Digital Signal Processor evolution over the last 30 years

François Charlot

IEEE Senior Member

April 2010, revised March 2013

Presentation outline

- DSP algorithms until the 80's
  - Filters, Fast Fourier Transform, Speech analysis and synthesis, GSM channel equalization
- The first decade of single-chip DSPs
- Early 90's: Emerging DSP markets and enablers
- The great divide:
  - Pervasive DSPs, Mobile DSPs, High-performance DSPs
- Some more DSP algorithms
  - CELP, MP3, JPEG, MPEG-2
- Digital Signal Processors 90's-00's
  - Pervasive DSPs: hybrid DSP/MCUs
  - Mobile DSPs: TI, competition
  - High-performance DSPs
- Specialized DSPs
- Configurable DSPs
- ARM attempts at DSP
- AnySP – The best mobile DSP?
- FPGAs
- Comparison of DSP implementations
- Some failed attempts
- DSP market
- Conclusion

DSP algorithms until the 80’s

Basic DSP algorithms

- The z-notation is used to represent a sampled signal and operations over it.

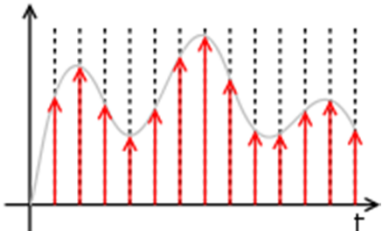

$$X(z) = \mathcal{Z}\{x[n]\} = \sum_{n=-\infty}^{\infty} x[n] z^{-n}$$

- The specification of early DSPs was to execute efficiently the following algorithms:
  - Finite-impulse response (FIR) filters
  - Infinite-impulse response (IIR) filters
  - Convolution and correlation
  - Fast Fourier Transform (FFT)

FIR and IIR filters

- FIR filter:
  - N coefficients (filter order)

$$H(z) = \sum_{n=-\infty}^{\infty} h[n] z^{-n}$$

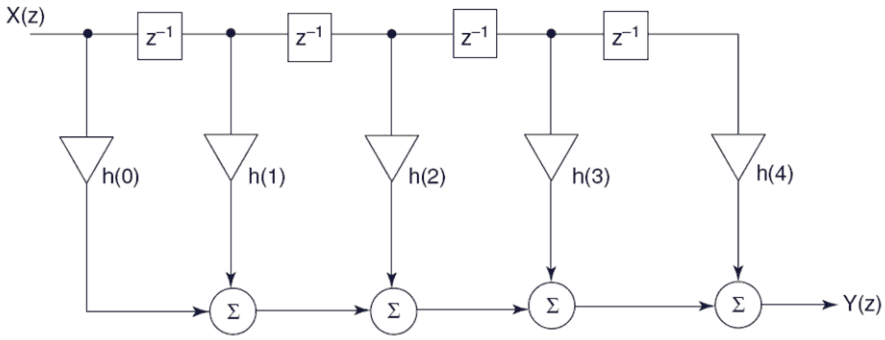

  - Stable, could be linear phase. Impulse response dies at sample N.
- IIR filter:
  - Much better roll-off than FIR filter for a given order.

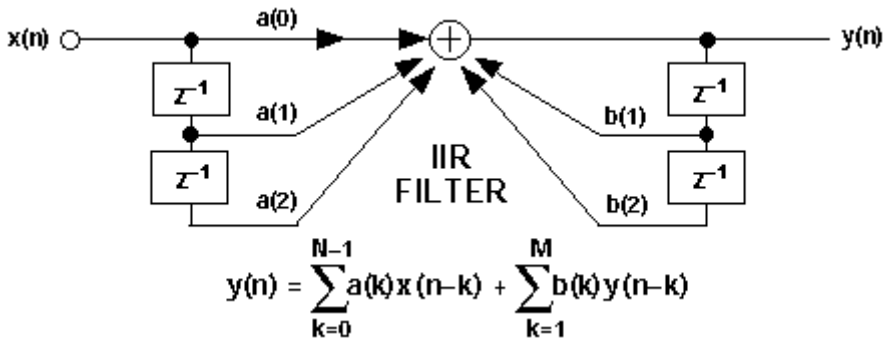

Definition of convolution

- Definition:

$$(f * g)[n] = \sum_{m=-\infty}^{\infty} f[m]\, g[n-m]$$

- Combines an input signal with the impulse response of the system.

