Cloudlet-based Wearable Cognitive Assistance

Mahadev Satyanarayanan (PI), Martial Hebert (co-PI), Dan Siewiorek (co-PI)
Yoshihisa Abe, Zhuo Chen, Da-Yoon Chung, Benjamin Gilbert, Kiryong Ha, Jan Harkes
Wenlu Hu, Ishan Misra, Padmanabhan Pillai†, Wolfgang Richter, Brandon Taylor (CMU, †Intel)

GOALS

- Let smart glasses (e.g. Google Glass) help human users be smarter
  - Sensors (especially camera) sense the world continuously, and only alert people when necessary
  - Leverages real-time computer vision
- Helpful particularly for
  - Aging people and people with brain injuries
  - Severely attention-challenged people (e.g. in an aircraft cockpit)

CLOUDLET-BASED APPROACH

- *Motivation:*
  - Computer vision tasks are resource-intensive; hard to work in real-time on mobile devices
  - Large latency between mobile devices and clouds hurts user experience in cognitive assistance
- Cloudlet ≅ "data center in a box"
  - Compute cluster distributed at the edge of Internet
  - Low end-to-end latency, high bandwidth
  - Powerful, well-connected and safe
  - Only soft state

SYSTEM ARCHITECTURE

- VM and PubSub Backbone based architecture supports diverse software building blocks
  - Response time for system infrastructure is ~50 ms
  - Context inference enables smart sensor control for energy saving and better user experience

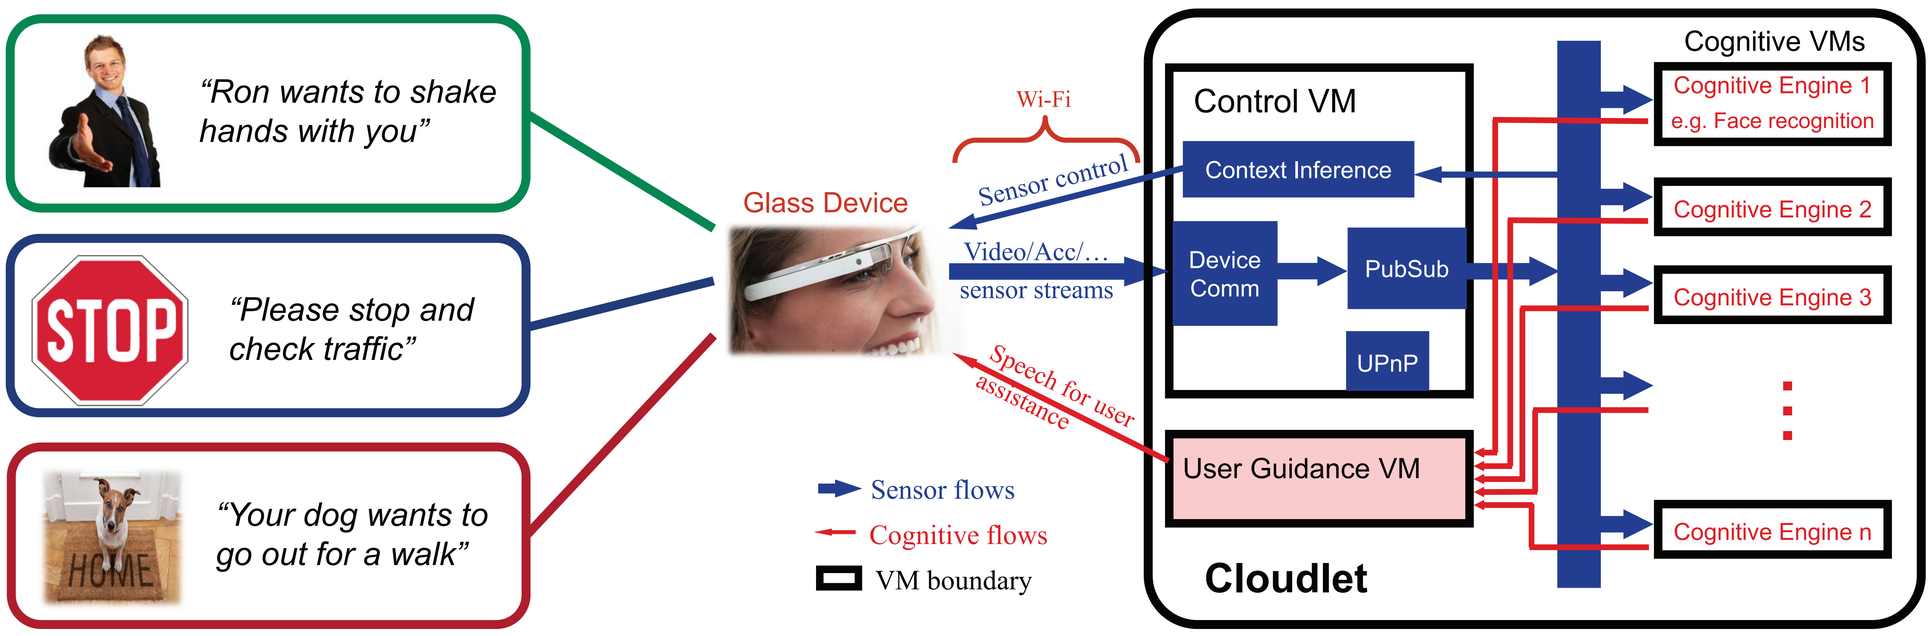

EXPERIMENTAL RESULTS AND FUTURE WORK

- Current prototype consists of seven cognitive engines
  - Running Windows/Linux; Developed with C, C++, or Python
  - Cloudlet consists of 4 advance desktop machines, running OpenStack

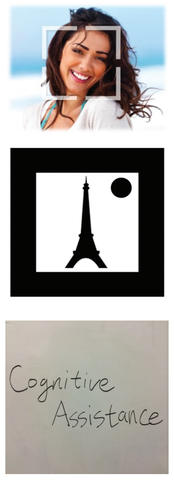

Cognitive Engine	FPS	Response time (ms)					Glass Life
		1%	10%	50%	90%	99%	
Face Recognition	4.4	196	389	659	929	1175	
Object (MOPED)	1.6	877	962	1207	1647	2118	
Object (STF)	0.4	4202	4371	4609	5055	5684	
OCR (Open)	14.4	29	41	87	147	511	~1 hour
OCR (Comm)	2.3	394	435	522	653	1021	
Motion Classifier	14.0	126	152	199	260	649	
Augmented Reality	14.1	48	72	126	192	498	

Speed improvement in detection and model creation

Longer battery, better thermal dissipation

Cloudlet provisioning and VM migration

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · Princeton University · UC Berkeley · University of Washington