Fundamentals of Embedded Computer Vision: Creating Machines That See

Day 3: Processor Choices for Embedded Vision

September 12, 2012

Jeff Bier
President, BDTI
Founder, Embedded Vision Alliance

What is Embedded Vision?

- “Embedded vision” refers to embedded systems that extract meaning from visual inputs
    - Embedded vision is distinct from multimedia
- Emerging high-volume embedded vision markets include automotive safety, surveillance, and gaming
    - The Xbox Kinect is the fastest-selling CE device to date: 10 million units in the first 6 months

$130 including game

$920 installed

$300 + $6/month

Why is Embedded Vision Proliferating Now?

1. It has the potential to create huge value
    - Applications in consumer, medical, automotive, entertainment, retail, industrial, aerospace, …
2. It's now possible
    - Sufficiently powerful, low-cost, energy-efficient processors are now emerging
3. Increasingly, it will be expected
    - As embedded vision becomes common in gaming, consumer electronics, and automotive equipment, consumers will expect it

The Processing Challenge

- Embedded vision applications typically require:
  - Very high performance
  - Programmability
  - Low cost
  - Energy efficiency

- Achieving all of these together is difficult
  - Dedicated logic yields high performance at low cost, but with little programmability
  - General-purpose CPUs provide programmability, but with weak performance or poor cost-, energy-efficiency

How is Embedded Vision Implemented?

Demanding embedded vision applications will most often use a combination of processing elements (similar to wireless baseband chips), e.g.:

- CPU for complex decision-making, network access, user interface, storage management, overall control
- High-performance DSP-oriented processor for real-time, moderate-rate processing with moderately complex algorithms
- Highly parallel engine(s) for pixel-rate processing with simple algorithms

Objectives of This Presentation

- Provide an introduction to embedded vision processor options and associated trade-offs

Processor Types for Embedded Vision

While any processor can in theory be used for embedded vision, the most promising types today are:

- High-performance embedded CPU
- Application-specific standard product (ASSP) + CPU
- Graphics processing unit (GPU) + CPU
- DSP processor + accelerators + CPU
  - Mobile “application processor”
- Field programmable gate array (FPGA) + CPU