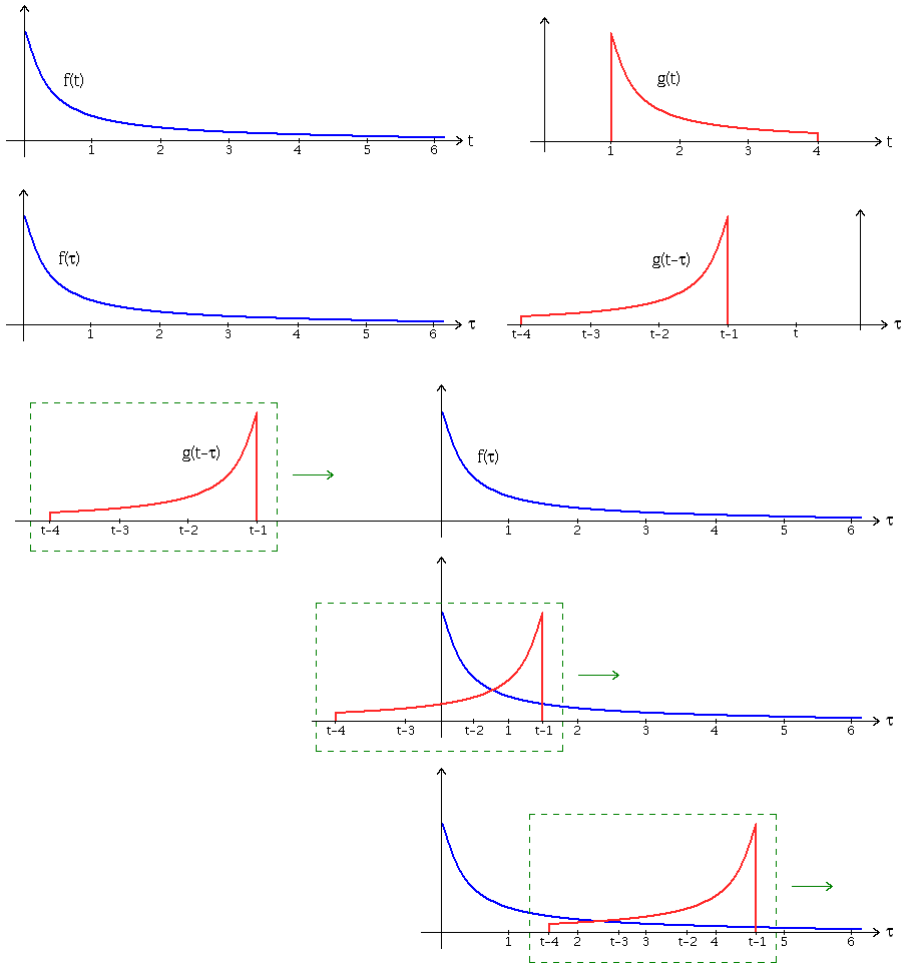

Image source: Wikipedia article on Convolution

Examples of convolutions

- Amplification and attenuation:

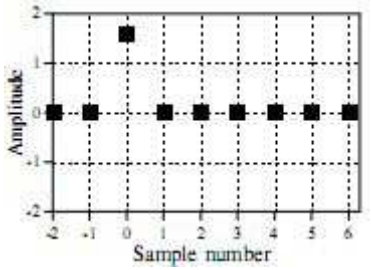

- Echo:

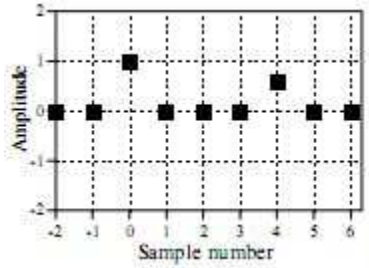

- Derivative:

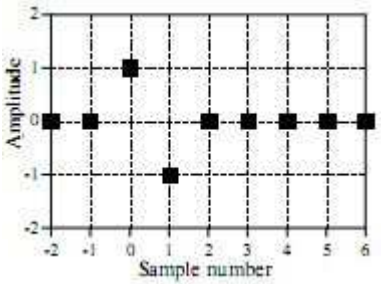

- Integral:

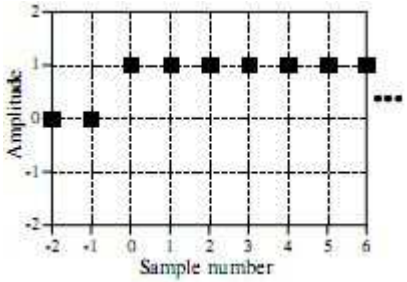

Source: The Scientist and Engineer's Guide to Digital Signal Processing, www.dspguide.com

Examples of convolutions – Filters

- Examples of impulse response of low-pass filters:

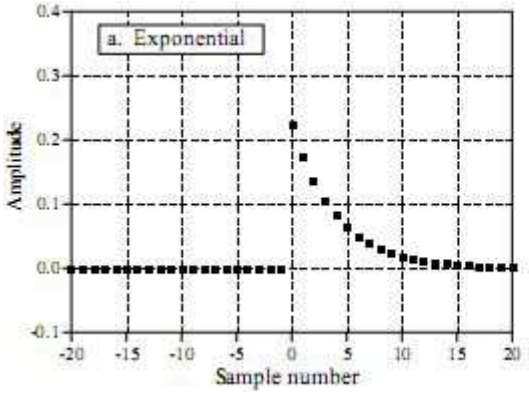

Simplest recursive filter

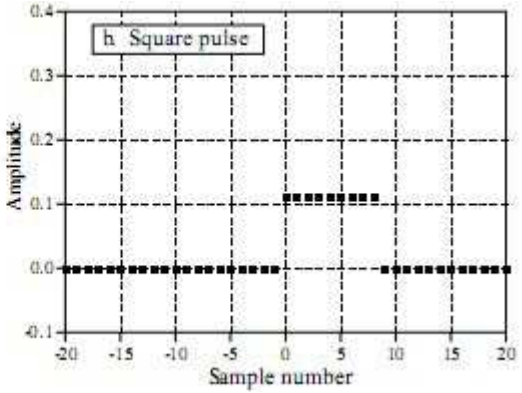

- Reduces noise
- Maintains edge sharpness

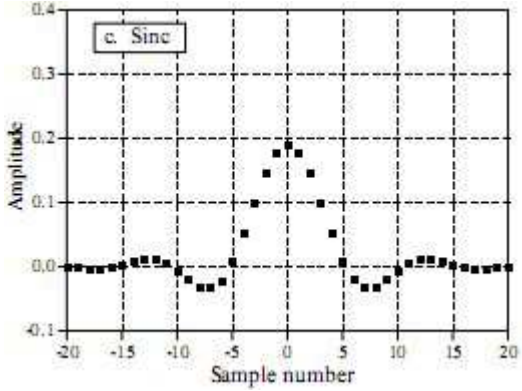

Separate bands of frequencies

Source: The Scientist and Engineer's Guide to Digital Signal Processing, www.dspguide.com

Some applications of convolution

- Radars: analyzing the measured impulse response
- Digital filter design
- Distance phones calls: echo suppression
  - Create an impulse response that counteracts that of the reverberation.
- Audio: apply the impulse response of a real environment on an audio signal.

