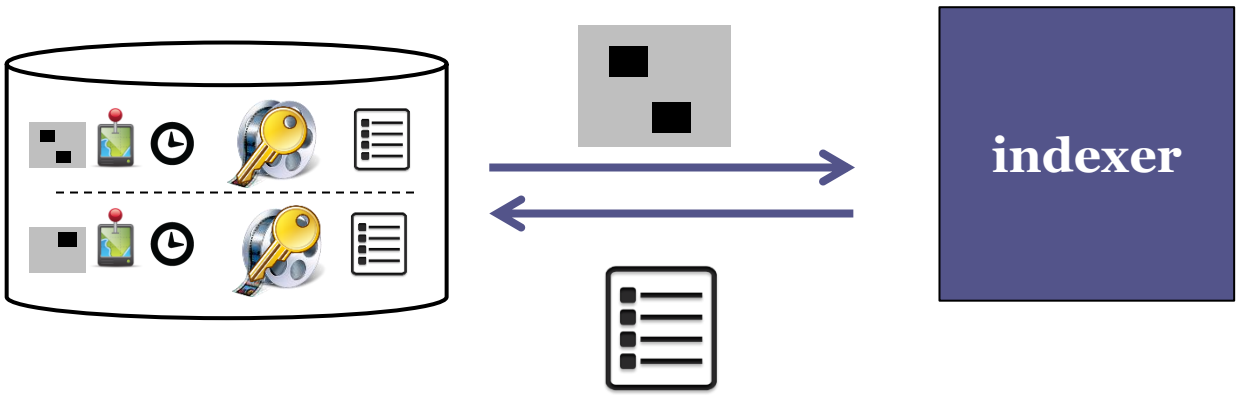

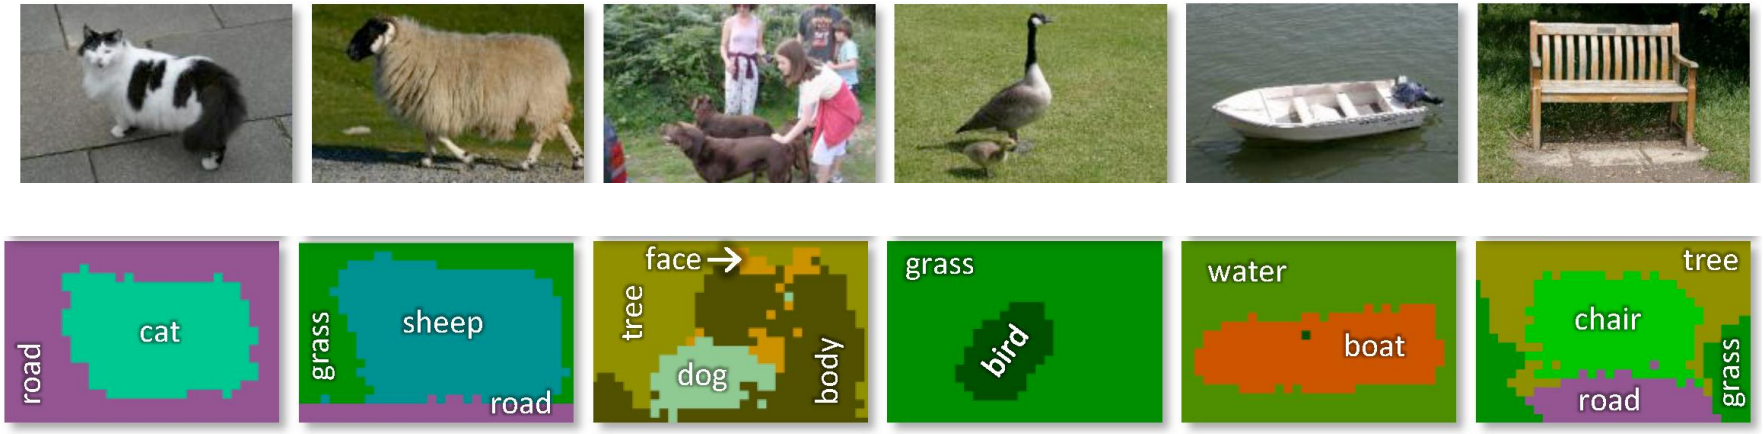

© Semantic Texton Forests for Image Categorization and Segmentation, J. Shotten et al., IEEE Conf on Computer Vision and Pattern Recognition, 2008

