

Cloudlets

At the Leading Edge of Cloud-Mobile Convergence

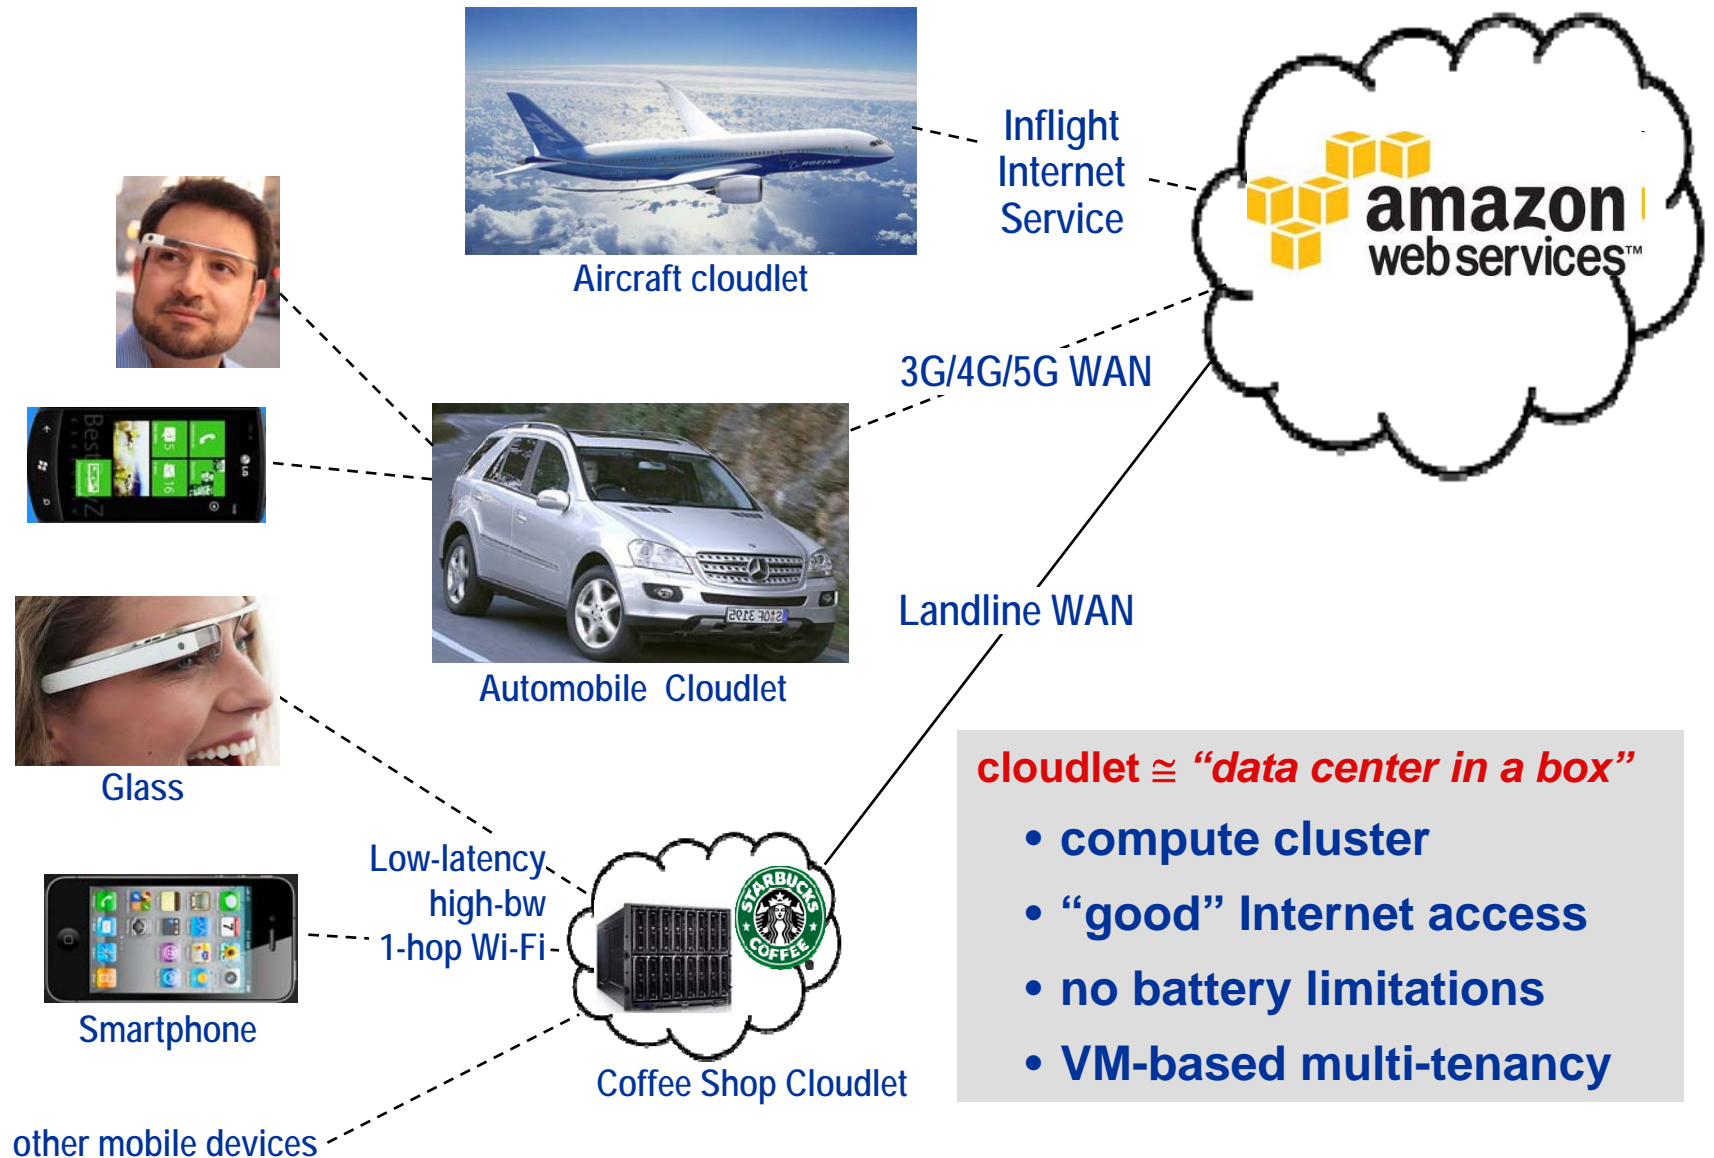
Mahadev Satyanarayanan
School of Computer Science
Carnegie Mellon University