Correlation

- Correlation is identical to the convolution with the difference no signal is reversed.
- Application examples:
  - Correlation between a radar transmitted and received signals.
  - Autocorrelation: white noise removal, pitch or tempo estimation...

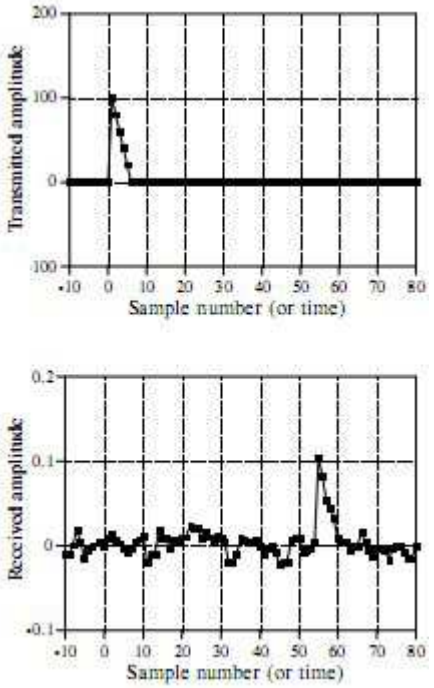

Source: The Scientist and Engineer's Guide to Digital Signal Processing, www.dspguide.com

Fast Fourier Transform (1/2)

- The Discrete Fourier Transform converts a series of time-domain values (e.g., a sampled signal) into the frequency domain.
- The FFT is an O(N log N) algorithm *vs*. O(N^2) for the DFT.
  - Invented by Cooley and Tukey in 1965.
  - Most often it is used to divide a N-point FFT into 2 N/2 ones (radix-2 FFT).
  - Basic element is a 2-point FFT called butterfly.
    - W_N are called the twiddle factors.

$$W_N = e^{-j\left(\frac{2\pi}{N}\right)}$$

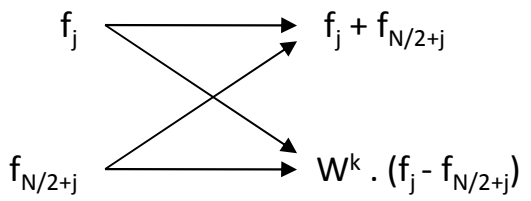

Fast Fourier Transform (2/2)

- Diagram of an 8-point radix 2 FFT:

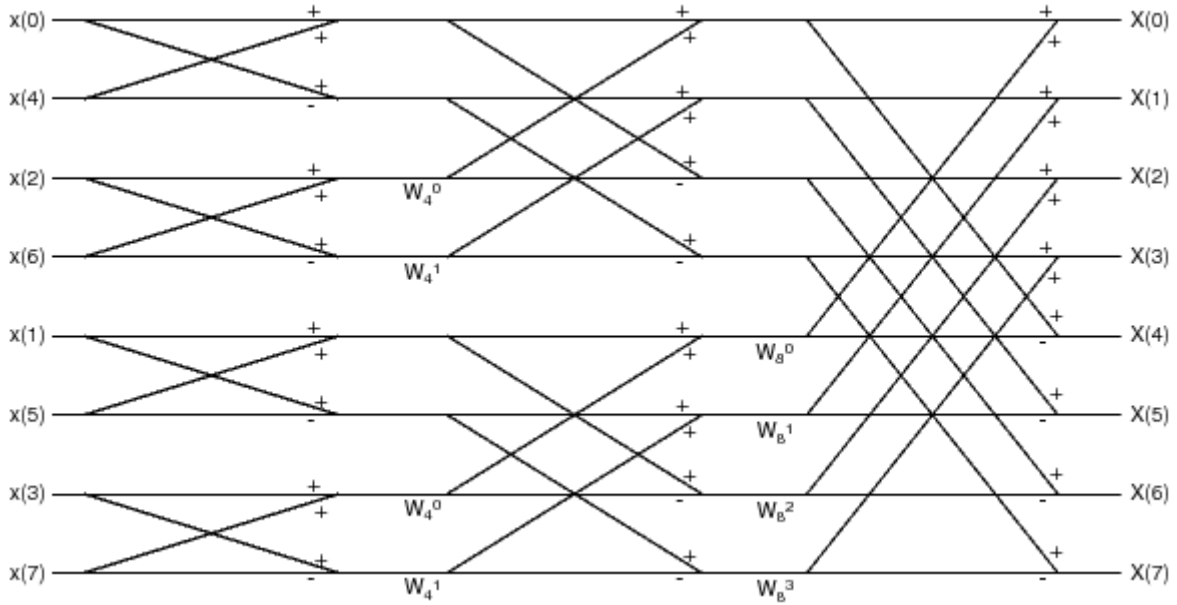

  - Note the pattern of the input samples indices.
- Higher radix (4 or 8) speed up the computation at the expense of larger code size.
  - ~ 25% fewer multiplications.

