



Computing in Memory: Speech Recognition

Spanansion, Nuance, and Objective Analysis
Session 302-B

Introduction

- Industrial design and UI are driving adoption of new technology products
- Natural HMIs involve extremely complex systems
- Speech is the latest HMI being adopted

What are the challenges and solutions being explored in improving speech recognition and Natural Language Understanding?



Industry Experts

- Gabi Artzi, Solutions Architect, Automotive Business Unit, Nuance
- Rich Fastow, Director, Device Technology Engineering, Spansion
- Jim Handy, Director and Chief Analyst, Objective Analysis

Moderator: Stephan Rosner, VP, Software and Systems, Spansion



Flash Memory Summit

[Session 302-B](#)

Computing in Memory: Speech Recognition

Gabi Artzi, Solutions Architect

August 2012

Nuance is Innovating in Major Consumer Markets

Cars, Phones, eReaders, GPS, Computers, Copiers, ...



A Tipping Point?

**Developer
Community**

**Computational
Resources**

**NLU
Innovations**

**Algorithmic
Advances**

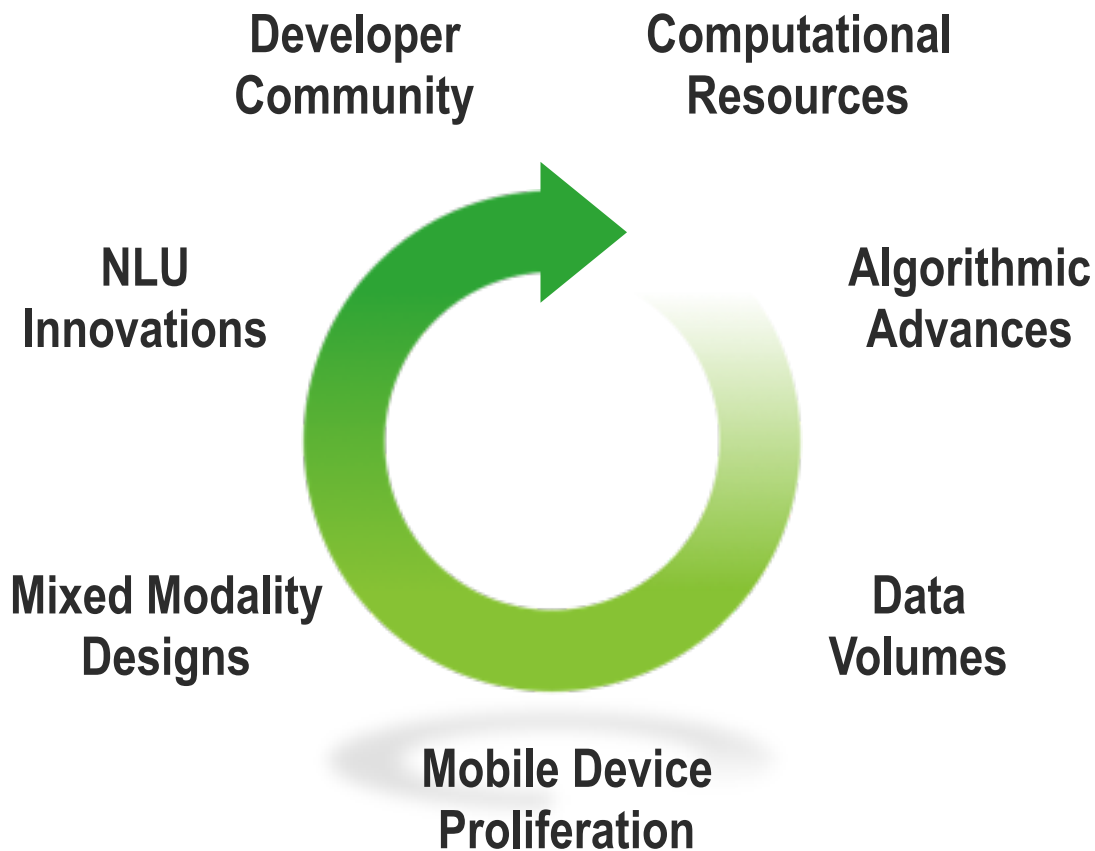
**Mixed Modality
Designs**

**Data
Volumes**

**Mobile Device
Proliferation**

- 2 billion mobile transactions annually
- 6 billion handsets deployed
- 70 million cars shipped
- 10 thousand mobile app developers
- 500 thousand physicians
- 51 languages supported

A Tipping Point?



- 2 billion mobile transactions annually
- 6 billion handsets deployed
- 70 million cars shipped
- 10 thousand mobile app developers
- 500 thousand physicians
- 51 languages supported

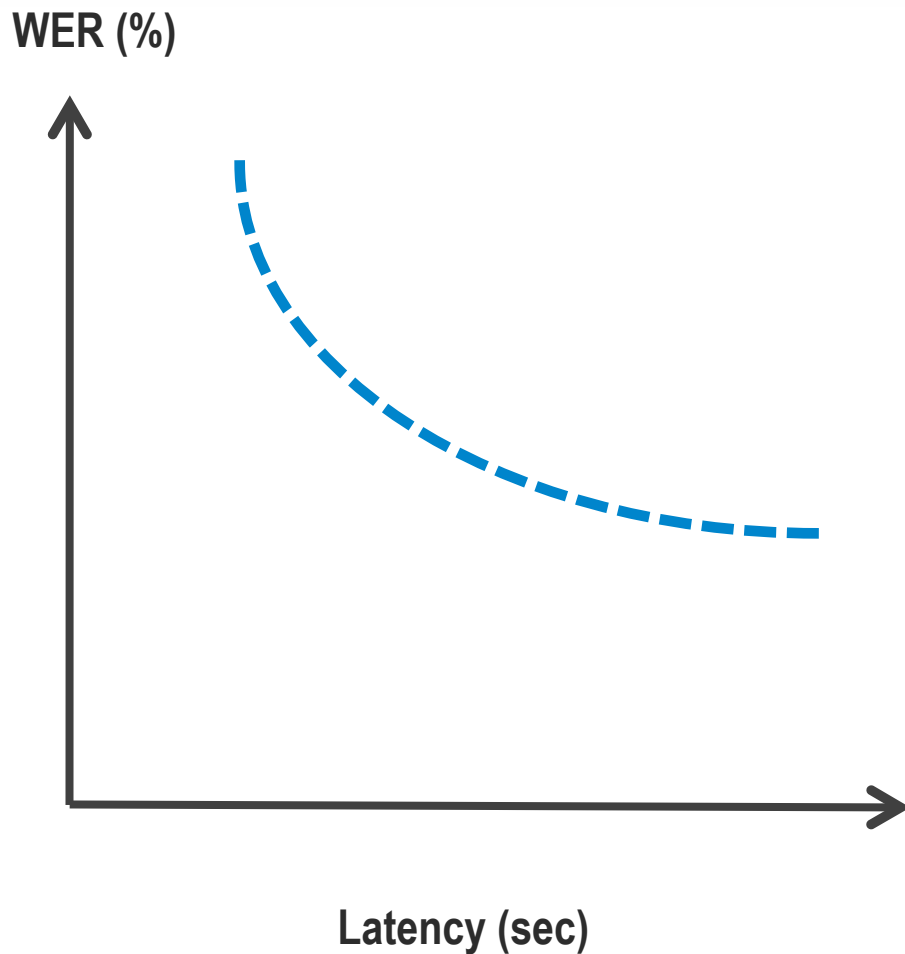
Why an Acoustic Coprocessor?