<http://elijah.cs.cmu.edu>

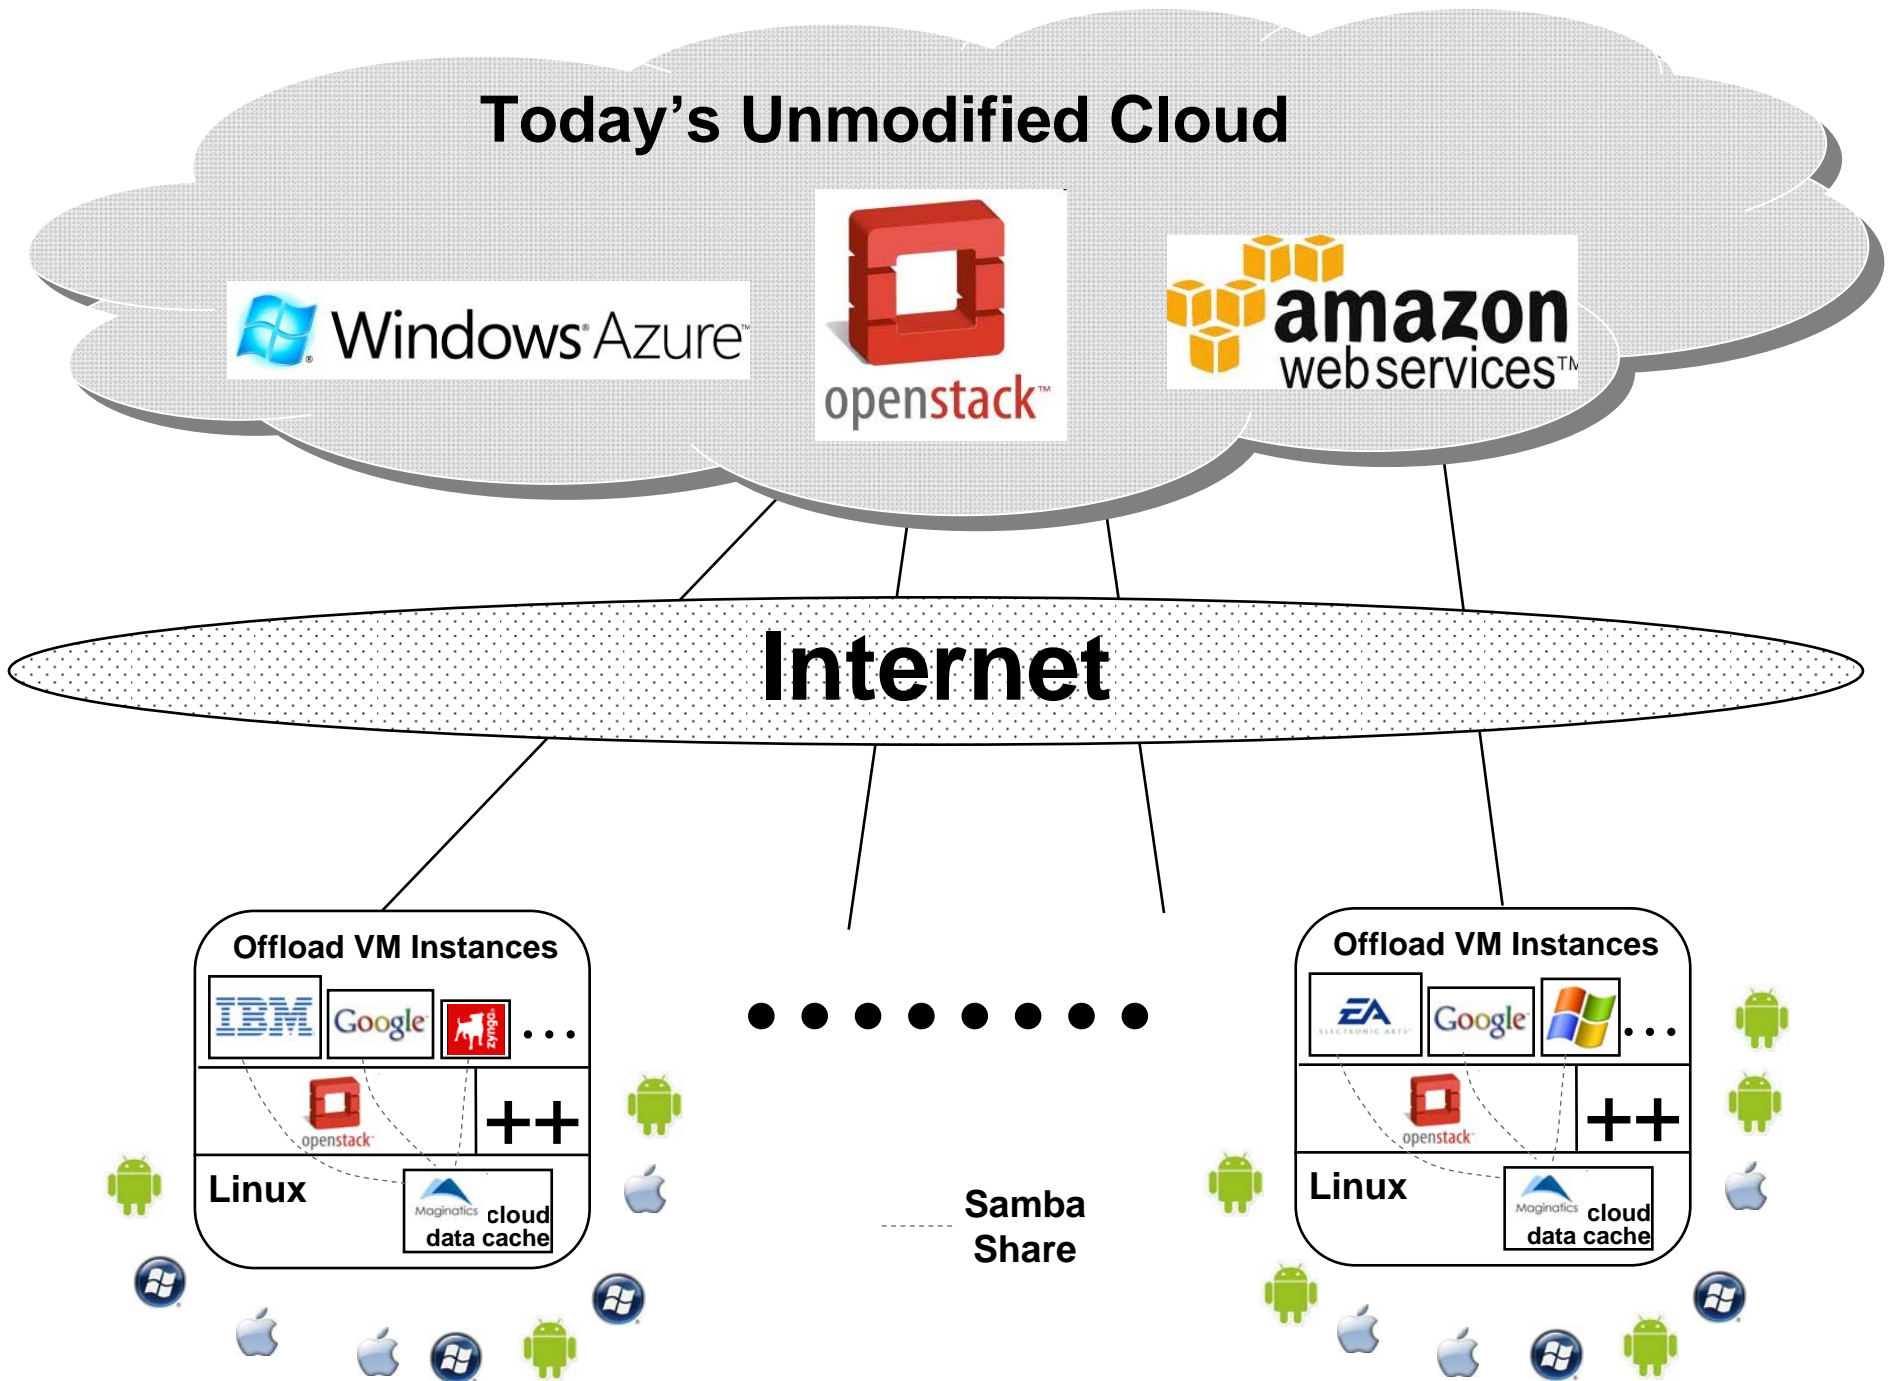
Joint work with: Yoshihisa Abe (CMU), Victor Bahl (Microsoft Research), Vas Bala (IBM Research), Jeff Boleng (CMU-SEI), Ramon Caceres (AT&T Research), Zhou Chen (CMU), Sarah Clinch (Lancaster University), Nigel Davies (Lancaster University), Roxana Geambasu (Columbia University), Benjamin Gilbert (CMU), Kiryong Ha (CMU), Jan Harkes (CMU), Martial Hebert (CMU), Wenlu Hu (CMU), Kaustubh Joshi (AT&T Research), Grace Lewis (CMU-SEI), Ed Morris (CMU-SEI), Padmanabhan Pillai (Intel Labs), Wolfgang Richter (CMU), Dan Siewiorek (CMU), Soumya Simanta (CMU-SEI), Pieter Simoens (CMU & University of Ghent), Roy Want (Google), Yu Xiao (CMU & Aalto University)