The Path of Least Resistance

Making Dinner on a Tuesday

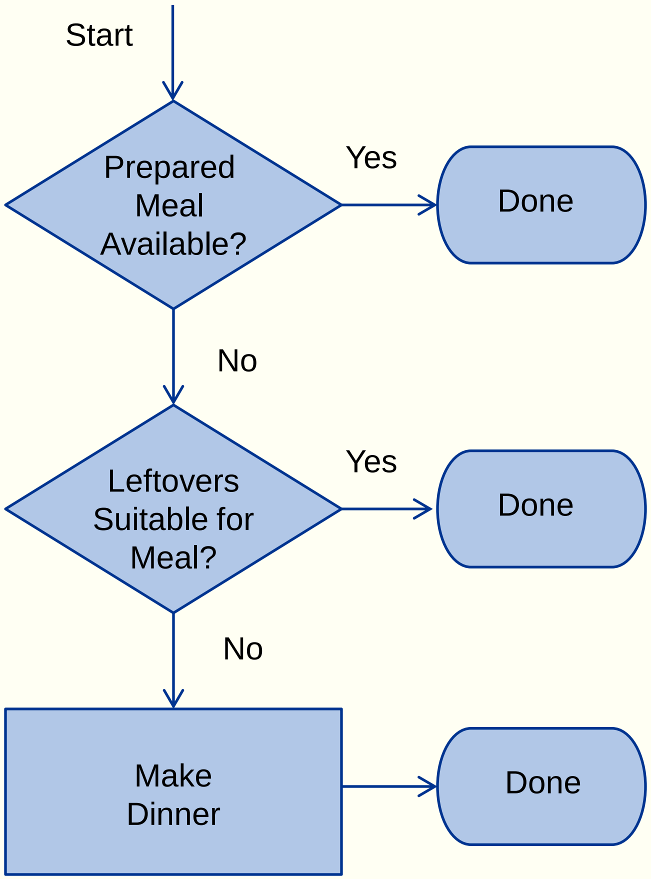

Selecting a Processor for an Embedded Vision Application

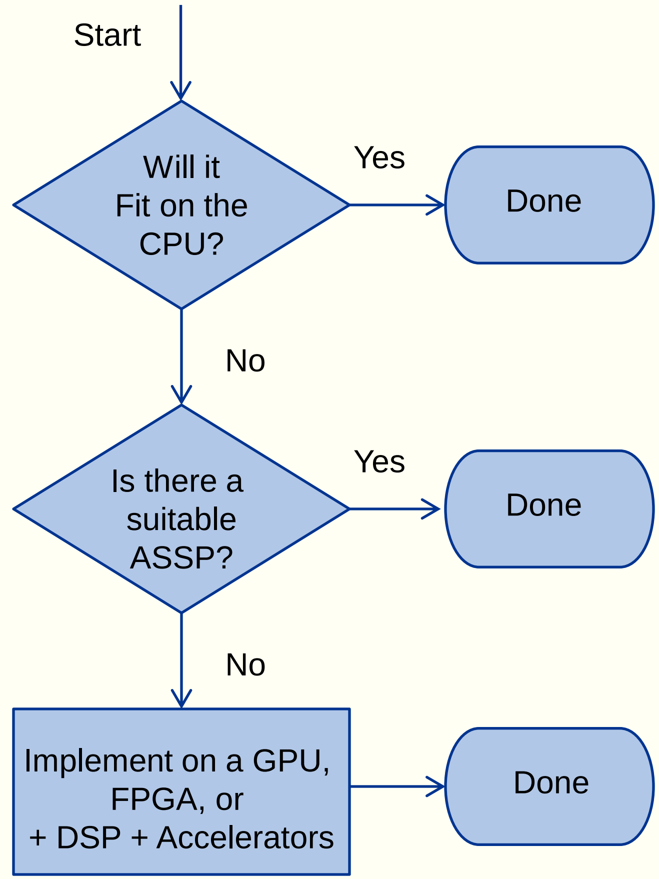

High-performance Embedded CPUs

Though challenged with respect to performance and efficiency, unaided high-performance embedded CPUs are attractive for some vision applications

- Vision algorithms are initially developed on PCs with general-purpose CPUs
- CPUs are easiest to use: tools, operating systems, middleware, etc.
- Most systems need a CPU for other tasks

However:

- Performance and/or efficiency is often inadequate
- Memory bandwidth is a common bottleneck

Example: Intel Atom (used in NI 1772C Smart Camera)

Best for: Applications with modest performance needs

Application-Specific Standard Product + CPU

Application-specific standard products (ASSPs) are specialized, highly integrated chips tailored for specific applications or application sets

- ASSPs may incorporate a CPU, or use a separate CPU chip
- 👍 By virtue of specialization, they tend to deliver superior cost- and energy-efficiency
- 👍 They usually include strong application-specific software development infrastructure and/or application software

However:

- 👎 The specialization may not be right for your particular application
- 👎 They may come from small suppliers, which can mean more risk
- 👎 They use unique architectures, which can make programming them, and migration to other solutions, more difficult
- 👎 Some are not user-programmable

Example: PrimeSense PS1080-A2 (used in Kinect)

Best for: Ultra-high-volume, low-cost applications

Graphics Processing Unit (GPU) + CPU

GPUs, mainly used for 3-d graphics, are increasingly capable of being used for other functions

- Referred to as "general-purpose GPU" or "GPGPU"
- Often used for vision algorithm development
- Widely available; easy to get started with parallel programming
- Well-integrated with CPU (sometimes on one chip)
- Typically cannot be purchased as a chip, only as a board, with limited selection of CPUs
- Low-cost, low-power GPUs (designed for smart phones, tablets) are not GPGPUs

Example: NVIDIA GT240 (used in GE NP240 rugged single-board computer)

Best for: Performance-hungry apps with generous size/power/cost budgets

DSP Processor + Co-processors + CPU

Digital signal processors ("DSP processors" or "DSPs") are processors specialized for signal processing algorithms