- Offload CPU
- Increase Acoustic Models Size
 - Drives Higher Recognition Accuracy
- Reduce RAM consumption
- Reduce Speech Recognition Latency
 - Drives Better User Experience



Speech Recognition Quality Metrics & User Experience


(Experiments done by Spansion)



- Quality of speech recognition solutions is determined by
 - Accuracy = Word Error Rate (WER)
 - Latency = system responsiveness
- Today's Systems use "Compact" models
 - Optimized for embedded systems
 - Trade-off between Accuracy and latency
- Next-Generation "Full" models
 - Needed for next-generation applications
 - 15-20% reduction in error rate
 - Unacceptable latencies

Speech Recognition Quality Metrics & User Experience

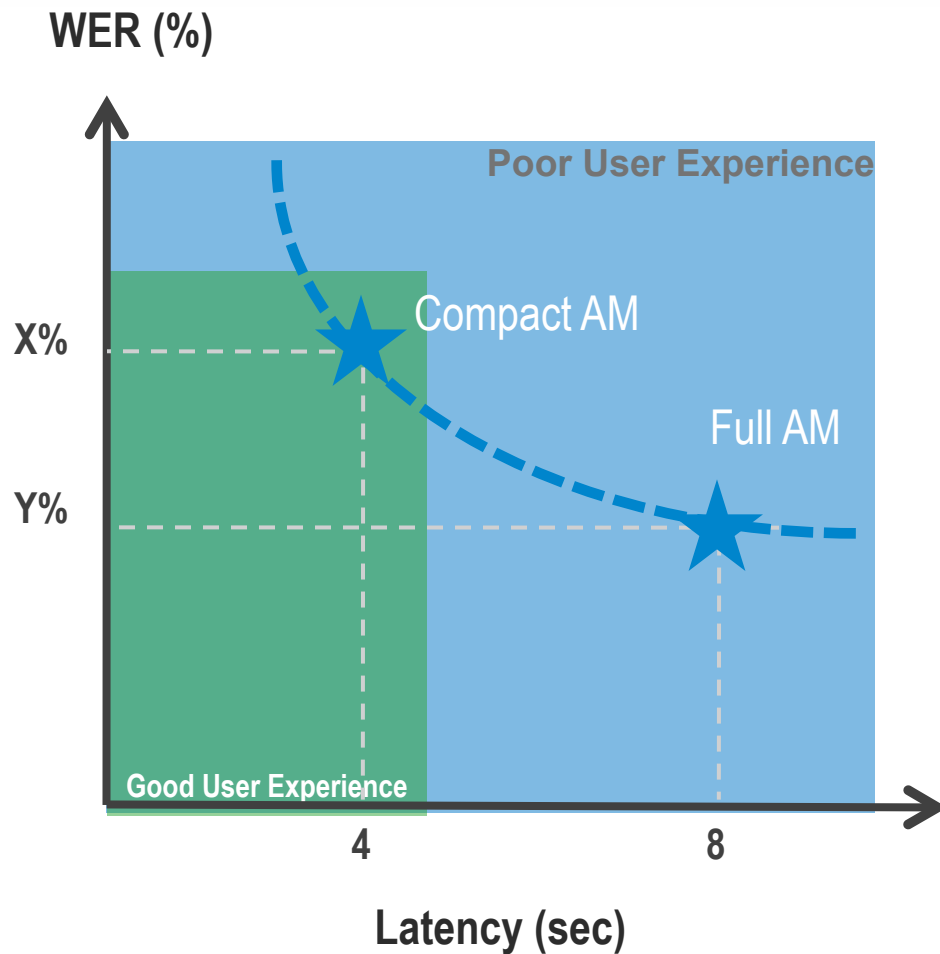
(Experiments done by Spansion)



- Quality of speech recognition solutions is determined by
 - Accuracy = Word Error Rate (WER)
 - Latency = system responsiveness
- Today's Systems use "Compact" models
 - Optimized for embedded systems
 - Trade-off between Accuracy and latency
- Next-Generation "Full" models
 - Needed for next-generation applications
 - 15-20% reduction in error rate
 - Unacceptable latencies

