Design of A Low Power SoC Testchip for Wearables and IoTs

May Wu, Ravi Iyer, Yatin Hoskote, Steven Zhang, Julio Zamora, German Fabila,

Ilya Klotchkov, Mukesh Bhartiya

August, 2015

Legal Notices and Disclaimers

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document..

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Wearables and IoT: Redefining computing

50 Billion

CISCO

>50 Billion

ERICSSON

200 Billion

IDC

75 Billion

Morgan Stanley

SIGHT and SOUND:

The New Frontiers

INTEGRATED INTELLIGENCE:

Always Listening. Always Watching.

We are limited by high power consumption

Power is the greatest barrier for intelligent “always on” devices

Power consumption is nowhere near where we need to be

1~2 Hours
Battery Life

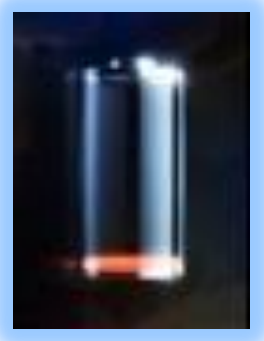

Wearable Device with A/V Capabilities Today	
Camera	~100mW
Processing	200~400mW
Transmission	200~250mW
Storage	Tens of mW
Display	Tens~Hundreds of mW

Lower MPixels Type

Looxcie measure Estimate

Best WiFi on publication

Rethinking enabling technologies for *always on* IoT and Wearable devices.

Better Ways for Data Capturing and Transmission

- Pushing intelligence close to sensing side
- Threshold / Buffer
- Activity triggers
- Context-Ware Encoding
- Adaptive and cooperative communication

Intelligent IoT and Wearables Demand a Better SoC

Intel Labs:
Testchip for Always-On Devices

Introduced in Q1 of 2015, undergoing further development

14nm Intel Process

From ~2mW keyword recognition to few 10s mW A/V processing

Key Features:

Testchip for Always-On Devices

Always-Watching: with a Vision Processing Engine (e.g. Gesture, Scene Detect)