Applications of FFT

- Spectral analysis of signals
  - Need samples over a full period of the signal.
  - Need to apply a window (e.g., Hann) for non-periodic signals so to attenuate the side spectral lobes.

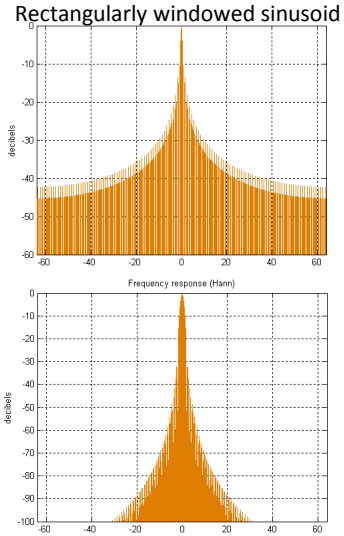

- Frequency response of a system

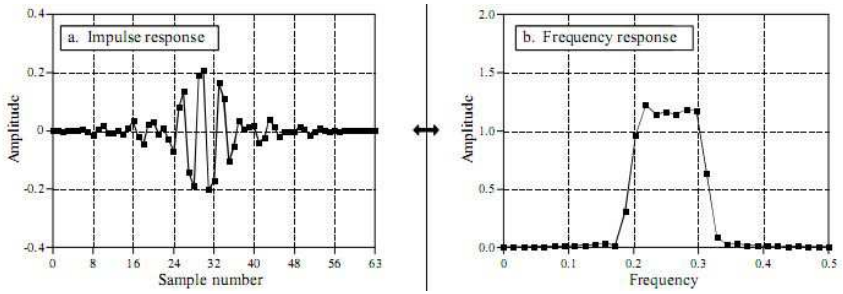

- Alternative to deconvolution
  - Determine input from impulse response and output
- Less computationally intensive than convolution

time domain → frequency domain → convolution → time domain

FFT x IFFT

Image source: Wikipedia article on Window functions and www.dspguide.com

Speech analysis and synthesis

- The human speech production system is modeled into a linear filter excited by an impulse train (voiced speech) or white noise (unvoiced speech).

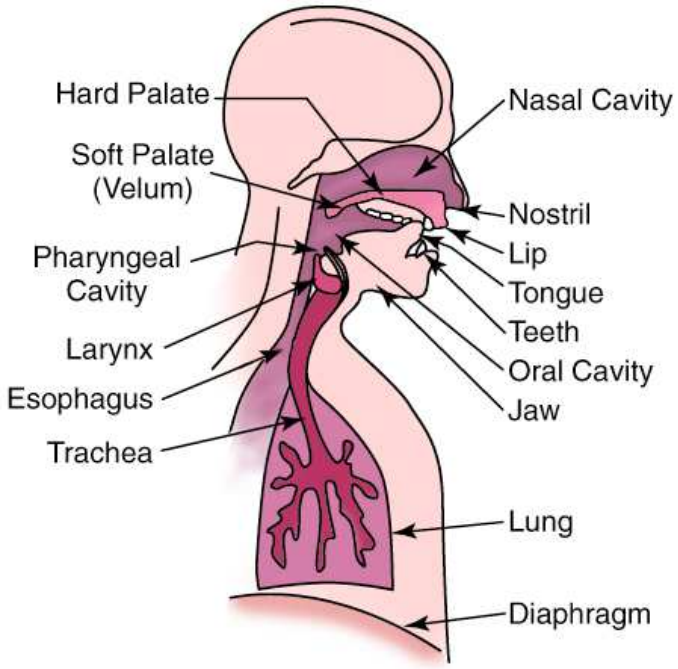

becomes...

(source-filter model of speech production)

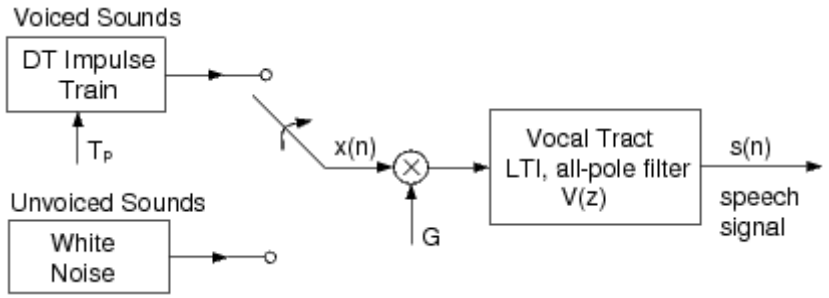

Source: Charles A. Bouman, Digital Signal Processing with Applications, Purdue course ECE 438, available at cnx.org

Linear Predictive Coding

- Linear Predictive Coding used since the mid-70's.
  - Analyze the signal over a time window during which the model parameters may be deemed constant (~20-25 ms).
  - Transmit the voice/unvoiced state, the frequency value, all parameters of the filter.
  - Some bits are more important than others, hence the need for channel encoding/decoding
- Examples:
  - LPC-10 in the TI Speak and Spell (1978)
    - ~1000 bits/second; pipelined hardwired implementation; Gene Frantz was one of the designers...
  - 2400 bits per second FS-1015 US DoD/NATO standard (1984)

GSM voice and channel codec

- The full rate GSM codec uses RPE-LTP-LPC:
  - LPC for short-term prediction (8-stage filter).
  - Long-Term Prediction:
    - LPC signal is reconstructed and correlated with the original one with a lag of 40-120.
    - The maximum correlation is kept, then the gain is processed so to get the loudness information.
    - The residual signal is down-sampled and encoded using APCM.
  - First demo on TMS320C50 at ICASSP in May 1989.
- Channel encoding:
  - 260 bits / 20 ms from the codec classified by importance.
    - Parity bits and/or duplication by convolutional encoding translate to 378 bits per frame.
  - Decoding through the Viterbi algorithm.