Speech Recognition Quality Metrics & User Experience

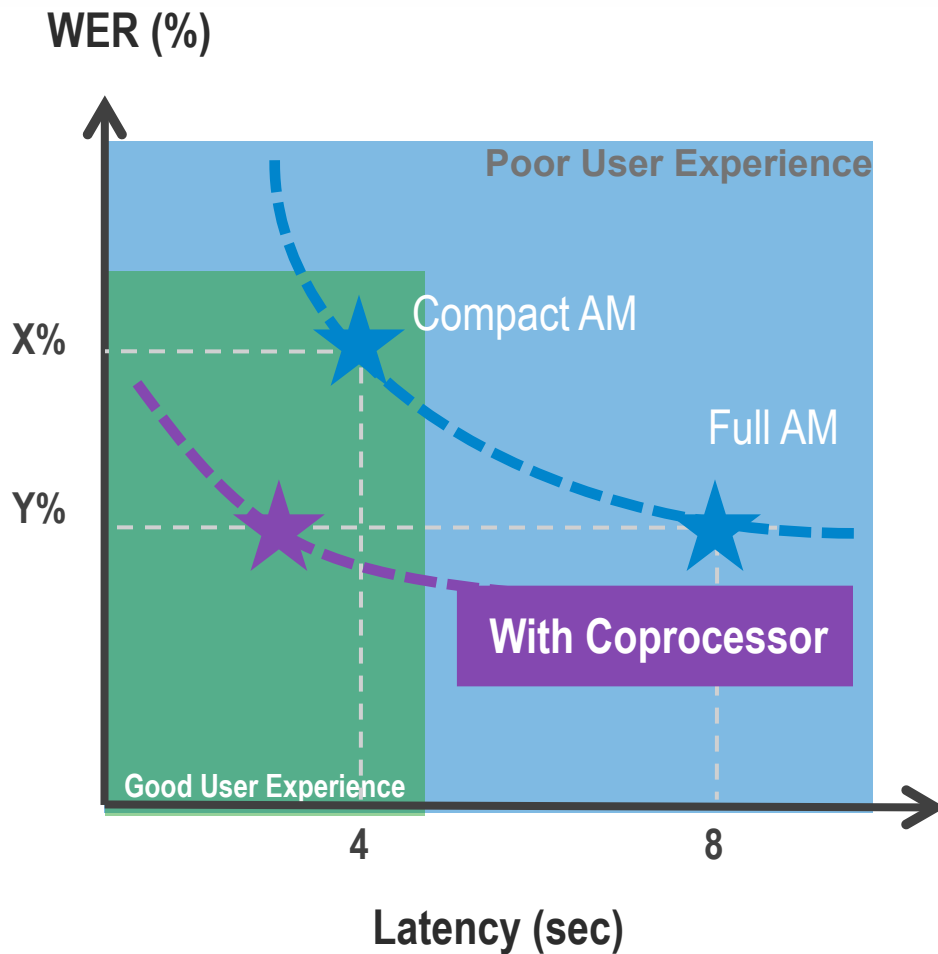
(Experiments done by Spansion)



- Quality of speech recognition solutions is determined by
 - Accuracy = Word Error Rate (WER)
 - Latency = system responsiveness
- Today's Systems use "Compact" models
 - Optimized for embedded systems
 - Trade-off between Accuracy and latency
- Next-Generation "Full" models
 - Needed for next-generation applications
 - 15-20% reduction in error rate
 - Unacceptable latencies

Speech Recognition Quality Metrics & User Experience

(Experiments done by Spansion)




- Quality of speech recognition solutions is determined by
 - Accuracy = Word Error Rate (WER)
 - Latency = system responsiveness
- Today's Systems use "Compact" models
 - Optimized for embedded systems
 - Trade-off between Accuracy and latency
- Next-Generation "Full" models
 - Needed for next-generation applications
 - 15-20% reduction in error rate
 - Unacceptable latencies

Spansion Acoustic Coprocessor enables optimal use of Full acoustic models

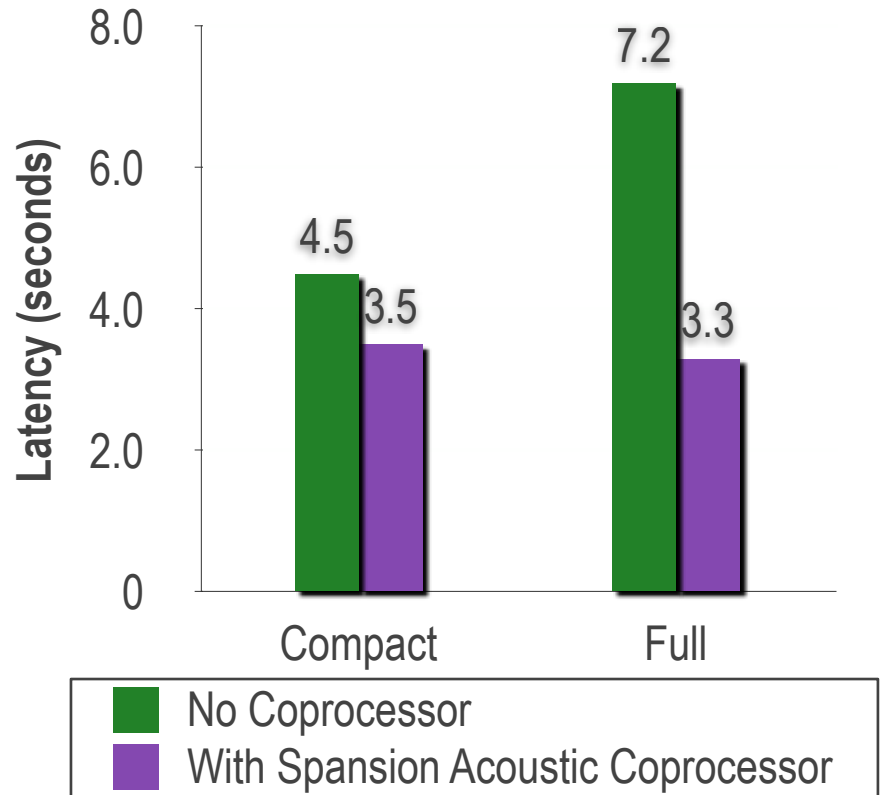
Benchmark results (ENU and GDE 1-shot VDE)

(Experiments done by Spansion)



CPU gain:

- >20% for Compact models
- >40% for Full models



Latency reduction:

- >25% for Compact models
- >50% for Full models



NUANCE



Flash Memory Summit

Computing in Memory: Speech Recognition
SpanSion® Acoustic Coprocessor

Rich Fastow
August 23, 2012

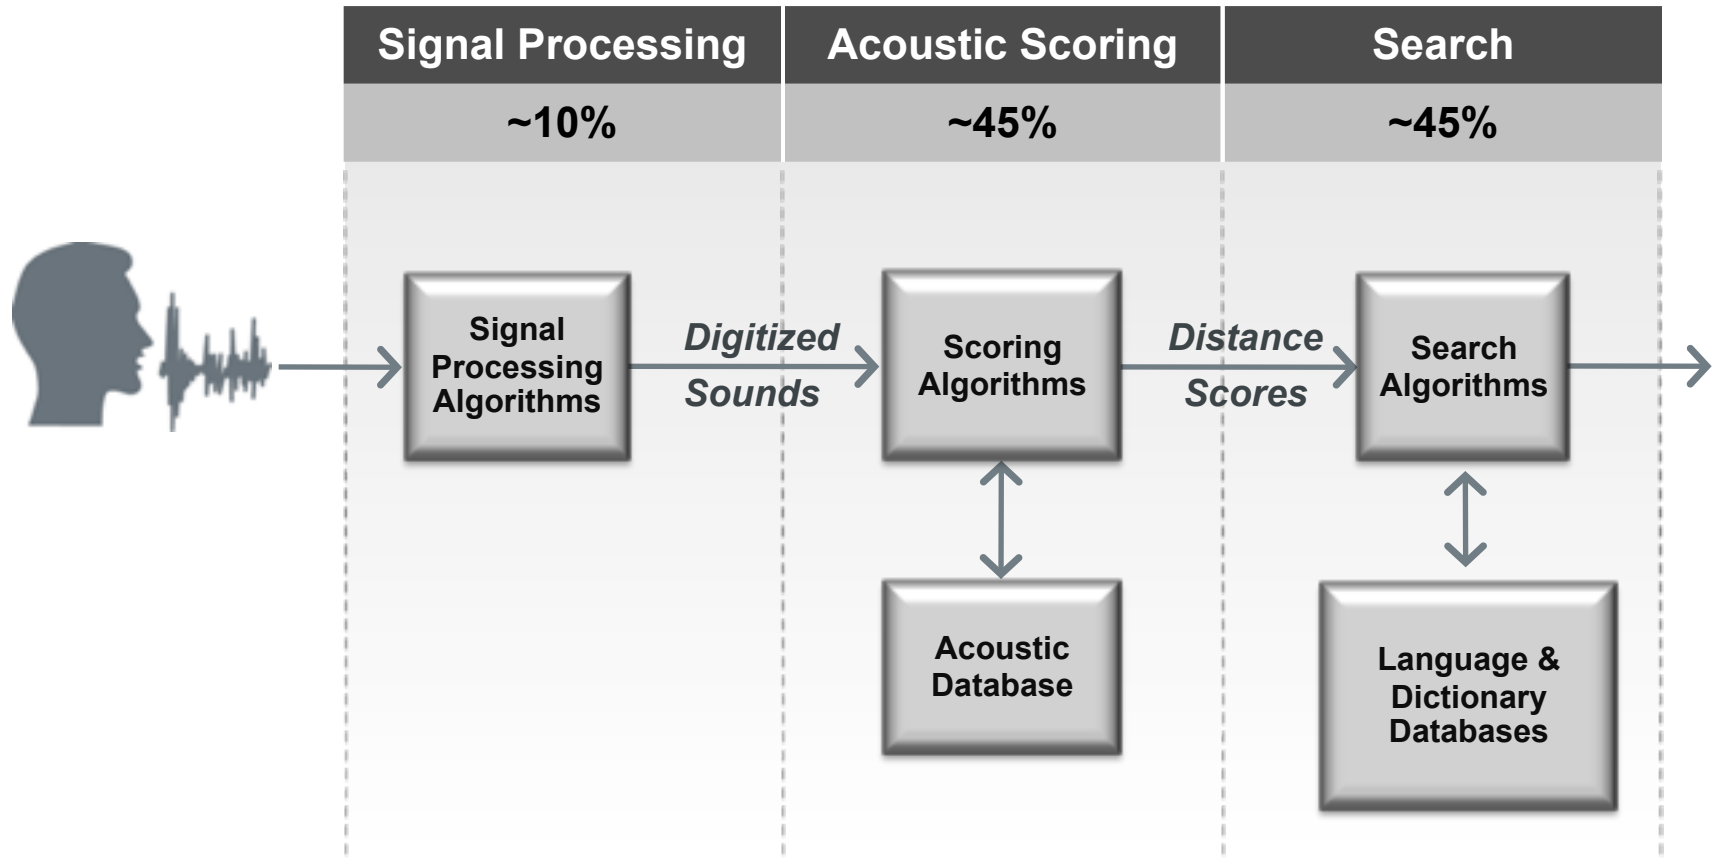
Spansion Acoustic Coprocessor

Spansion® Acoustic Coprocessor

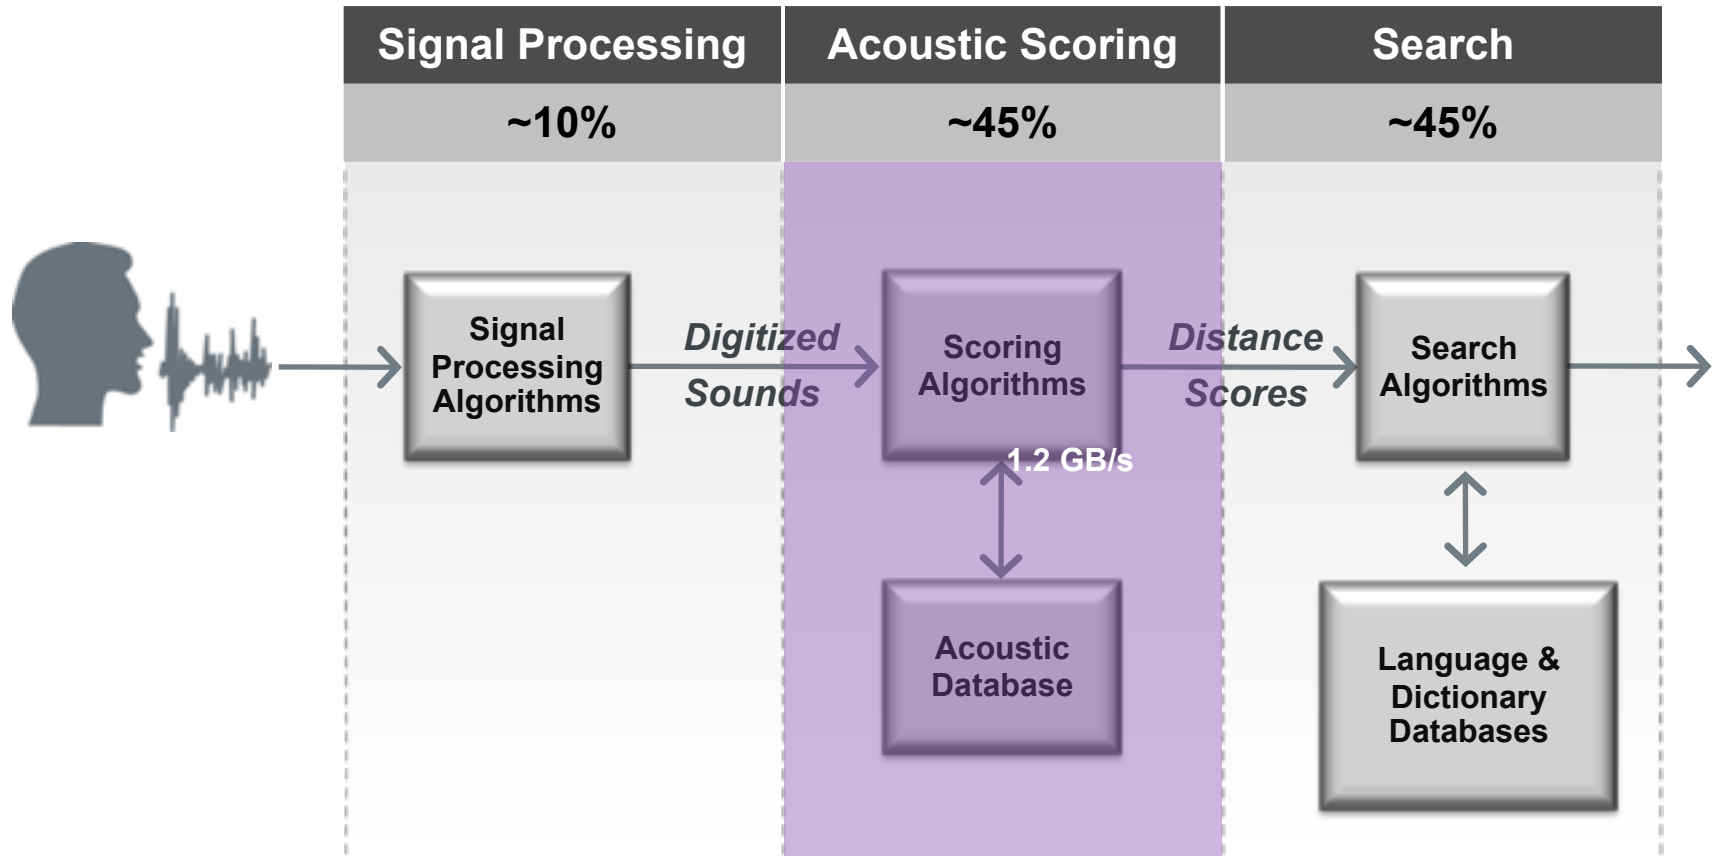