Always-Listening: Voice Activity Detect; Short phrase(s) Recognition; Speaker ID

Low Power Embedded Communication

Light-Weight Security Framework and Processing

High Level Architecture Overview of the Testchip

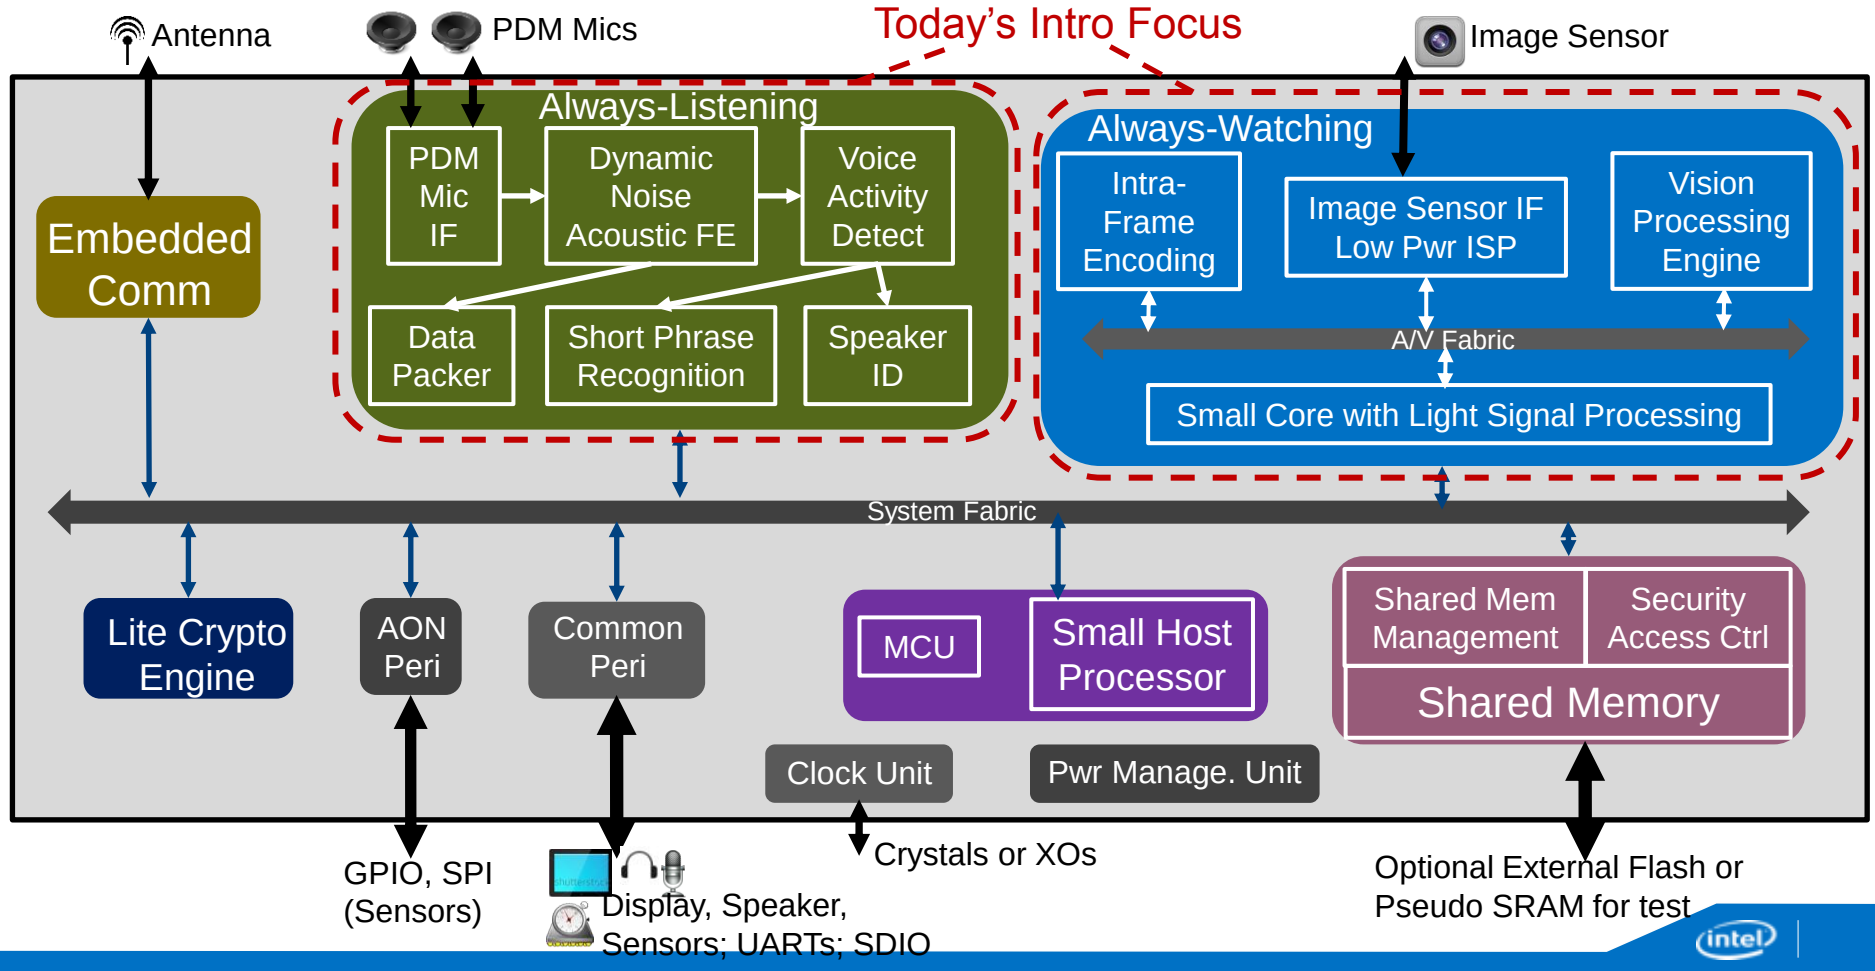

Our Design Strategy for “Always-Watching” Devices

Two Key Advances

1. **Vision-Driven LP Imaging**
   - Very aggressive image sensor power gating
   - Race to halt
   - Light-detect assisted auto-exposure processing
   - Intra-frame and data analysis driven encoding

2. **Optimized Common Neural Network Processing for Multiple Applications**
   - Shifted Neural Network for the classification
     - Shift operations, fixed point, approx sigmold/hyberbolic tangent functions, etc
   - Memory optimized convolutional layers for vision recognition feature extraction