- This makes them more efficient than CPUs for the kinds of signal processing tasks that are at the heart of vision applications
- DSPs are relatively mature and easy to use compared to other kinds of parallel processors

However:

- DSPs often lack sufficient performance, and aren't as easy to use as CPUs
- Hence, DSPs are often augmented with specialized co-processors and a CPU on the same chip

Example: Texas Instruments DaVinci (used in Archerfish Solo consumer smart surveillance camera)

Best for: Apps with moderate performance needs and moderate size/power/cost budgets

Mobile "Application Processor"

A mobile "application processor" is a highly integrated system-on-chip, typically designed primarily for smart phones but also used for other applications

- Typically comprise a high-performance CPU core and a constellation of specialized co-processors: GPU, VPU, 2-d graphics, image acquisition, etc.
- 👍 Energy efficient
- 👍 Often have strong development support, including low-cost development boards, Linux/Android ports, etc.

However:

- 👎 Specialized co-processors are usually not user-programmable

Example: Qualcomm QSD8650 (used in HTC Incredible)

Best for: Apps with moderate performance needs, wireless connectivity, and tight size/power/cost budgets

FPGA + CPU

FPGA flexibility is very valuable for embedded vision applications

- Enables custom specialization and enormous parallelism
- Enables selection of I/O interfaces and on-chip peripherals

However:

- FPGA design is hardware design, typically done at a low level (register transfer level)
- Ease of use improving due to:
  - Platforms
  - IP block libraries
  - Emerging high-level synthesis tools
- Low-performance CPUs can be implemented in the FPGA; high-performance integrated CPUs on the horizon

Example: Xilinx Spartan-3 XC3S4000 (used in Eutecus Bi-i V301HD intelligent camera)

Best for: High performance needs with tight size/power/cost budgets

CONCLUSIONS

Conclusions

- There are many types of processors available for embedded vision, and many examples of each type
- Which processor is best depends on details of the application
- There are tough trade-offs among performance, power consumption, price, flexibility, ease-of-use, and risk
- Parallel processing of some type is usually necessary
- Heterogeneous multiprocessors are often a good fit
- Ease of use depends not only on the complexity of the architecture and quality of the tools, but also on other development resources, such as software libraries and technical support

Next Up: Algorithms, Tools, and Techniques

In tomorrow's presentation, we will:

- Explore fundamental building-block vision algorithms such as line detection
- Illustrate these algorithms with downloadable demos
- Introduce OpenCV, a free, open source computer vision software library

Next Up: Download BDTI OpenCV Demo Package from the Embedded Vision Academy (Free Registration Required)

BDTI OpenCV Executable Demo Package—No programming required
www.embeddedvisionacademy.com/opencvdemo