Combines custom logic and high-speed NOR flash memory to deliver faster, more accurate voice recognition

Enhances the user interface in embedded applications, creating a more compelling voice experience for consumers

Voice Recognition – Custom Logic + Memory Solution

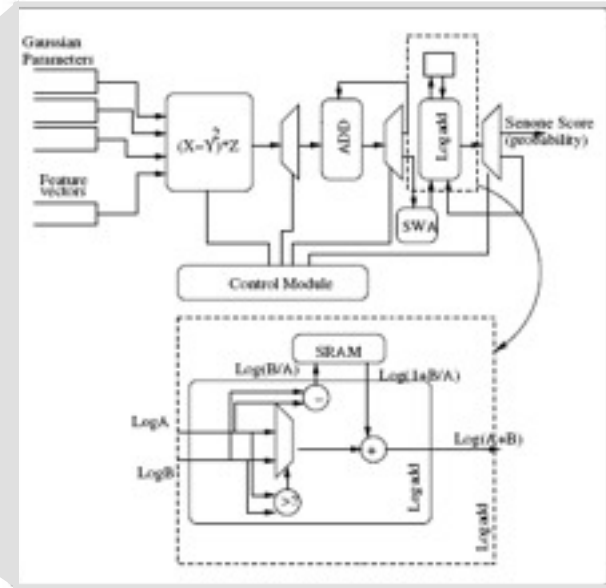
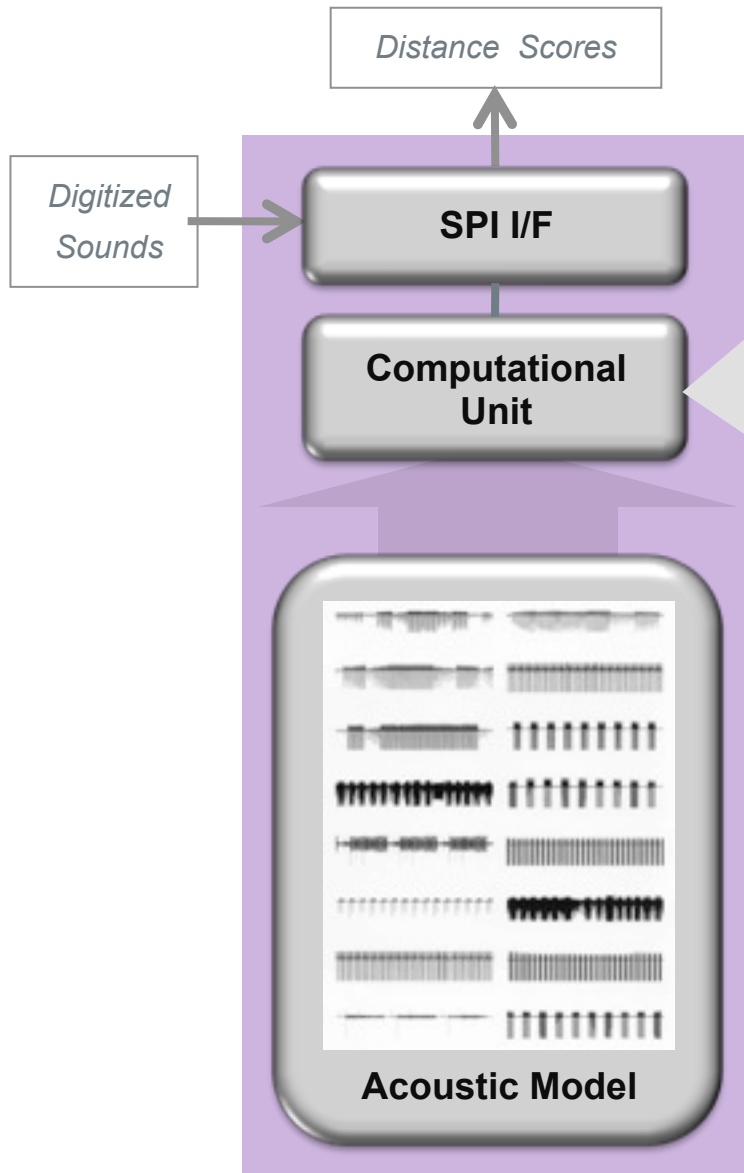