Always-Watching – Multiple Applications Common Solutions

Application

Tailored Preprocessing

Sobel + Convolutional layers

Haar or convolutional layers

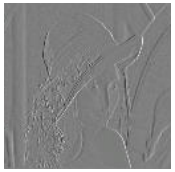

Otsu + Convolutional layers

HOG

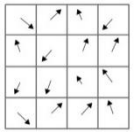

Neural Network

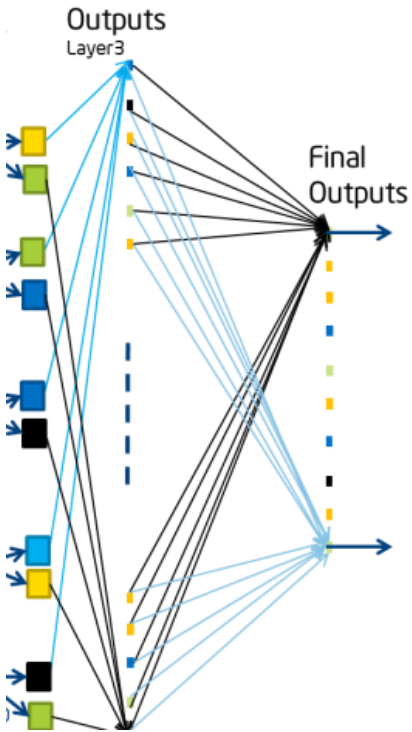

One NN engine shared
Changing weights and topology

Output

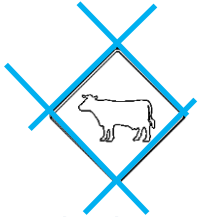

Animal crossing

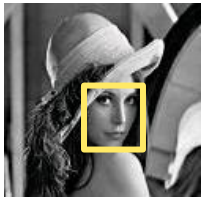

Face at (50,30)

Text:"7"
(Input)

Action: "Close Application"
(Command)