Search

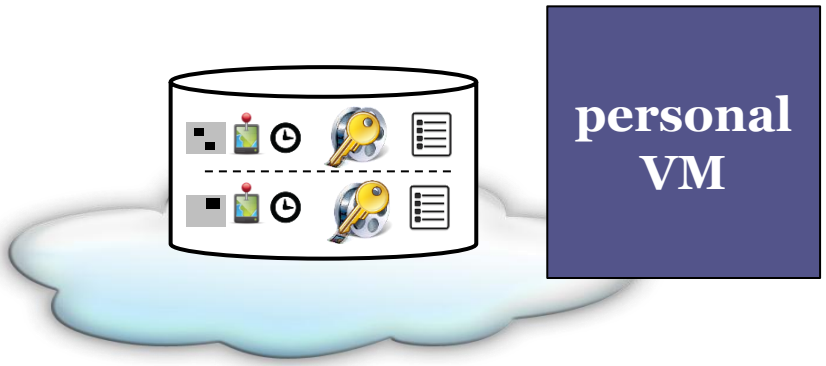

ad-hoc denaturing

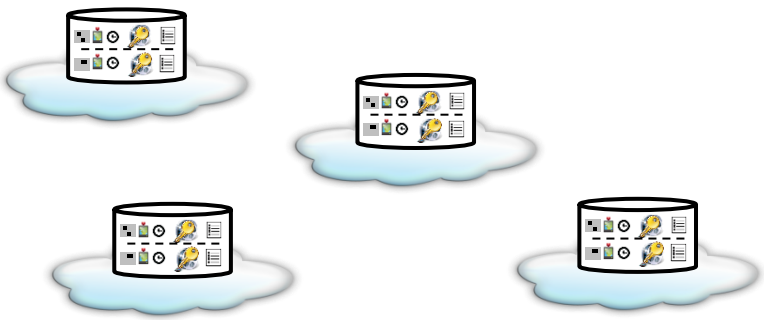

tags
location
time

Find me all red cars in Pittsburgh today between 11 AM and 2 PM

GigaSight

(1) search on subset of denatured frames

(2) view full length of denatured video

Results on wireless throughput

Individual throughput limited by mobile hardware and software stack
Cumulative throughput limited by channel collisions