Voice Recognition – Custom Logic + Memory Solution



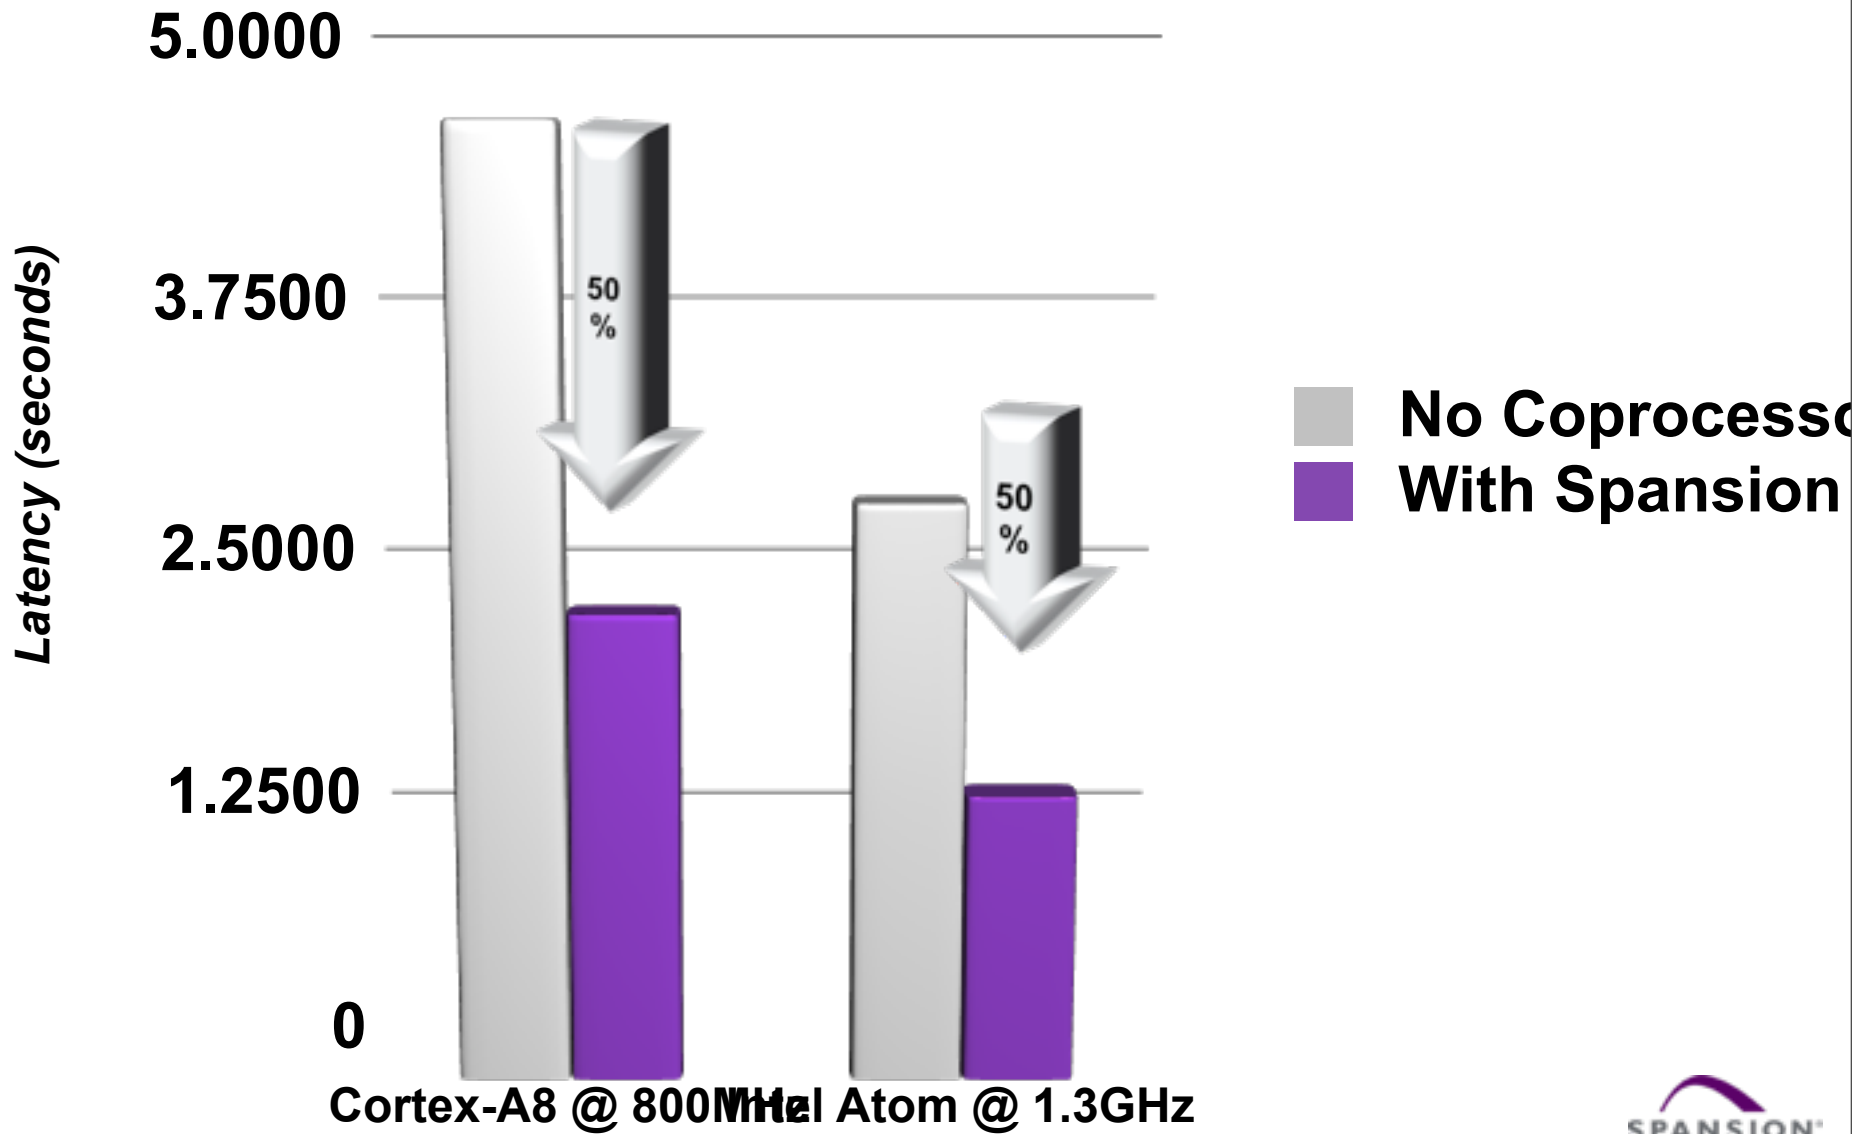
Spansion Acoustic Coprocessor accelerates the speech process by combining custom logic and high speed on-chip NOR flash memory