Bring the Cloud Closer

Create a Small Cloudlet Nearby



Today's Unmodified Cloud



Three Questions

- 1. Does Latency Really Matter?**
- 2. Can Cloudlets Self-Manage?**
- 3. Can Cloudlets Reduce Bandwidth Demand?**

1. Does Latency Really Matter?

“The Impact of Mobile Multimedia Applications on Data Center Consolidation”



Ha, K. Pillai, P., Lewis, G., Simanta, S., Clinch, S., Davies, N., Satyanarayanan, M.
IEEE International Conference on Cloud Engineering, San Francisco, CA, March 2013

“How Close is Close Enough? Understanding the Role of Cloudlets in Supporting Display Appropriation by Mobile Users”

Clinch, S., Harkes, J., Friday, A., Davies, N., Satyanarayanan, M.
Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom 2012), Lugano, Switzerland, March 2012.



Face Recognition

CDF of Response Times (milliseconds)
Amazon EC2 offload



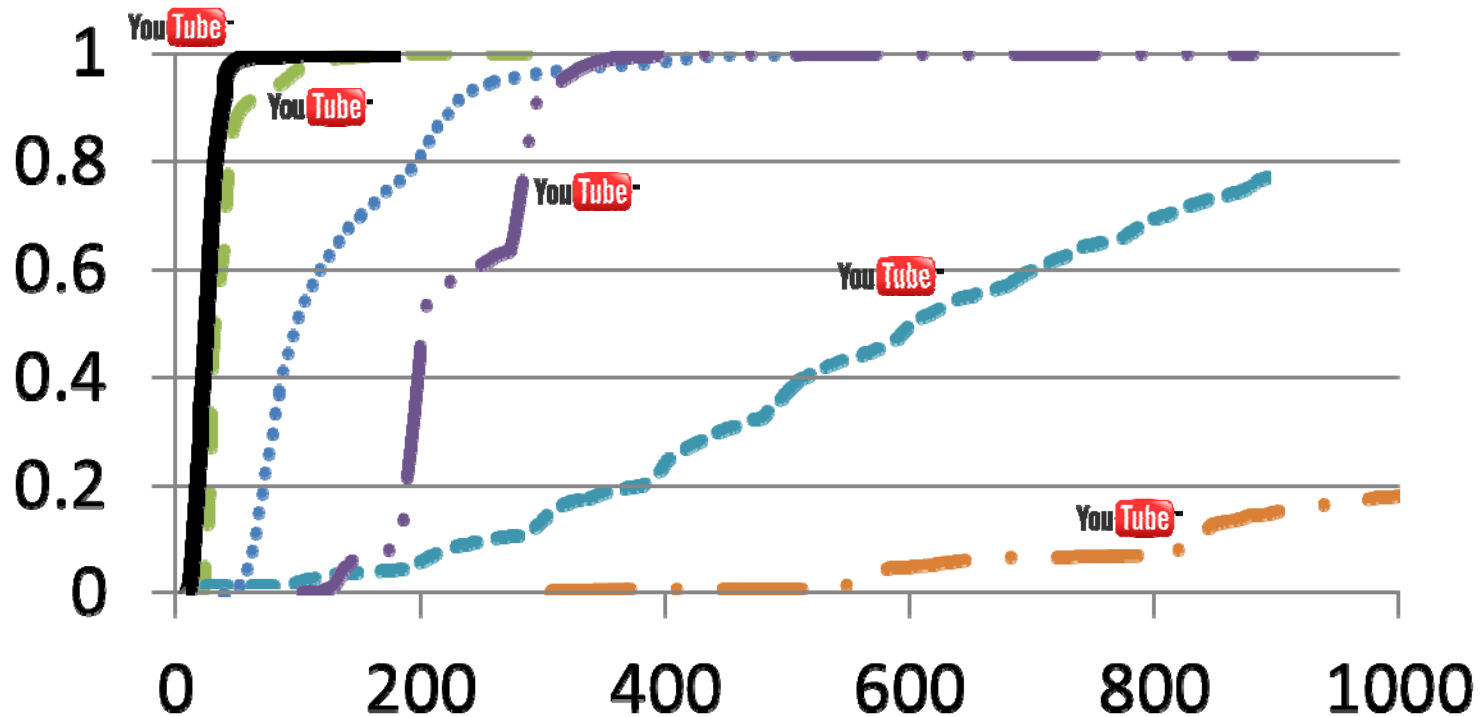

Augmented Reality