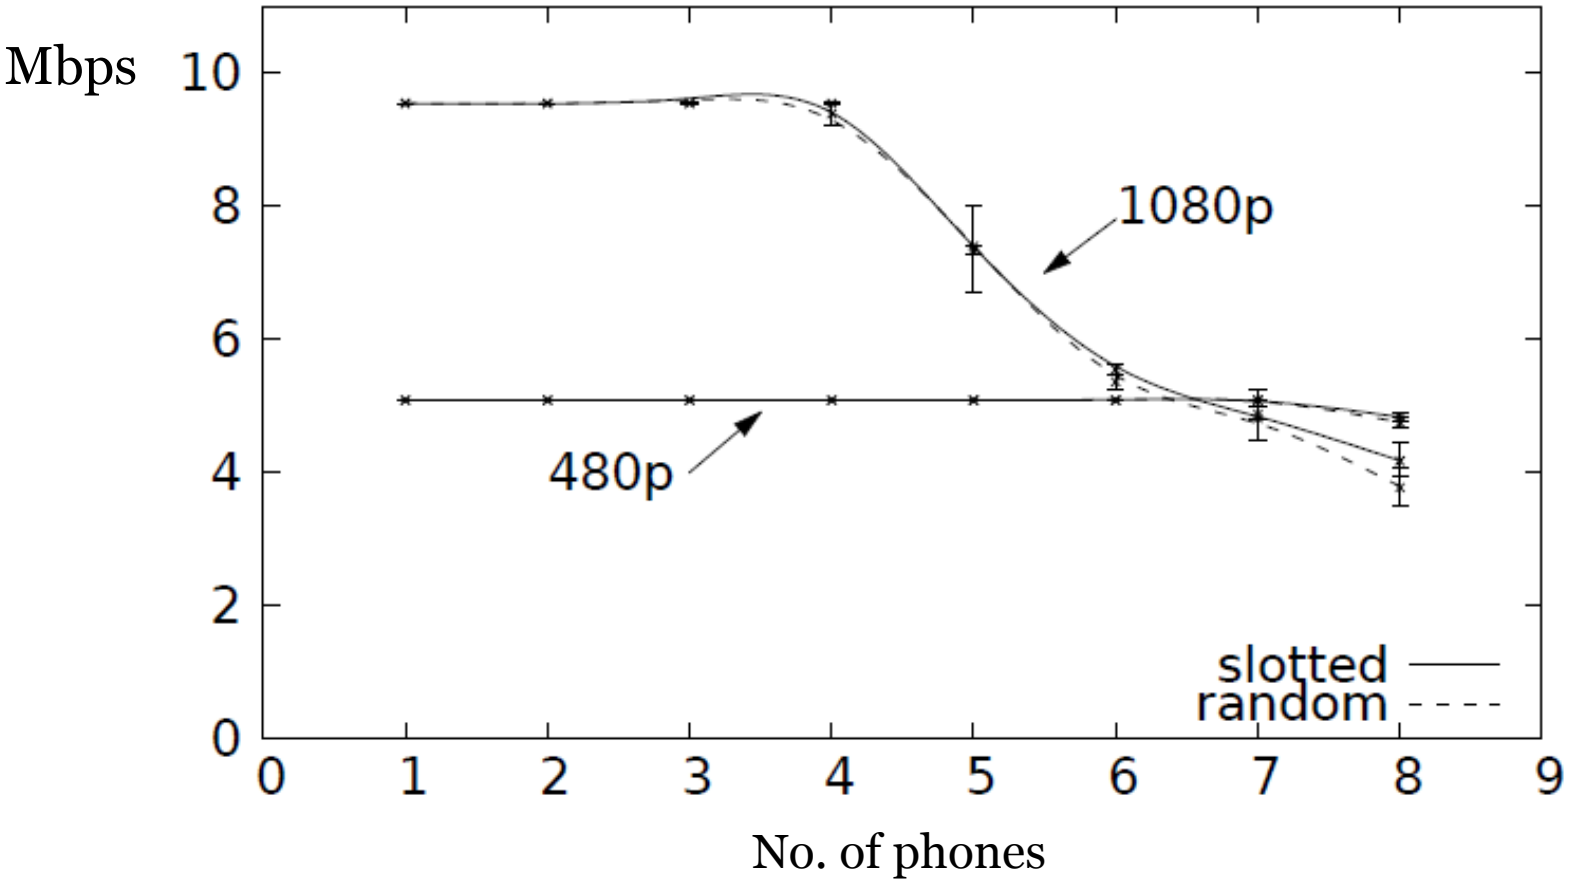

Denaturing

Computer vision algorithms are CPU bound
Resolution balances throughput and accuracy