Spancion Acoustic Coprocessor Concept



- 65nm NOR flash technology
- Computational Unit consists of 8 ALUs, customized to calculate the distance between the incoming sound and the Acoustic Database
- Acoustic Database stored in on-chip flash
- High BW, on-chip data bus
- Low BW, SPI host interface

Infotainment System Characterization (Latency)



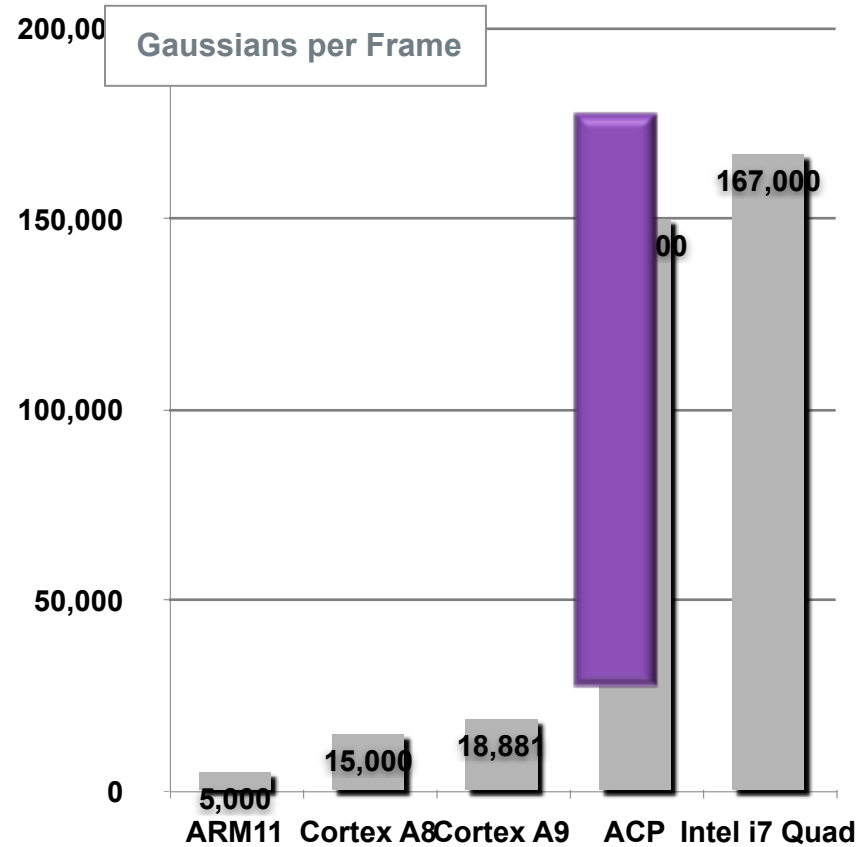
The Future of Embedded Voice Recognition

SpanSion Acoustic Coprocessor:

- Improves Voice Recognition Response Time
 - Up to 50% latency reduction with today's acoustic databases
- Enhances Overall System Performance
 - Up to 50% CPU load reduction
- Enables an Improved Voice Recognition Experience
 - Natural Language Understanding (NLU)
 - Multi-lingual /multi-accent support



Spancion Acoustic Coprocessor - Performance Benchmark



Executing Nuance Acoustic Model scoring algorithms on single processors versus implementing in dedicated HW in Spancion Acoustic Coprocessor



www.spansion.com

SpanSION®, the SpanSION logo, MirrorBit®, MirrorBit® Eclipse™ and combinations thereof are trademarks and registered trademarks of SpanSION LLC in the United States and other countries. Other names used are for informational purposes only and may be trademarks of their respective owners.

This document is for informational purposes only and subject to change without notice. SpanSION does not represent that it is complete, accurate or up-to-date; it is provided "AS IS." To the maximum extent permitted by law, SpanSION disclaims any liability for loss or damages arising from use of or reliance on this document.