Trellis diagram

- Example trellis diagram:

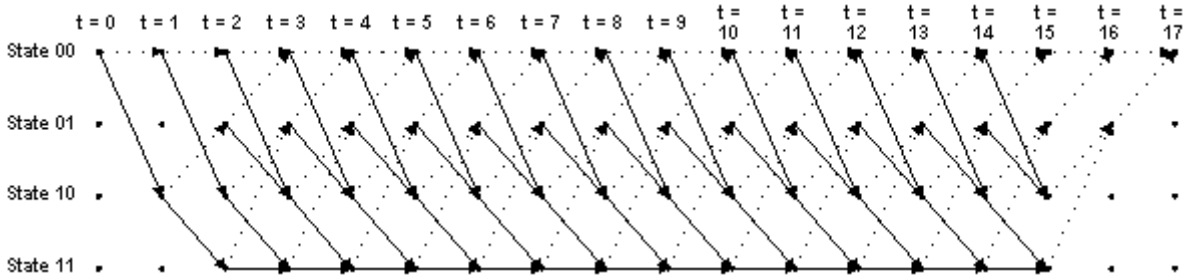

- Path with an example message:

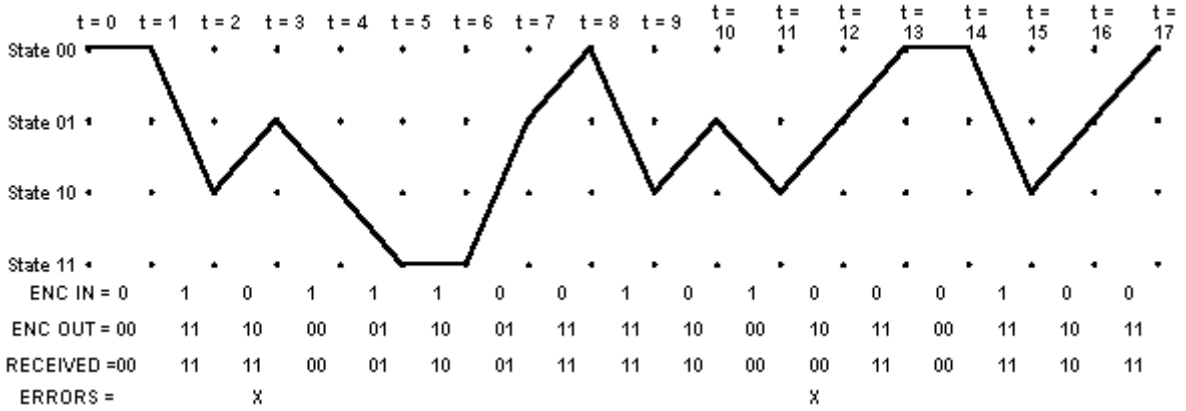

Source: Chip Fleming, Tutorial on Convolutional Coding with Viterbi Decoding, http://home.netcom.com/~chip.f/Viterbi.html

GSM channel equalization

- Some training bits are added to the GSM frames to allow learning of the channel response:

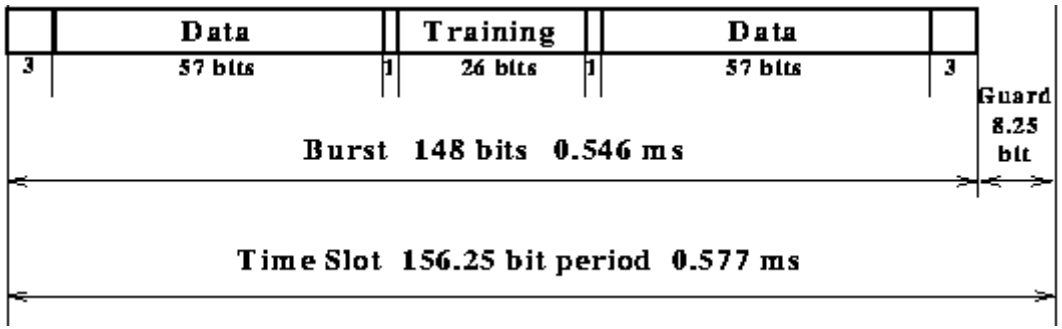

- Diagram of the GSM channel equalizer:

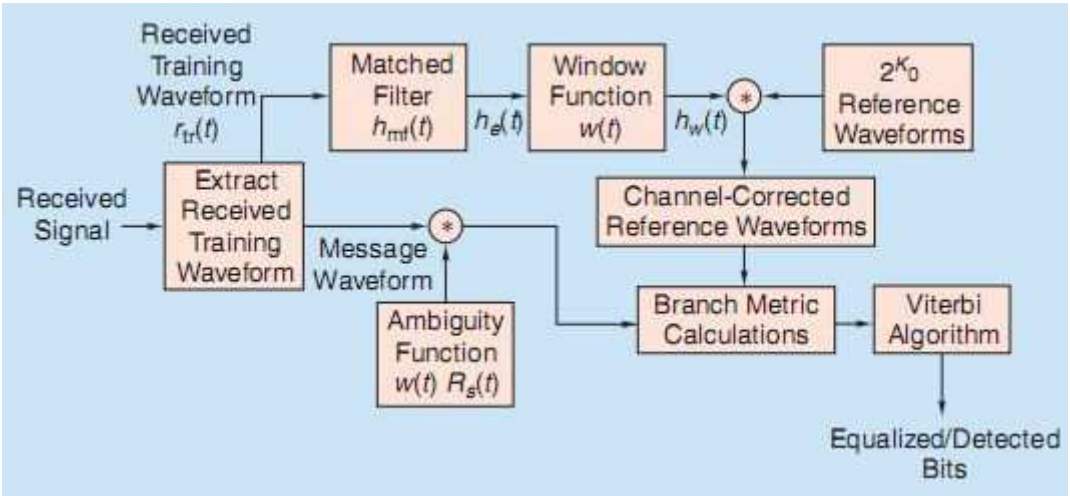

The Viterbi algorithm

- Add-compare-select:
- Accumulate path metrics
- Decode most likely path

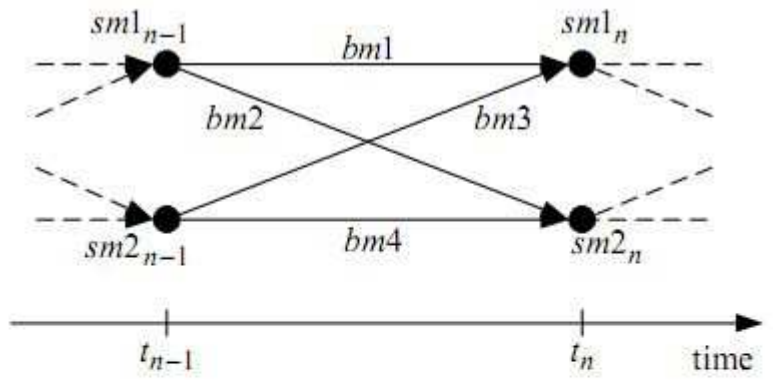

$$sm1_n = \min\left(sm1_{n-1} + bm1,\ sm2_{n-1} + bm3\right)$$

$$sm2_n = \min\left(sm1_{n-1} + bm2,\ sm2_{n-1} + bm4\right)$$

Add Add

Select Compare

Source: Prof. Vladimir Stojanovic, 6.973 Communication System Design course lecture notes, http://ocw.mit.edu

Numbers representation

- 16-bit integer: -32768..32767
- Q15 notation scales -1..+1 up to -32768..+32767
- Limits the re-scaling needs during the computations since multiplications stay within the -1..+1 interval.
  - Q15 x Q15 = Q30
- Additions may require extra guard bits.
  - May also feature automatic saturation.
- Alternative is floating-point representation
  - e.g., ±0.mantissa . 2^{exp}

The first decade of single-chip DSPs

Early times of DSP'ing

Time Frame	Approach	Primary Application	Enabling Technologies
Early 1970's	Discrete logic	• Non-real time processing • Simulation	• Bipolar SSI, MSI • FFT algorithm (1965)
Late 1970's	Building block	• Military radars • Digital Comm.	• Single chip bipolar multiplier • Flash A/D
Early 1980's	Single Chip DSP µP	• Telecom • Control	• µP architectures • NMOS/CMOS
Late 1980's	Function/Application specific chips	• Computers • Communication	• Vector processing • Parallel processing
Early 1990's	Multiprocessing	• Video/image processing	• Advanced multiprocessing • VLIW, MIMD, etc.
Late 1990's	Single-chip multiprocessing	• Wireless telephony • Internet related	• Low power single-chip DSP • Multiprocessing

Source: Prof. Kurt Keutzer, Berkeley course CS-252, 2000

DSP context in the early 80's

- Digital telephony standards
  - 1980: studies on end-to-end fully digital link
  - 1984: G.721 (32 kb/s ADPCM)
  - 1984: CCITT Red Book (ISDN standard)
- Performance of general-purpose processors:
  - Intel 8086 (1978): 70/118-cycle MPY, 5-10 MHz
  - Motorola 68000 (1979): 70-cycle MPY, 4-12.5 MHz
  - Intel 80286 (1982): 24-cycle MPY, 6-12.5 MHz

DSP requirements, 80's to early 90's

- Multiply-accumulate
- Scaling (fixed-point DSPs)
- Z^{-1}
- FFT
- Keep MAC fully busy
- Handling of data frames, bitwise convolution
- Viterbi support
- Hardwired MAC
- Shifters
- Circular buffers
- Add-sub-mpy, bit-reverse
- 2 data accesses per cycle, 0-overhead software loops
- Fast logical operations
- Add-compare-select

The very first devices

- Intel 2920 (1978): AD/DA, shift-and-add (no MPY), 192-word program memory, no program flow control. 4.5 µm NMOS.
- AMI S2811 (1979): 16-bit ALU, 12x12 MPY, not a standalone µP. 4.5 µm VMOS.
- AT&T DSP1 (1979): true DSP but not for merchant market. 4.5 µm NMOS.
- NEC 7720 (1980): 32-bit ALU, 16x16 MPU, dual accumulator, 4 MHz instruction rate. 3 µm NMOS.