CDF of Response Times (milliseconds)
Amazon EC2 offload



Simulation-based Graphics

CDF of Response Times (milliseconds)
Amazon EC2 offload



Won't Mobile Devices Get Better?

Year	Typical Server		Typical Handheld or Wearable	
	Processor	Speed	Device	Speed
1997	Pentium® II	266 MHz	Palm Pilot	16 MHz
2002	Itanium®	1 GHz	Blackberry 5810	133 MHz
2007	Intel® Core™ 2	9.6 GHz (4 cores)	Apple iPhone	412 MHz
2011	Intel® Xeon® X5	32 GHz (2x6 cores)	Samsung Galaxy S2	2.4 GHz (2 cores)
2013	Intel® Xeon® E5	64 GHz (2x12 cores)	Google Glass OMAP 4430	2.4 GHz (2 cores)

Adapted from: J. Flinn, "Cyber Foraging: Bridging Mobile and Cloud Computing via Opportunistic Offload." Morgan & Claypool Publishers, 2012.

Mobility demands a premium

- weight, size, battery life, heat dissipation, ...
- Moore's Law applied differently in this world

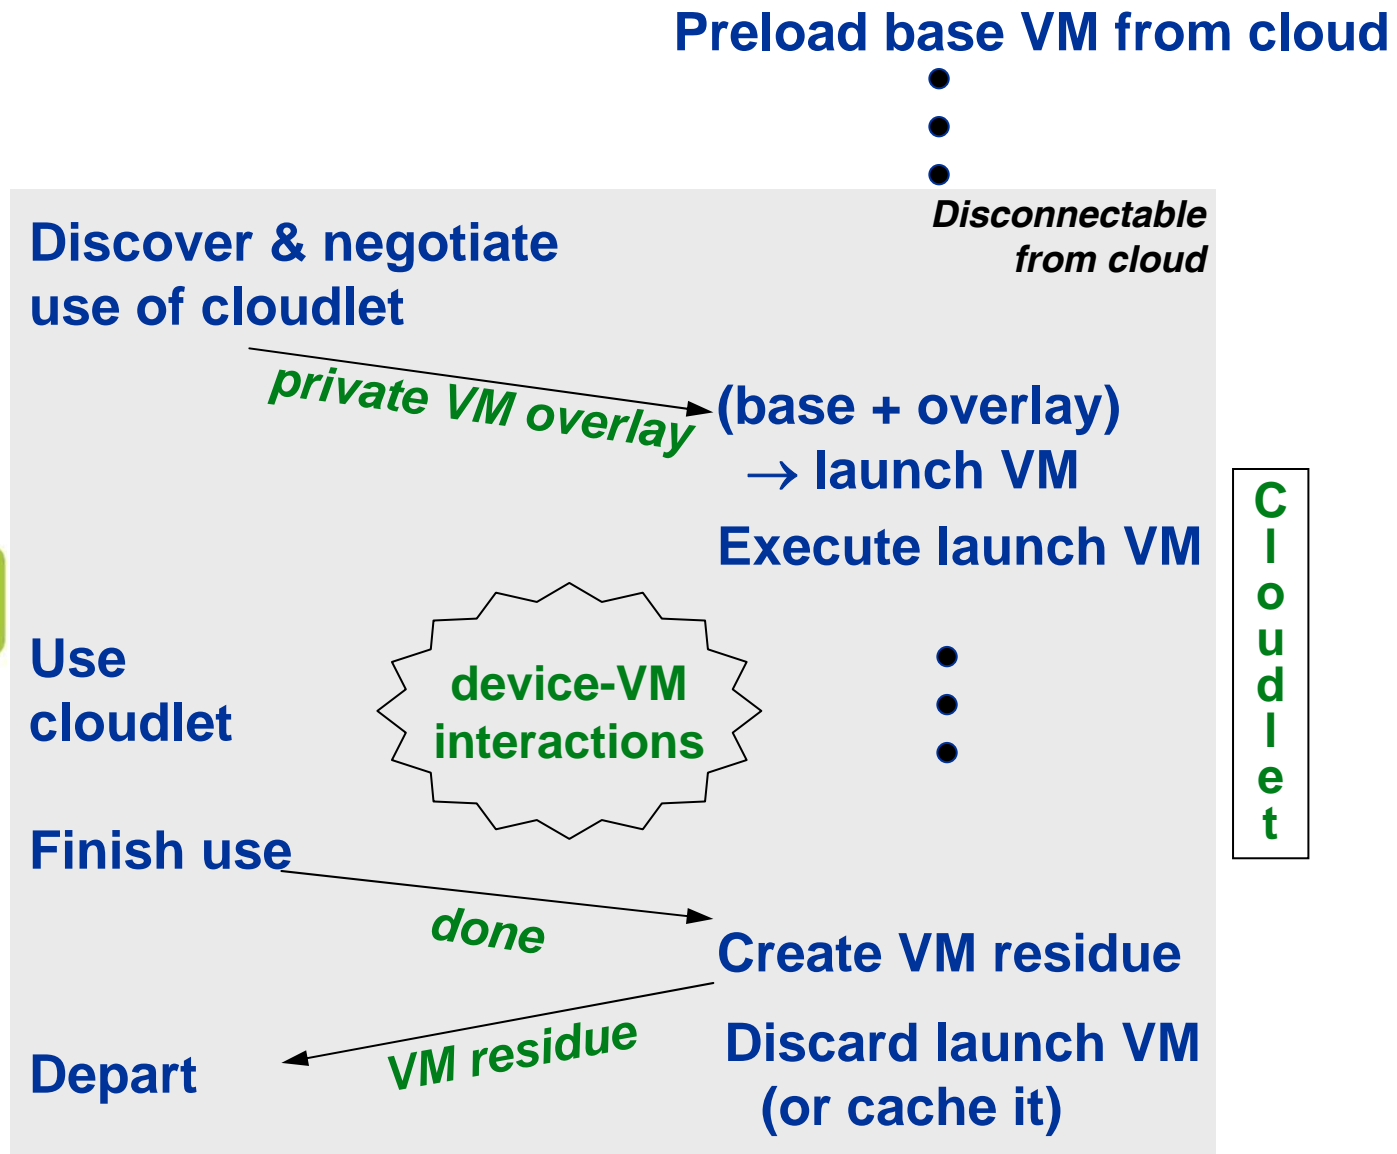
2. Can Cloudlets Self-Manage?