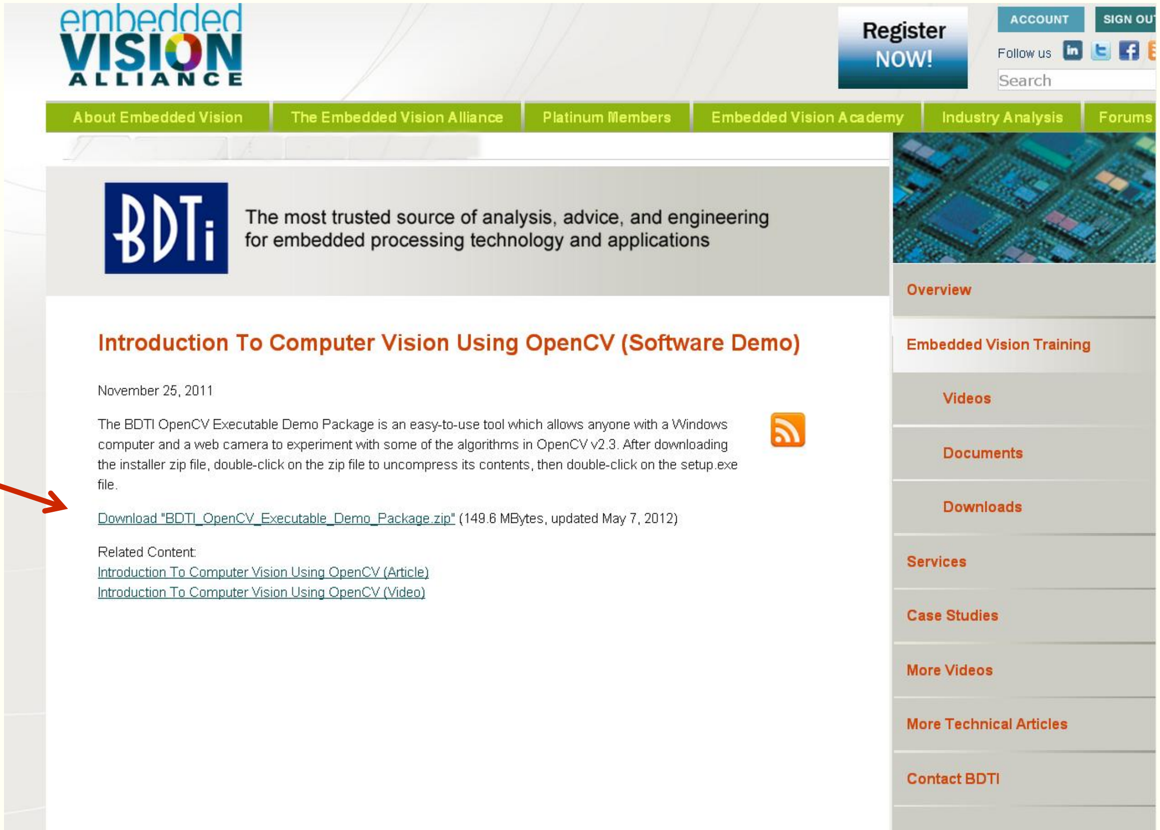

RESOURCES

The *Embedded Vision Summit*
A Free Educational Event for Engineers—Boston, September 19th

Learn how to use the coolest new technology in the industry to create "machines that see"

- Technical presentations on sensors, processors, tools, and design techniques
- Keynotes by Prof. Rosalind Picard, MIT Media Lab and Gary Bradski, CEO, OpenCV Foundation
- Cool demonstrations and opportunities to meet with leading vision technology suppliers

Part of UBM Electronics' DESIGN East event

- DESIGN East also includes the Embedded Systems Conference, Sensors in Design, DesignMED, Android Summit, LED Summit, and exhibits

The Summit is free, but space is limited. To begin the registration process, send an email to summit@Embedded-Vision.com

For more info: www.embedded-vision.com/embedded-vision-summit

Selected Resources: The Embedded Vision Alliance

The Embedded Vision Alliance is an industry partnership to transform the electronics industry by inspiring and empowering engineers to design systems that see and understand

Free Resources from the Embedded Vision Alliance

The Embedded Vision Alliance web site, at www.Embedded-Vision.com covers embedded vision applications and technology, including interviews and demonstrations

The Embedded Vision Academy, a free service of the Alliance, offers free in-depth tutorial articles, video "chalk talks," code examples and discussion forums: www.EmbeddedVisionAcademy.com

The Embedded Vision Insights newsletter provides updates on new materials available on the Alliance website. Sign up at www.Embedded-Vision.com/user/register

Embedded vision technology and services companies interested in becoming sponsoring members of the Alliance may contact info@Embedded-Vision.com

Embedded Vision Insights
The Latest Developments on Designing Machines that See

Follow us

What is BDTI?

BDTI is a group of engineers dedicated to helping the electronics industry effectively use embedded digital signal processing technology

BDTI performs hands-on, independent benchmarking and evaluation of chips, tools, algorithms, and other technologies

BDTI helps system designers implement their products through specialized engineering services

BDTI offers a wealth of free information for engineers

Additional Resources

BDTI's web site, www.BDTI.com, provides a variety of free information on processors used in vision applications.

BDTI's free "InsideDSP" email newsletter covers tools, chips, and other technologies for embedded vision and other DSP applications. Sign up at www.BDTI.com.