Always-Watching - Vision Processing

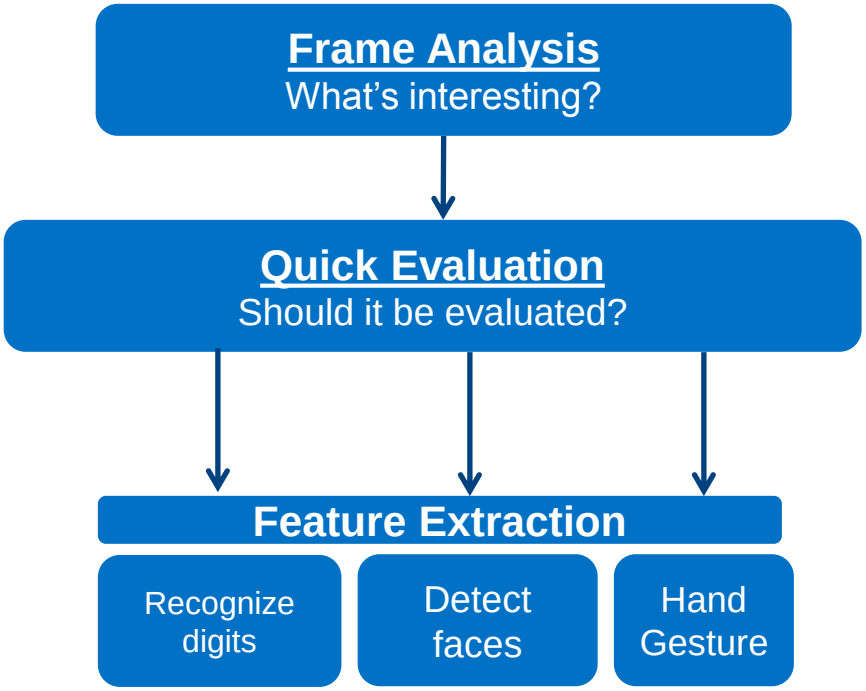

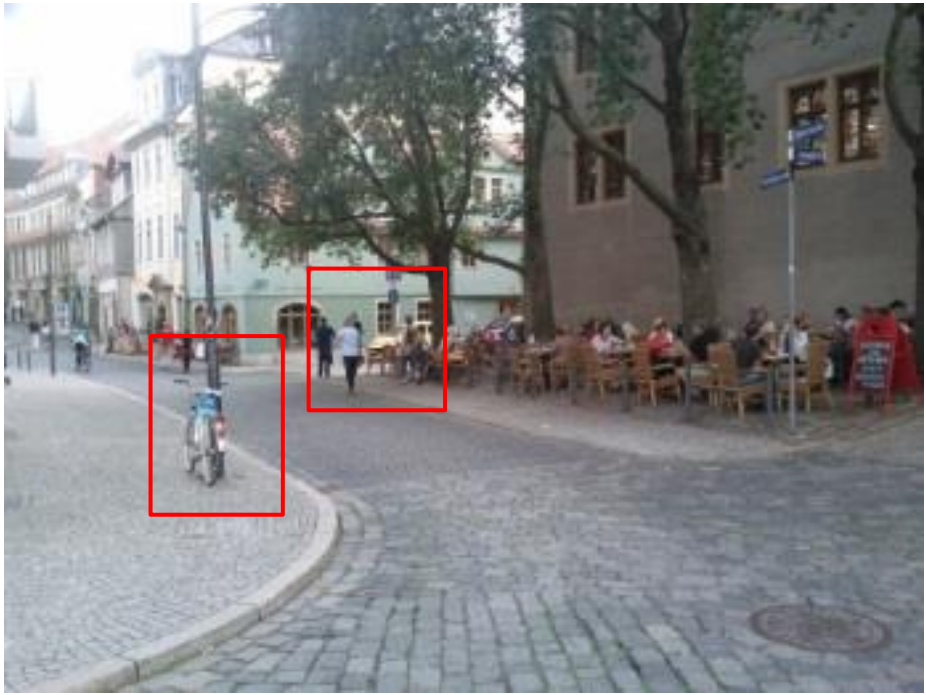

Identify what is interesting to reduce NN evaluations

Always-Watching - SubSystem Overview

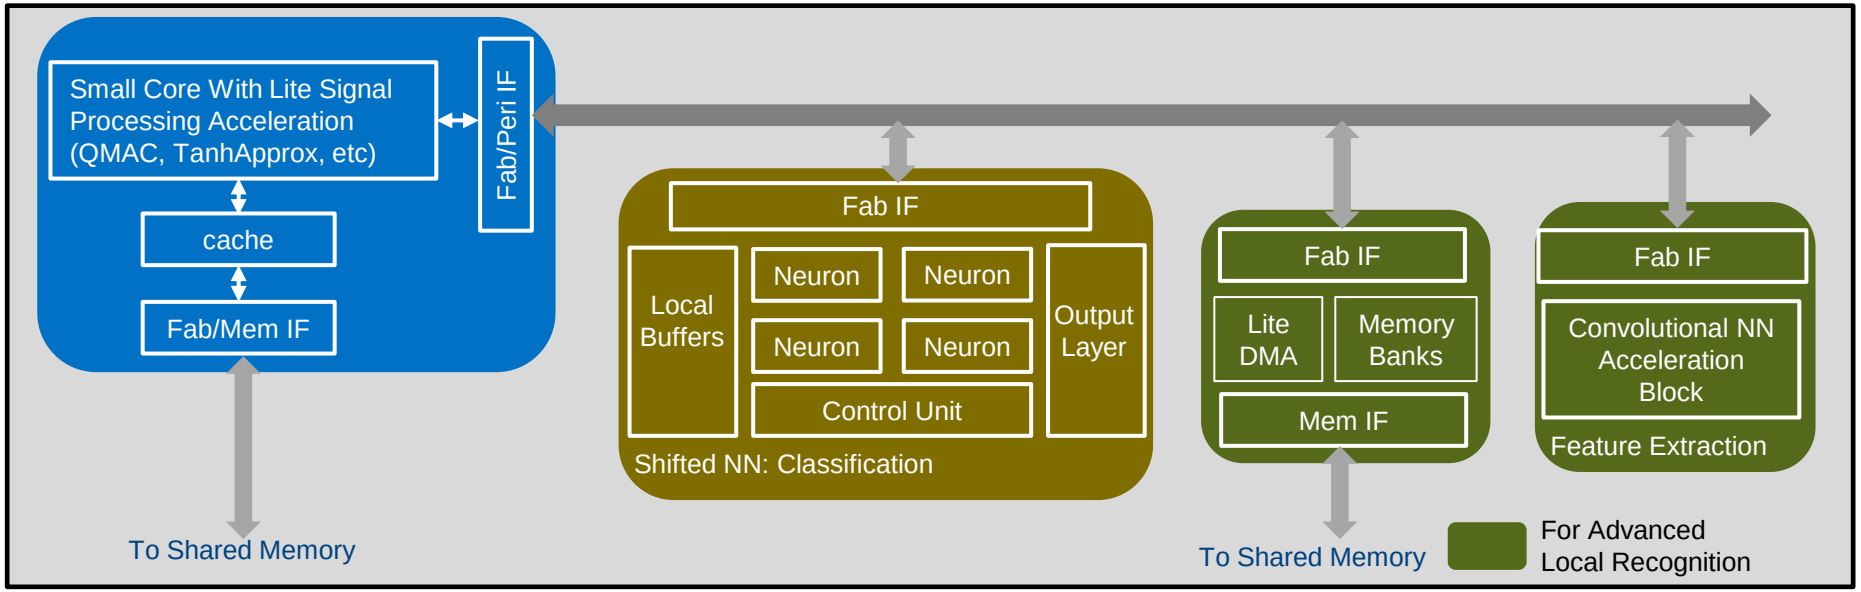

Processing Phase	Processing Element
Frame Analysis and Segmentation	Small Core with Lite Signal Processing Acceleration
Quick Evaluation	Small Core with Lite Signal Processing Acceleration
Feature Extraction	• Small core with Lite Signal Processing Acceleration • CNN Acceleration IPs for feature extraction
Classification	Highly optimized Shifted NN