“Just-in-Time Provisioning for Cyber Foraging”

Ha, K., Pillai, P., Richter, W., Abe, Y., Satyanarayanan, M.

Proceedings of the Eleventh International Conference on Mobile Computing Systems, Applications and Services (MobiSys 2013), Taipei, Taiwan, June 2013

Dynamic VM Synthesis



Optimizations

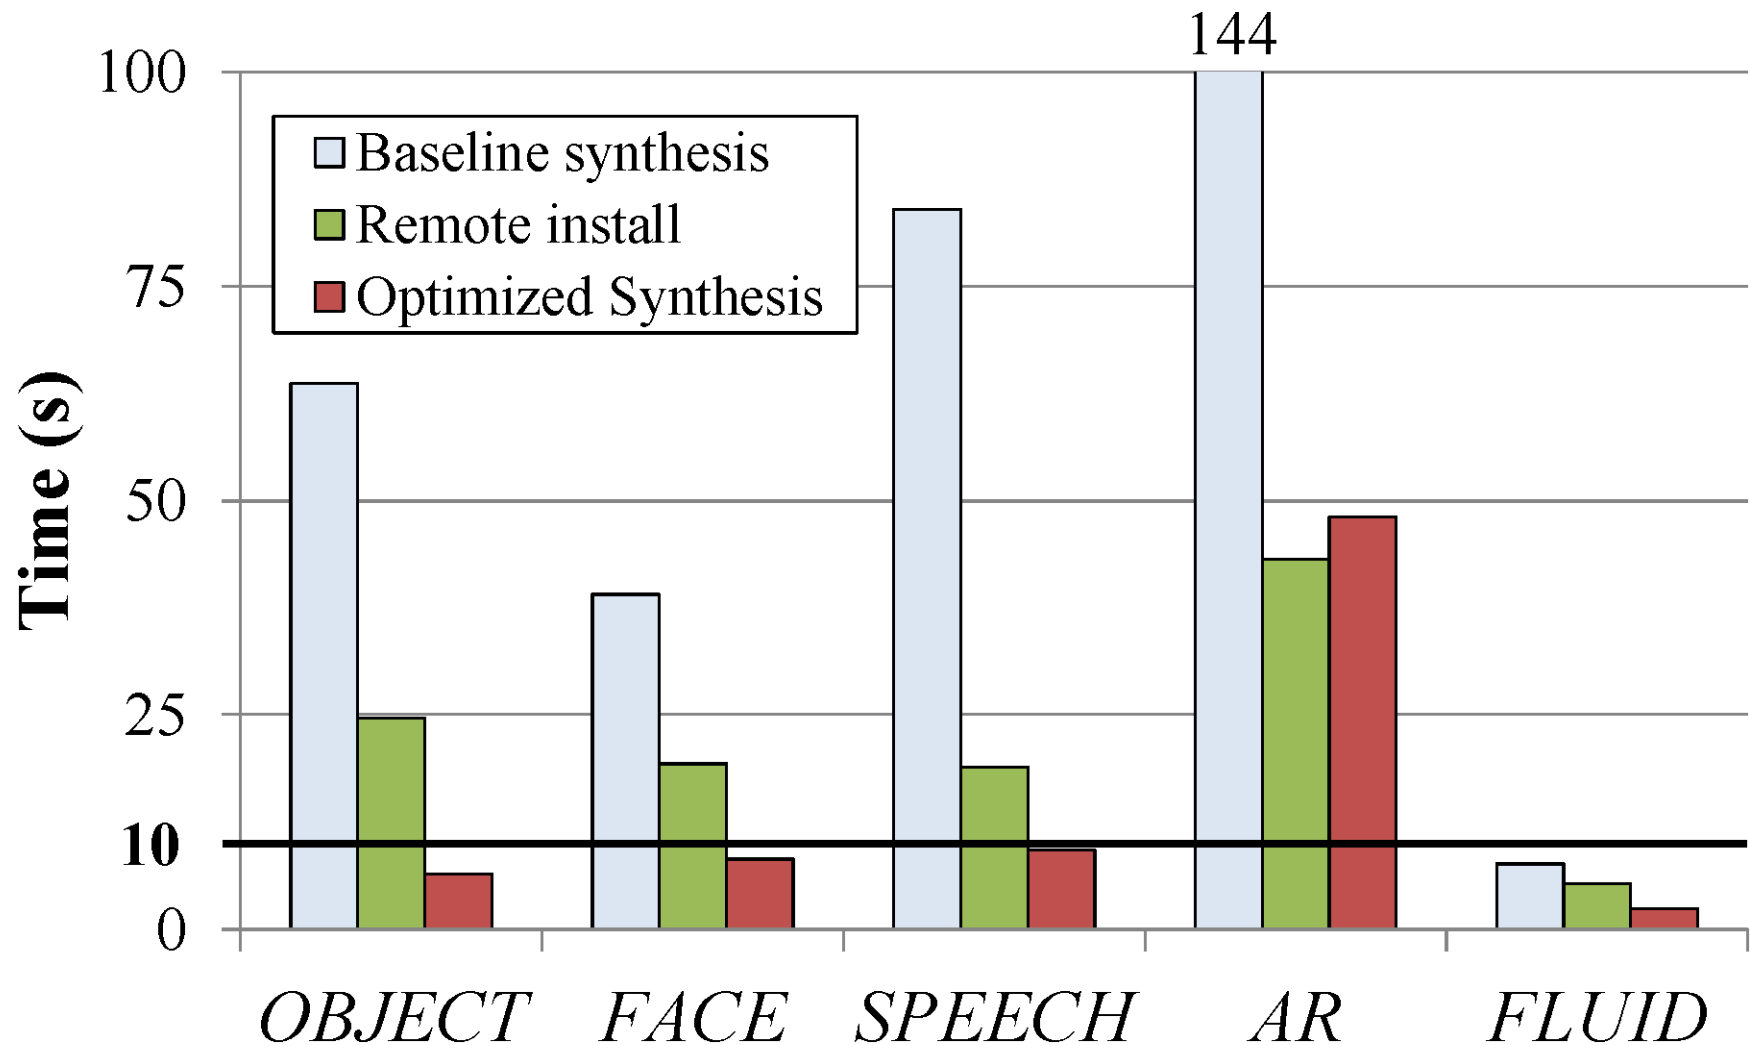
Goal is to minimize “time to value”