Dates are those of first prototypes

Enter the TI TMS32010...

- 1976: TMS9900 16-bit µP, ahead of its time.
- 1978: TI Management realizes it is not making the needed headway in the µP market. Look for next big thing.
- 1978: Harvey Cragon suggests a DSP device.
- 1979: Intel 2920 paper for ISSCC convinces TI Management to go ahead. Signal Processing Computer architecture defined. Dr Surendar Magar (DSP PhD) hired from Plessey UK.
- Sep. 1981: start of first samples fabrication.
- Feb. 1982: Magar presents paper at ISSCC.

Source: A.W. Leigh, The TMS32010. The DSP chip that changed the destiny of a semiconductor giant.

Some TMS32010 facts and figures

- 3 (2.7) µm NMOS, 57,000 transistors, 43.8 mm^2.
- 5 MHz instruction rate, 2-cycle MAC.
- 144-word RAM, 1.5 Kw ROM.
- RAM size fits a 64-point FFT.
- Initially no hardware multiplier – Magar put it in 1979.
- Initially internal ROM only, 24-pin package.
  - Microprocessor mode added early 1981 (TMS7000 feedback).
- Features added to make the chip self-emulating (extra stack level...).
- Initially was to be named TMS10010 (from the TMS1000).
  - 32010 made up from 32 (bits), 01 (first device in series) and 0 (extra performance like TMS9900/TMS99000).