Always-Watching Vision Processing: Hand Gesture Experiment

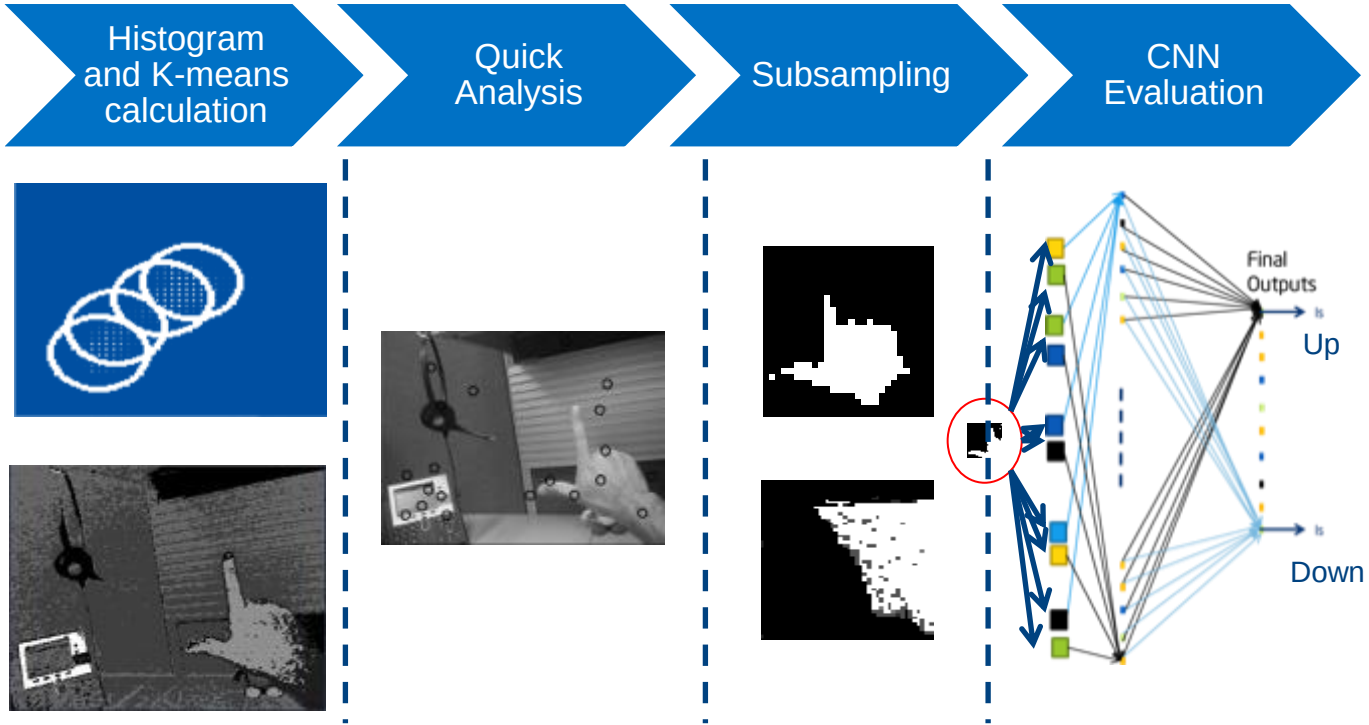

2 fps or lower; QVGA or lower; YUV; Distance 20cm~1m;
Response time 200ms or lower; Recognition processing power <1mW to several mW