- (Overlay transfer + decompression + applying delta + VM resume)
- excludes time for association with cloudlet
- includes time to produce first result (1-2 minutes)

4 optimizations

1. *Deduplication*
 2. *Bridging semantic gaps*
 3. *Pipelining*
 4. *Early Start*
- } Reduce overlay size
(at overlay creation time)
- } Expedite synthesis
(at runtime)

Optimizations Win Big!



3. Can Cloudlets Reduce Bandwidth Demand?

“Scalable Crowd-Sourcing of Video from Mobile Devices”

Simoens, P., Xiao, Y., Pillai, P., Chen, Z., Ha, K., Satyanarayanan, M.

Proceedings of the Eleventh International Conference on Mobile Computing Systems, Applications and Services (MobiSys 2013), Taipei, Taiwan, June 2013

Effortless Video Capture

Google Glass Hits Runway At New York Fashion Week (PHOTOS)



CARL FRANZEN - SEPTEMBER 10, 2012, 11:28 AM | 3571

37

Tweet

17

Like

20

Updated 12:16 p.m. EDT, Monday, September 10

Google's experimental hi-tech glasses, "Google Glass," made an **unexpected appearance at New York Fashion Week** on Sunday when models for designer **Diane Von Furstenberg (DVF)**'s Spring 2013 collection took to the runway wearing the computerized specs, only to be **joined later** by the bespectacled designer herself and Google co-founder and Glass project leader Sergey Brin.

Opportunistic Sensing

New Jersey family's picture catches theft in the making

By **Jamie Guzzardo**, CNN
August 25, 2010 9:49 a.m. EDT

Twitter Share Email Save Print
Recommend 1608 recommendations. Sign Up to see what your friends recommend.



The picture shows a man allegedly taking John Myers' bag as his family gets photographed. Police later arrested a suspect.

STORY HIGHLIGHTS

- A New Jersey man snapped a photo of his family during a trip to Wisconsin
- The photo also caught the image of someone allegedly making off with the man's bag
- Among other valuables, the man's wallet and car keys were in the bag
- Wisconsin Capitol Police quickly apprehended a suspect and returned the bag

Read more on the story on CNN affiliate *WISN-TV* in Madison, Wisconsin.

(CNN) -- In today's technology-laden society, hearing of crimes solved or cold cases cracked with state-of-the-art tech tools has become commonplace. But for one New Jersey family all it took to catch an alleged thief was a camera and a little luck.

John Myers and his family from Bloomfield, New Jersey, were visiting Madison, Wisconsin, to attend a friend's wedding Saturday at the state Capitol.

According to Myers, the family went outside after the ceremony to take pictures. He decided it would be a good idea to take a family

THE MARRIOTT REWARDS® PREMIER CREDIT CARD FROM CHASE

5 points per \$1 spent at all Marriott® locations

Over 3,400 Marriott locations worldwide

Start with 30,000 Marriott Rewards bonus points

LEARN MORE >>

ADVERTISEMENT

NewsPulse >>

Lost Child at Large Public Event



Found!



Challenge: Bandwidth

High cumulative data rate could overwhelm

- metro area networks
- ingress Internet paths into centralized cloud infrastructure

Today, 1 hour of video is uploaded to YouTube each second

- equivalent of only 3,600 users simultaneously streaming
- ***decentralization*** essential for one million users or more

Verizon recently upgraded metro area networks to 100 Gbps link

- one link supports 1080p streams from only 12,000 users
(at YouTube's recommended upload rate of 8.5 Mbps)
- supporting a million users will require 8.5 Tbps
- ***even one YouTube data center per city may be too sparse!***

Challenge: Denaturing


One possible incentive model: like book publishing

(royalties shared by service provider for images of value to others)

Continuous analysis and editing of privacy-sensitive parts of images

- combines meta-data (e.g. location) and content (e.g. specific faces)
- automatic but user-specific, possibly context sensitive
- cannot require manual intervention by user (only exceptions)
(no incentive model can sustain continuous manual user effort)