Detection accuracy normalized against 1080p resolution

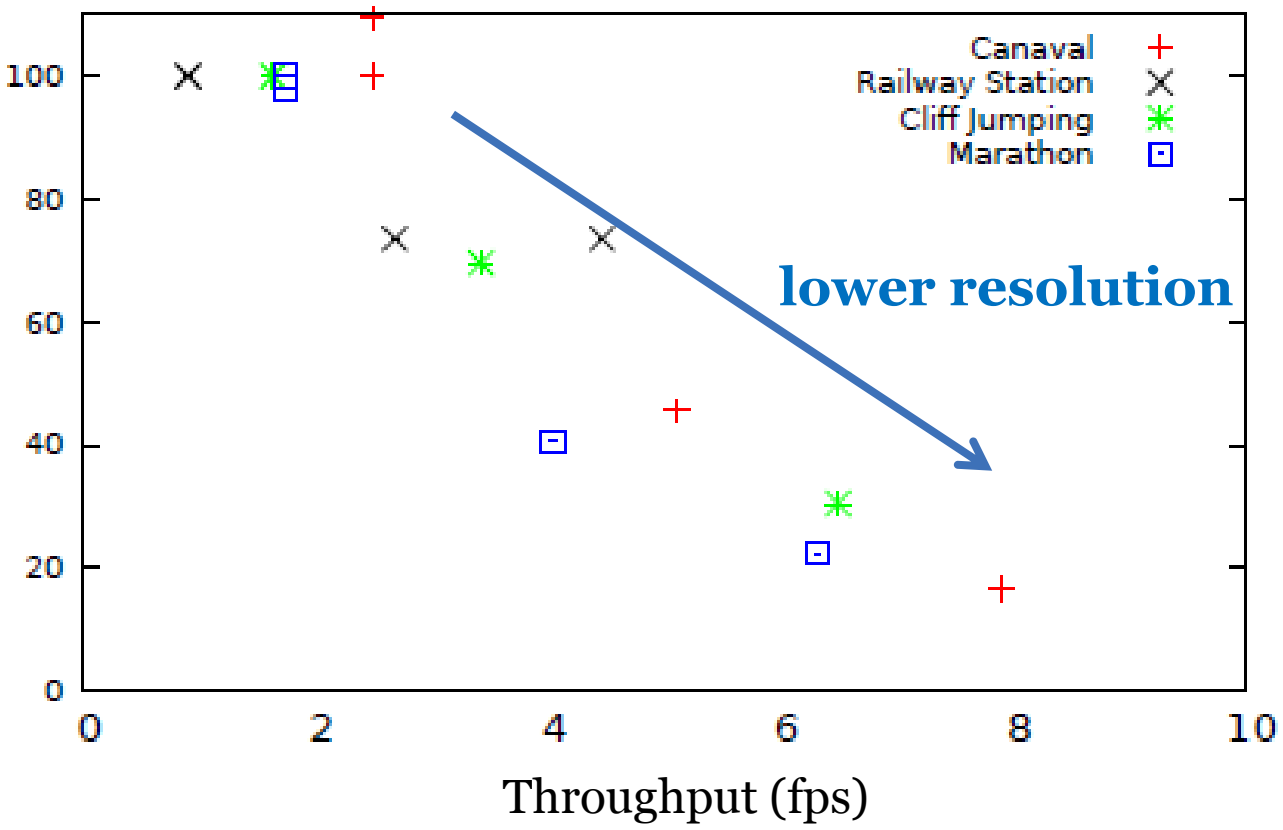

Indexing and search

Resolution has less to no impact on accuracy

1080p

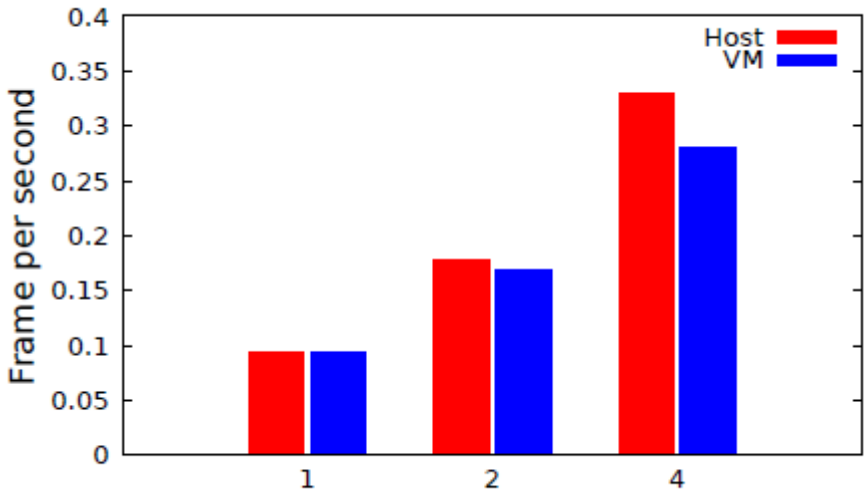

480p

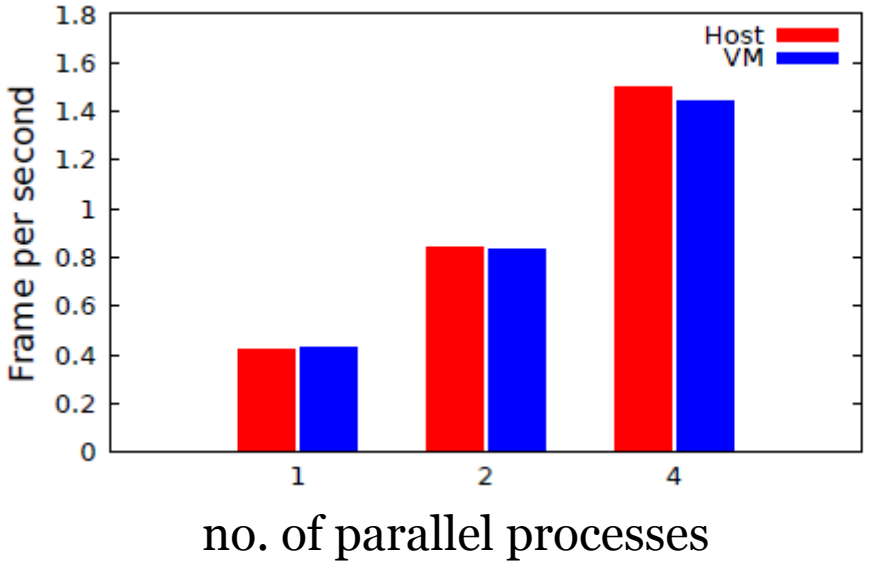

Conclusions

GigaSight is framework for content-based search on crowd-sourced, denatured videos.

Key architectural design choices

- effortless capture AND privacy preservation to stimulate sharing
- distributed at the edge
  - personal VM for privacy
  - bandwidth

Scalability

- frame rate is key parameter
- ad-hoc denaturing
- CPU and bandwidth bounded
- computer vision algorithms still need a lot of improvement

Thank you for your attention

psimoens@cs.cmu.edu

yuxiao@cs.cmu.edu

padmanabhan.s.pillai@intel.com

krha@cs.cmu.edu

zhuoc@cs.cmu.edu

satya@cs.cmu.edu

Back Up

Contributors

- Pieter Simoens, Ph.D., Ghent University - CMU
- Yu Xiao, Ph.D., Aalto University - CMU
- Kiryong Ha, Carnegie Mellon
- Zhuo Chen, Carnegie Mellon
- Babu Pillai, Intel Labs Research
- Satya, prof., Carnegie Mellon