Always Listening Speech Processing Pipeline: Design for Power Reduction

End-to-End H/SW Partitioning with Low power Always-Listening on Device

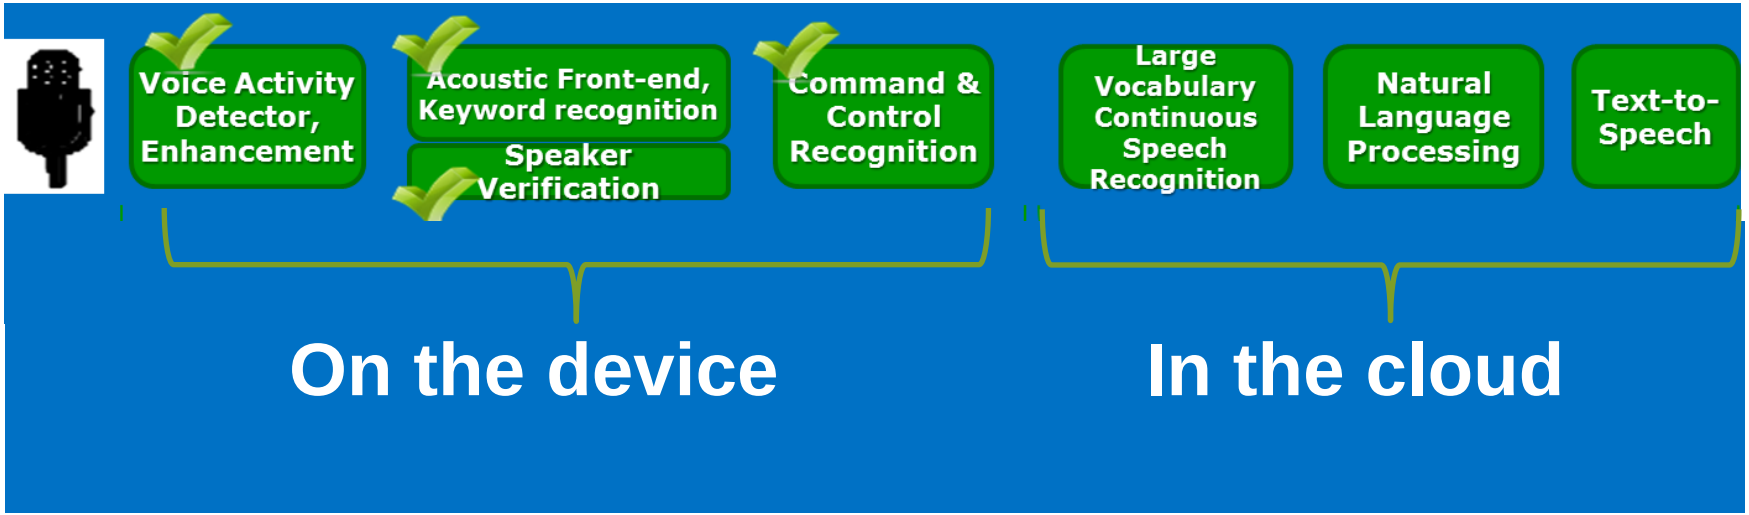

Always Listening Block Overview

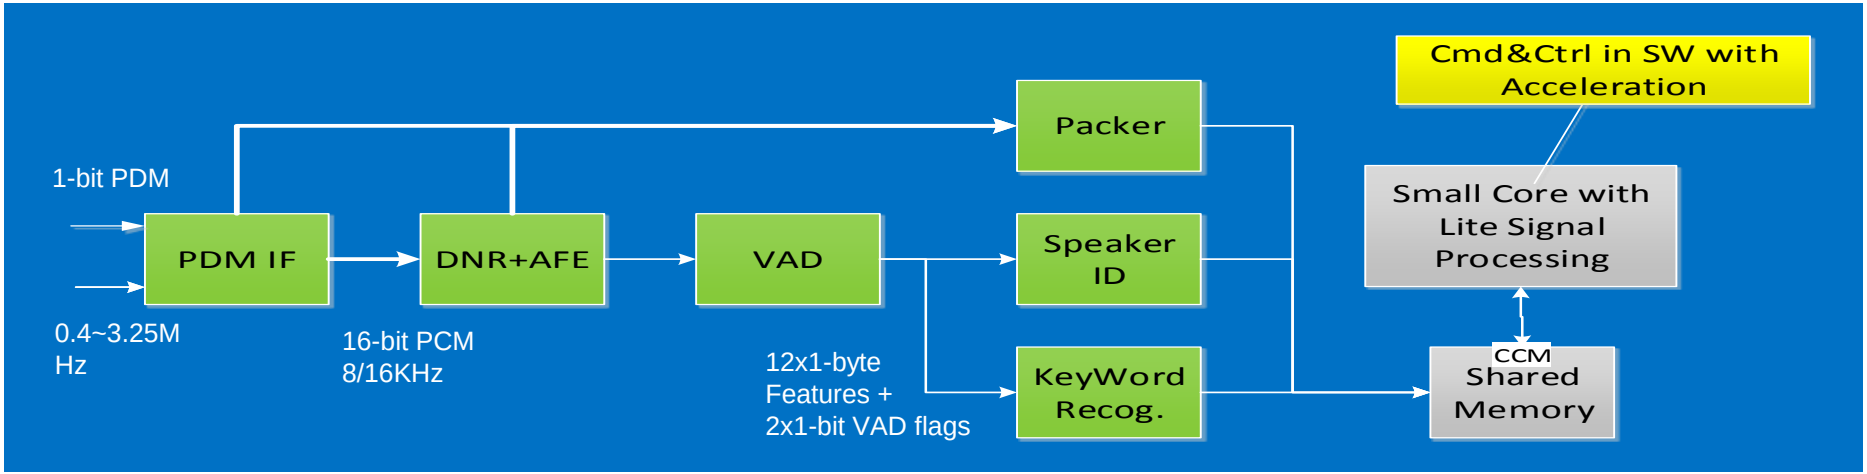

Key Advances

- **Noise reduction tailored for speech recognition**
- **0.1~0.2mW DNR+AFE+VAD**
- **single digit mW for tens of command and control recognition, with accurate VAD support**

- **2 audio channels** can be on/off independently
- **Audio sampled** as 16-bit@16kHz/8KHz for voice activity detect and short phrase recognition
- **Acoustic Front End and Voice Activity Detect process** 1 audio channel in 160-sample frames, 100 frames/s, and produce frames of 12 features + voice&quiet flags

The Testchip On Intel 14nm

4mmX8mm
Shared Die
(low utilization)

Embedded Comm