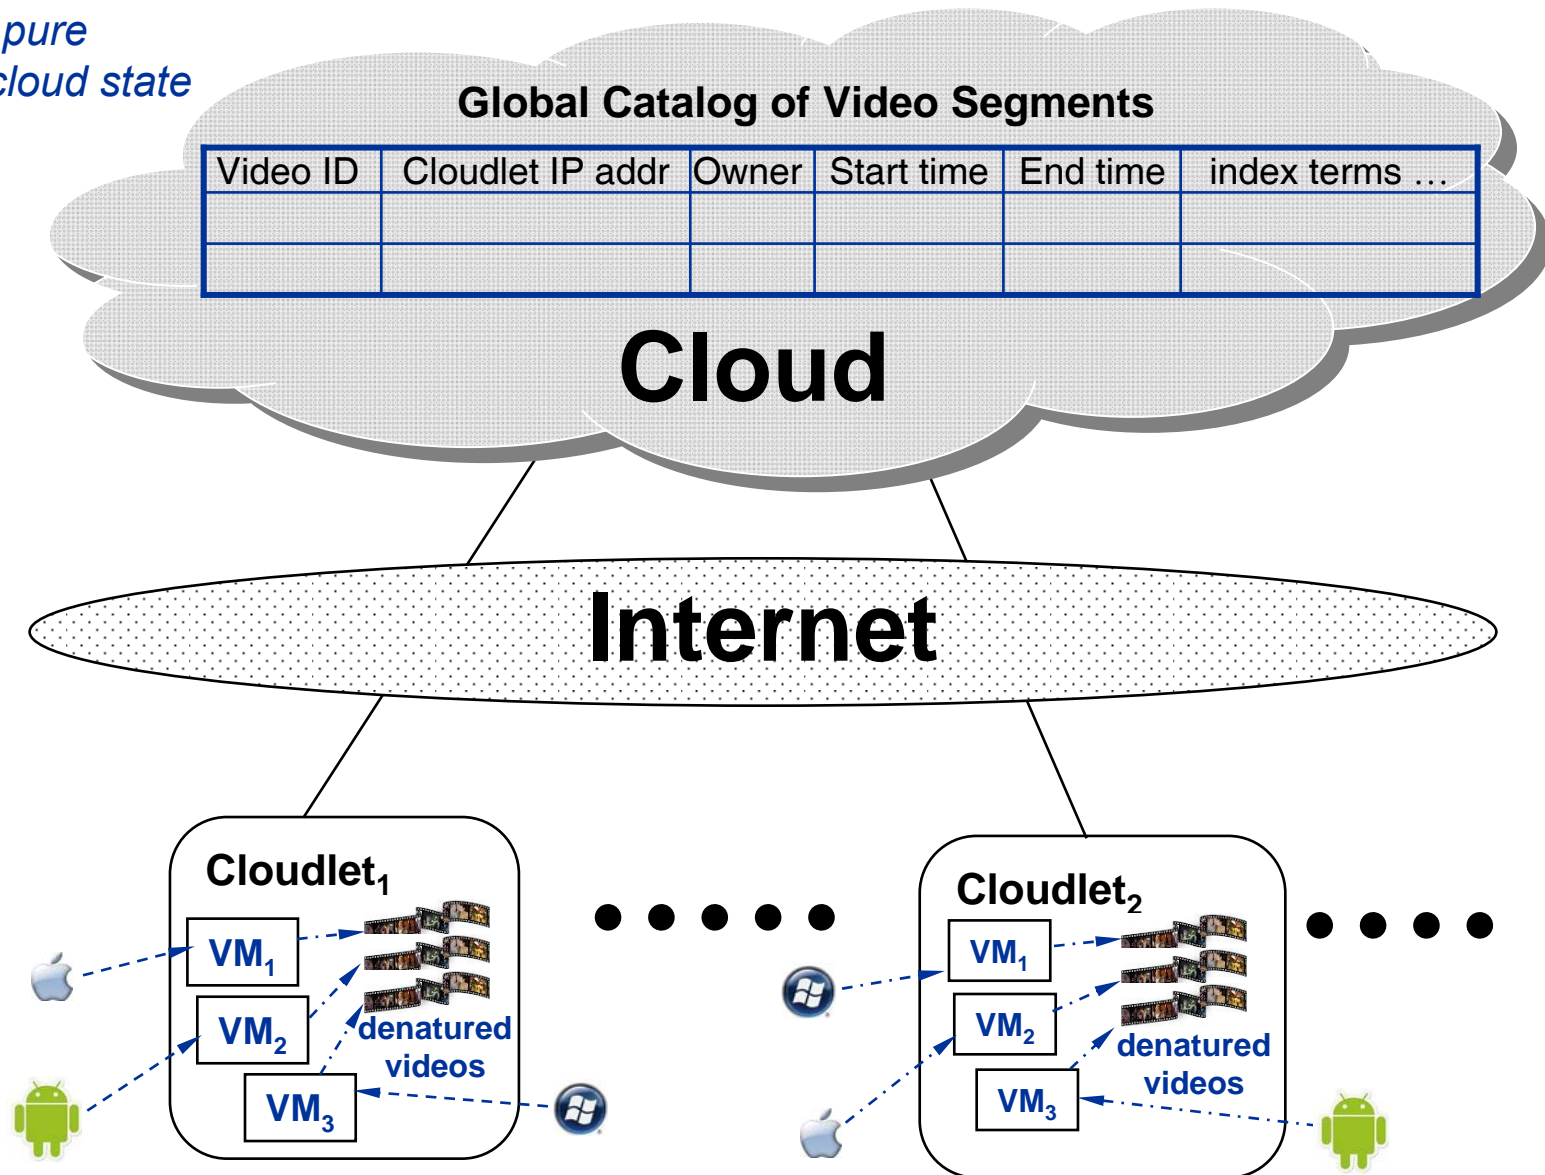
Image processing and denaturing is computationally expensive!