Why Speech is Big

Jim Handy

**OBJECTIVE
ANALYSIS**



Moore's Law Shrinks Chips



OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

I/O Limits Product Size



OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

...and Usability



OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Visual Interfaces Distract Drivers



OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Speech: The Smallest UI



OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Many Potential Speech Markets



- 2.5 trillion voicemails
- 2 billion mobile transactions annually
- 6 billion handsets
- 300 million computers
- 70 million automobiles
- 10 thousand mobile app developers
- 500 thousand physicians

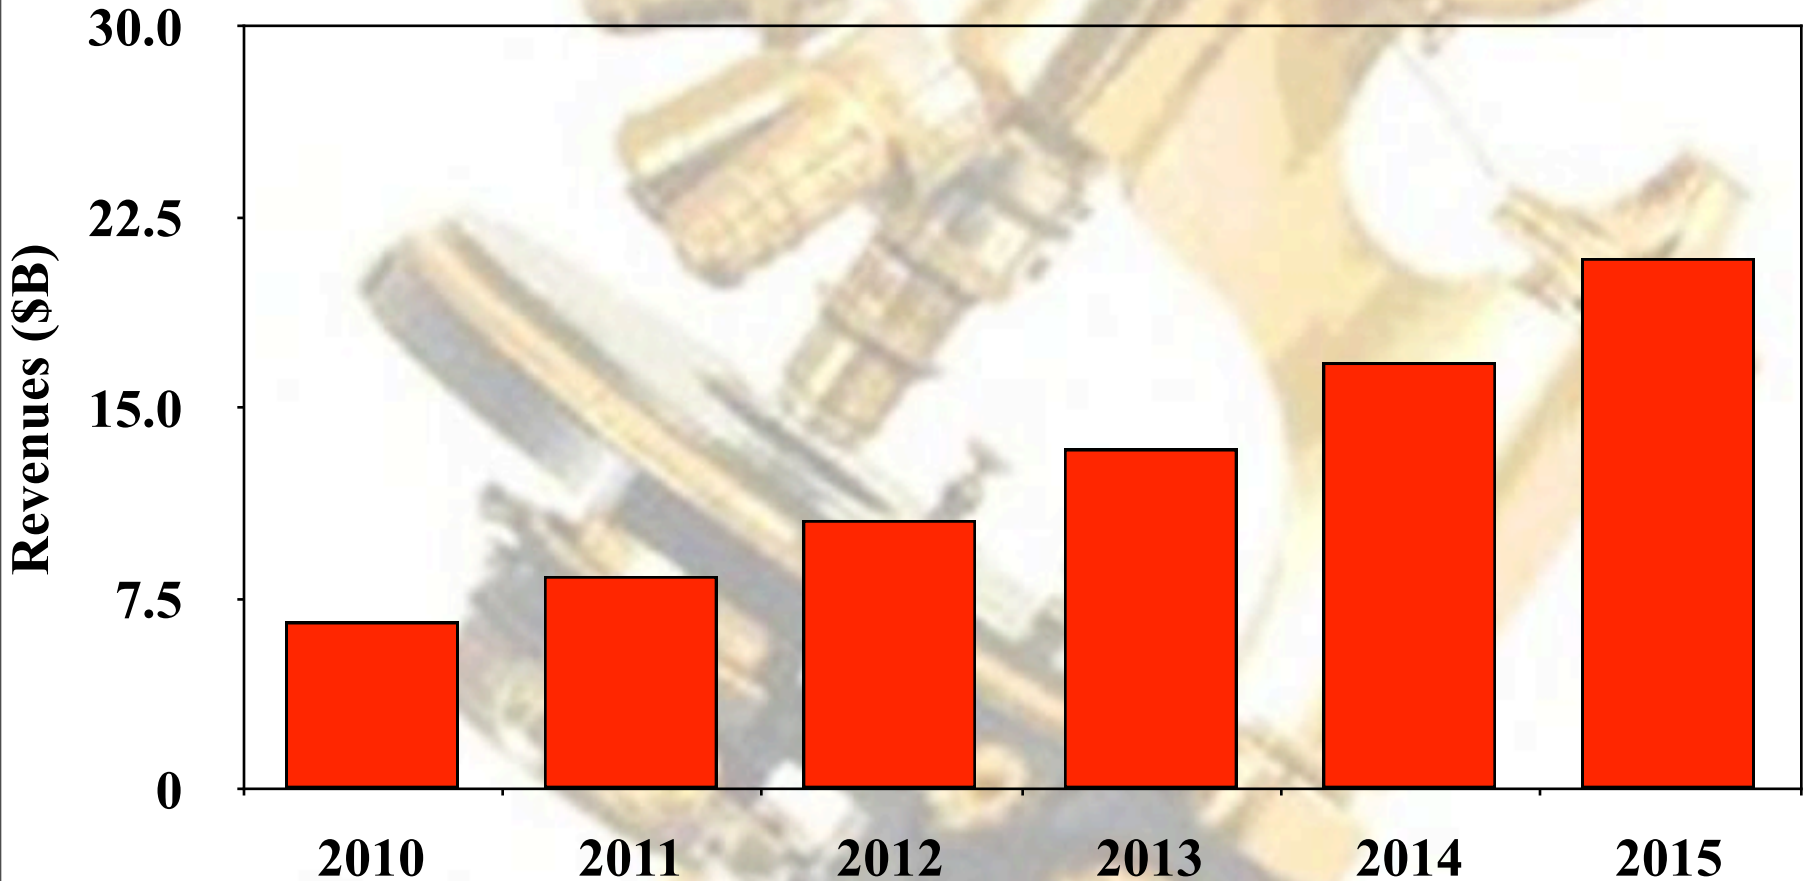
OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Many Potential Speech Markets

- 2.5 trillion voicemails
- 2 billion mobile transactions annually
- 6 billion handsets
- 300 million computers
- 70 million automobiles
- 10 thousand mobile app developers
- 500 thousand physicians

OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Result: Speech Interface Growth



Source; GIA 2010

OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com



Thank You!

Jim Handy

**OBJECTIVE
ANALYSIS**



Session Q&A

Santa Clara, CA
August 2012