(Including RF and
Other Analog
Circuit

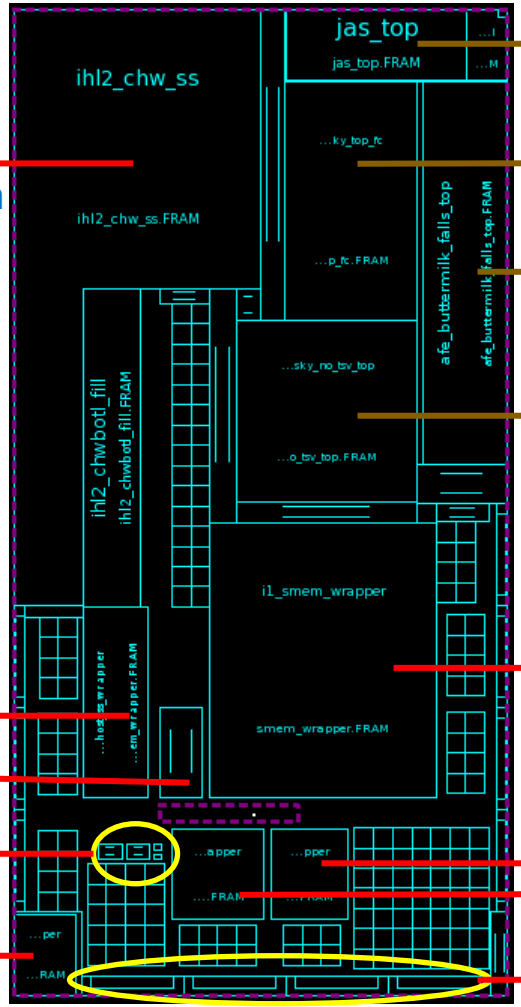

Other
testchip
Projects
shared
the same
Die
Area

Host Core
SubSystem

Audio/Speech

MCU & macros

Clock/
Ring Oscillator

Shared MEM

Vision&
Small core with
signal processing

Io family

Test 1: Always-Listening VAD and Keyword Recognition

	Voice Activity Detect Stage	Keyword Recognition Stage
With 1 digital Mic at low performance mode (customized Mic)	~1mW (including Mic)	~1.9mW (including Mic)
With 1 digital Mic at standard performance mode (regular Mic)	~1.5mW (including Mic)	~2.5mW (including Mic)

Test 2: On Chip ~22mW A/V capturing, Hand posture and Always-listening

Functions	Pwr (mW)
Audio/Speech	0.9~0.946
Imaging	1.65~3.322
Vision Recognition	5.478~5.566
Host Core with Memory	5.082~7.04
Shared Memory	3.586~3.674
Fabric&Peri	2.002