GigaSight Architecture

CDN in Reverse

Cloudlets now contain hard state – no longer pure cache of cloud state



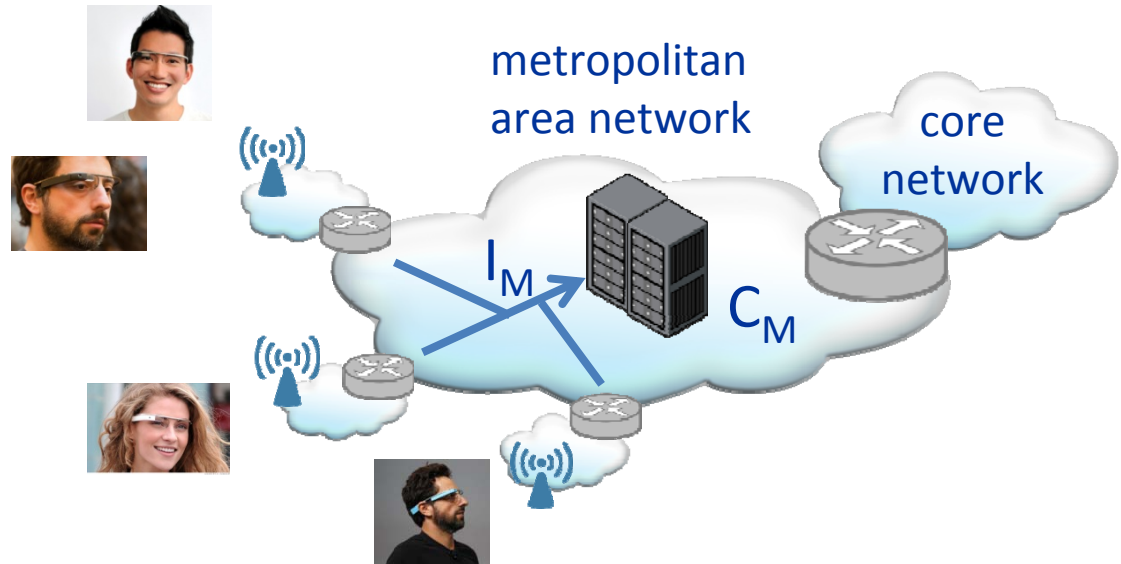
Core versus Edge Deployment



Less scalable, even with future networks



Bigger cloudlet deep inside



Broader Context of GigaSight


Near real-time edge analytics on high-data rate sensors



GE Aircraft Engines

1 TB/engine/day of sensor data

Google Car 750 MB/s



Boeing 787s to create half a terabyte of data per flight, says Virgin Atlantic

Internet of things will create a wide range of opportunities and challenges for airline

By Matthew Finnegan | Computerworld UK | Published 14:27, 06 March 13



The Path Forward

2008

Early ideas on transient customization (Kimberley project at CMU)


“**Transient Customization of Mobile Computing Infrastructure**”, Wolbach et al,
Proc. of the MobiVirt 2008 Workshop on Virtualization in Mobile Computing, June 2008, Breckenridge, CO

Discussions leading to cloudlet concept for futuristic mobile apps

Satya, Victor Bahl (Microsoft), Ramon Caceres (AT&T), Nigel Davies (Lancaster), Roy Want (Intel)



Natural Language Translation




Face Recognition

Year	Computer worse than human (%)	Computer better than human (%)	Indeterminate (%)	Worse/Better
1999	87.5	4.2	8.3	21.0
2001	87.5	8.3	4.2	10.5
2003	45.8	16.7	37.5	2.75
2005	37.5	33.3	29.2	1.13
2006	29.2	37.5	33.3	0.78

2009

“**The Case for VM-based Cloudlets in Mobile Computing**” published in October 2009

Very well received, over 370 citations in 4 years according to Google Scholar!



2010-
2013

Outside
Carnegie
Mellon

Research community embraces cloudlet concept

(many publications, partial list below)

- **“Cloudlet-screen computing: A multi-core-based, cloud-computing-oriented, traditional-computing-compatible parallel computing paradigm for the masses”**
Lin, T., IEEE International Conference on Multimedia and Expo (ICME 2009). July 2009
 - **“Impact of Cloudlets on Interactive Mobile Cloud Applications”**
Fesehaye et al, IEEE 16th International Enterprise Distributed Object Computing Conference, 2012 [Univ. of Illinois]
 - **“Cloud-Vision: Real-time face recognition using a mobile-cloudlet-cloud acceleration architecture”**
Soyata, T., IEEE Symposium on Computers and Communication, (ISCC) July 2012
 - **“Experimental framework of integrated cloudlets and wireless mesh networks”**
Khan et al, Telecommunications Forum (TELFOR), Nov. 2012
 - **“Tactical Cloudlets”**
Keynote by Dr. Raju Namjuru, Chief of Computational Sciences, US Army Research Laboratory, IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, Washington D.C., August 2012
 - **“Optimal Admission Control Policy for Mobile Cloud Computing Hotspot with Cloudlet”**
Hoang et al, IEEE Wireless Communications and Networking Conference (WCNC), April 2012
 - **“Adaptive Application Configuration and Distribution in Mobile Cloudlet Middleware”**
Verbelen et al, Mobile Wireless Middleware, Operating Systems, and Applications, Springer Lecture Notes, Volume 65, 2013
[Ghent Univ.]
- ... *many others* ...

Cloudlets now a mainstream concept
in the research community

Industry Investment Has Begun

NETWORKWORLD

News | Blogs | Newsletters | Videos | Events | **INSIDER** | More








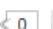

Security | LANs & WANs | UC / VoIP | Cloud | Infrastructure Mgmt | Wireless | Software | Data Center | SMB |
Broadband | Collaboration | Equipment | Mobile | Networks | Security | Storage | White Papers | Webcasts | Tests

News

Nokia Siemens to merge cloud, base-station computing to boost performance

The company's Liquid Applications platform will use computing power in the cloud and in base stations, based on conditions

By *Stephen Lawson*, IDG News Service
February 24, 2013 04:06 PM ET

 Add a comment  Print  Share   +1   Like  0  More

IDG News Service - Nokia Siemens Networks will expand the role of cellular base stations with a new platform that will store and deliver some application data locally, while tapping into information about subscribers and traffic to improve the process.

The company announced the system, called Liquid Applications, at an event in Barcelona on the eve of Mobile World Congress. Liquid Applications can improve consumers' mobile experience but cutting delays as well as delivering more relevant content, CEO Rajeev Suri



The screenshot shows the IBM News room page for the announcement of the world's first mobile edge computing platform. The page features a navigation menu with options like Home, Solutions, Services, Products, Support & downloads, and My IBM. The main content area includes a breadcrumb trail (News room > News releases > IBM and Nokia Siemens Networks Announce World's First Mobile Edge Computing Platform) and a news release summary. The summary states that on February 25, 2013, at the Mobile World Congress in Barcelona, Spain, Nokia Siemens Networks and IBM (NYSE: IBM) announced a collaboration to deliver the world's first mobile edge computing platform. This platform allows mobile operators to create a truly unique mobile experience, relieve the ever-increasing strain on network infrastructure, and bring completely new solutions to market. The new platform can accelerate the delivery of media-rich services by delivering content directly from the base station, ensuring enhanced quality of experience for consumers in the face of ever-increasing data traffic growth. The platform also enables a new generation of

cloudlet-based apps & services



early validation apps

An Open Cloudlet Ecosystem

open cloudlet API

early implementations



cloudlet infrastructure