Experiment Platform and Usage Examples

Gesture-Based Control

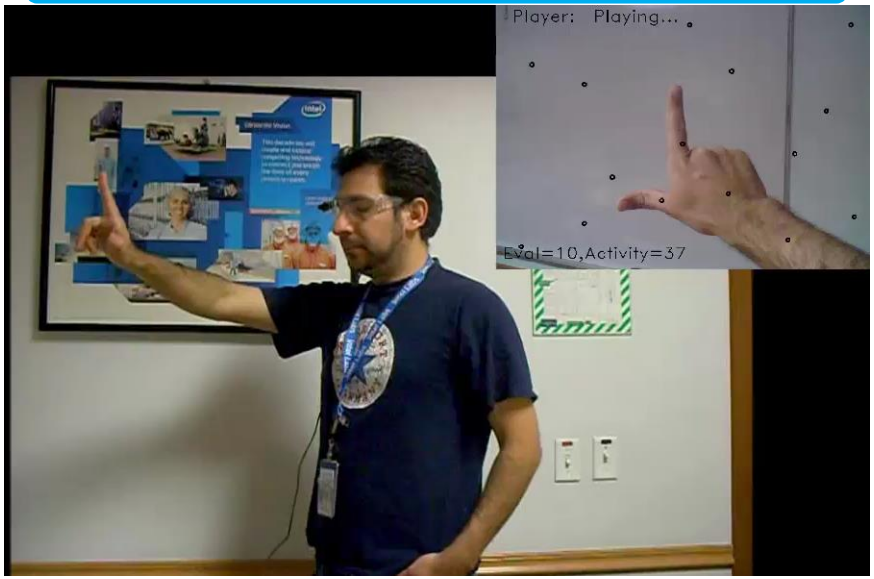

The Testchip Form Factor Test Board

Digit Recognition Example

Speech /Vision

Integrated Pipeline

Gesture Recognition Flow

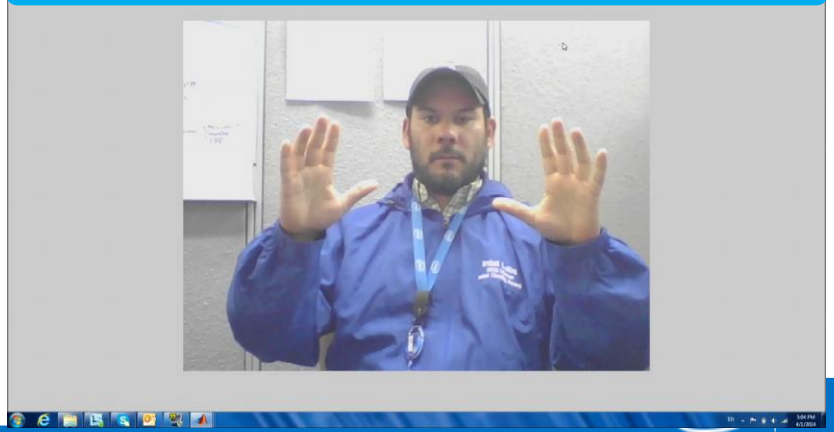

Immediate Next steps & Longer Term Directions

- Usages trend to require systems to make “human-like” decisions (bots, drones, kids play,…)
- Adaptive vision+speech+sensor capabilities for ULP recognition & understanding (VU/SU)
- Autonomous radio technologies (ULP wideband radio for sensing, wake-up radios, etc

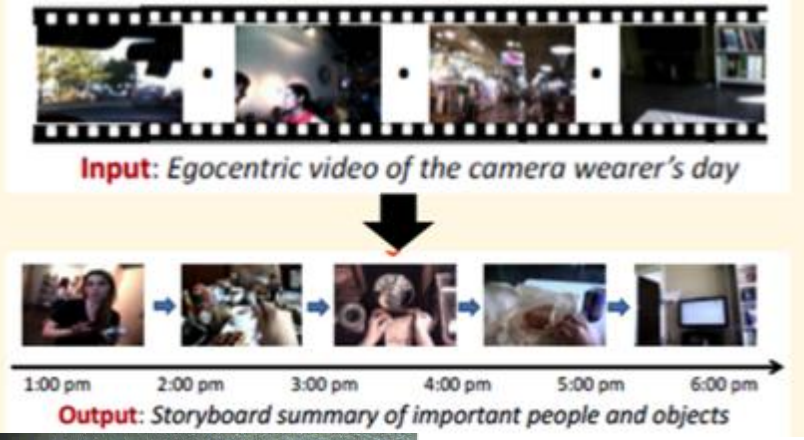

Drive the Always-On Revolution

IoT and Wearable Usage Tailored Power Efficiency

Reducing data transmitted

New SoC to open unprecedented drops in power consumption


intel