Block diagram and die photograph

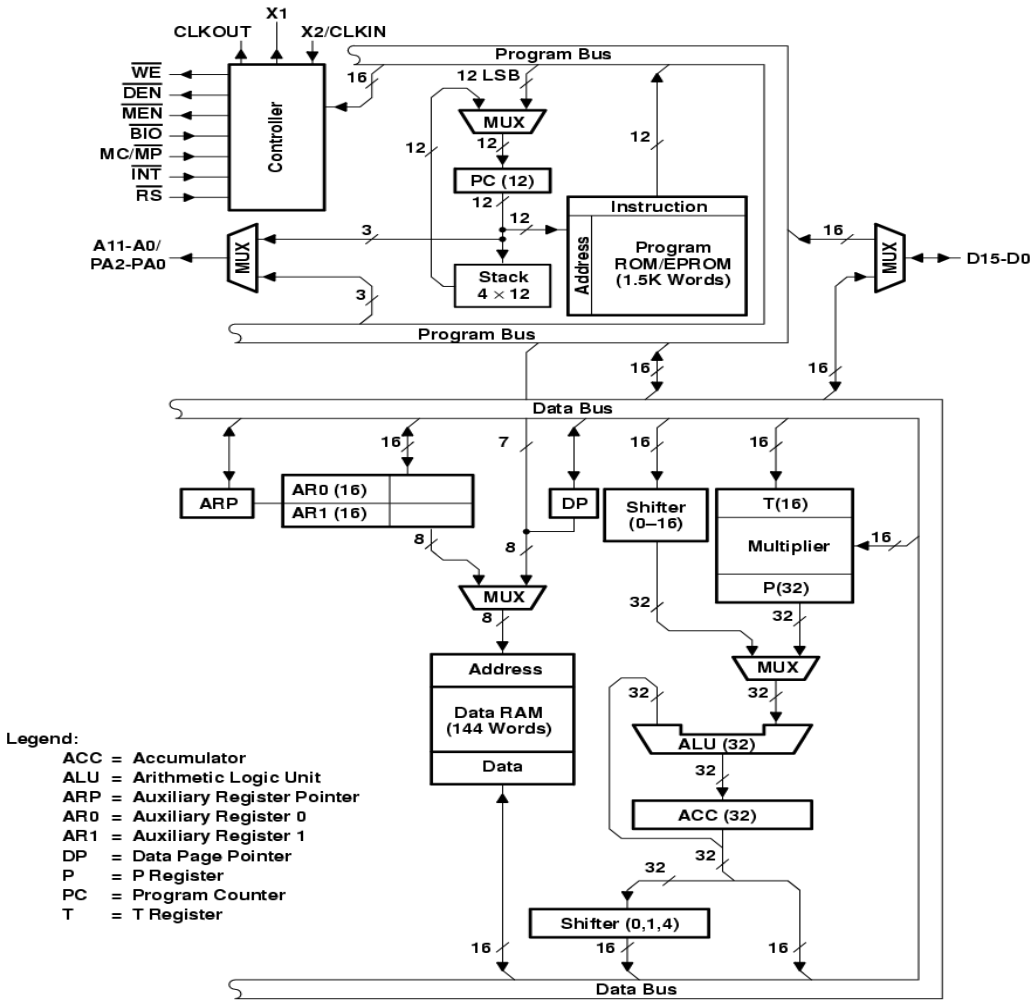

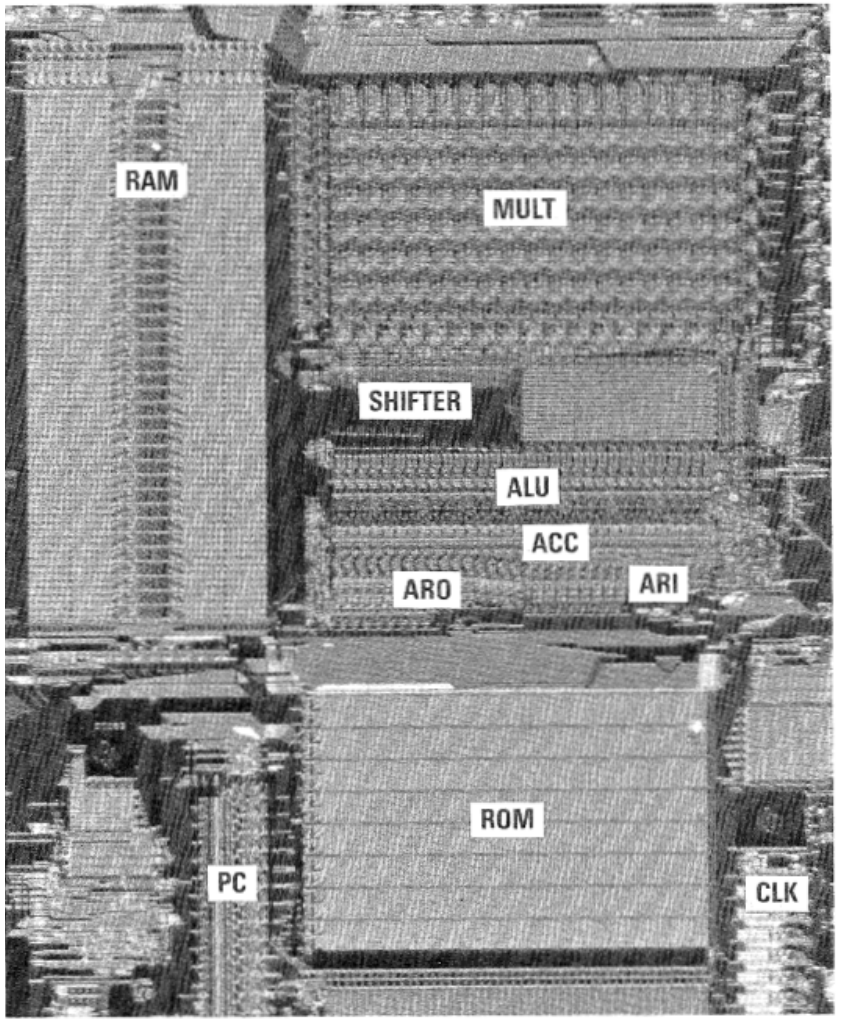

TMS320C1x, C25, C5x features

	‘C1x (1982)	‘C25 (1986)	‘C5x (1989)
MAC	2 cycles @ 5 MHz	1 cycle @ 10 MHz	1 cycle @ 20 MHz
ALU	No guard bits, 1 acc	No guard bits, 1 acc	No guard bits, 1 acc+buf.
Z^{-1}	Physical data move	Physical data move	2 circular pointers
FFT	No bit reverse	Bit reverse	Bit reverse
Bit manipulation	ALU/accumulator	ALU/accumulator	Parallel logic unit
Memory	4 Kw program 144 w data RAM	3 x 64 Kw space 544 w data RAM	3 x 64 Kw space 1 Kw DARAM 9 Kw SARAM
Addressing	1 data acc./cycle 2 pointers + ARP Immed.: load/mpy	2 data acc./cycle 8 pointers + ARP Immed.: all instr.	(same)
Program control	Branch on not zero	Repeat single instr.	Repeat block (1 level) Delayed instructions, XC
Power modes Clock	None Ext. clock / 4	External bus hold (stops CPU) Ext. clock / 4	Idle1 (CPU stopped) Idle2 (CPU+periph) PLL up to x9

Main competition in the 80's

ADI 2101 (1988)	AT&T DSP16A (1988)	Motorola 56001 (1987)
10 MHz	20/30 MHz	10 MHz
16/40 bits, 2 acc	16/36 bits, 2 acc	24/56 bits, 2 acc, add-sub-mpy (FFT)
2 x 16 Kw space 2 Kw P/D DARAM 1 Kw D SARAM	64 Kw P/D space 1 or 2 Kw RAM	3 x 64 Kw P/X/Y space 2 x 256 w X/Y RAM
Circular pointers, bitrev	1 circular pointer, no bitrev	Modulo addressing, bitrev
Hardware loops (4 levels) I-cache All opcodes conditional	Hardware loop (15 words) I-cache	Hardware loops (stack)
Still being produced and sold (25 MHz), although not at a competitive price. Last family member (ADSP-21990) introduced in 2002.	AT&T → Lucent → Agere → LSI Logic, DSP products discontinued	Symphony line of 24-bit audio DSPs (56xxx). New products introduced until 2009. Announced in 2012 there would be no new standalone DSPs.

Floating-point vs. fixed-point

- Wider dynamic range, same precision through range.
- Balance of cost, performance (cycle time), power consumption, ease of programming and time-to-market.
- Floating-point DSPs better suited to:
  - High-end audio
  - Medical
  - Radar
  - Industrial control and robotics

Floating-point DSPs in the 80's

- Hitachi HD61810 (1982): very first one, 16-bit numbers (m12e4).
- NEC 77230 (1985): first practical one; 32-bit storage, 55-bit mpy, 250 ns cycle time.
- TMS320C30 (1988), AT&T DSP32C (1988), Motorola 96002 (1989)
- Some TMS320C30 features:
  - 60 ns cycle time, 16 MW address space, 8 data registers, regular instruction set, instruction cache, DMA controller, guard bits...

Ease of programming – ’C30/’C50

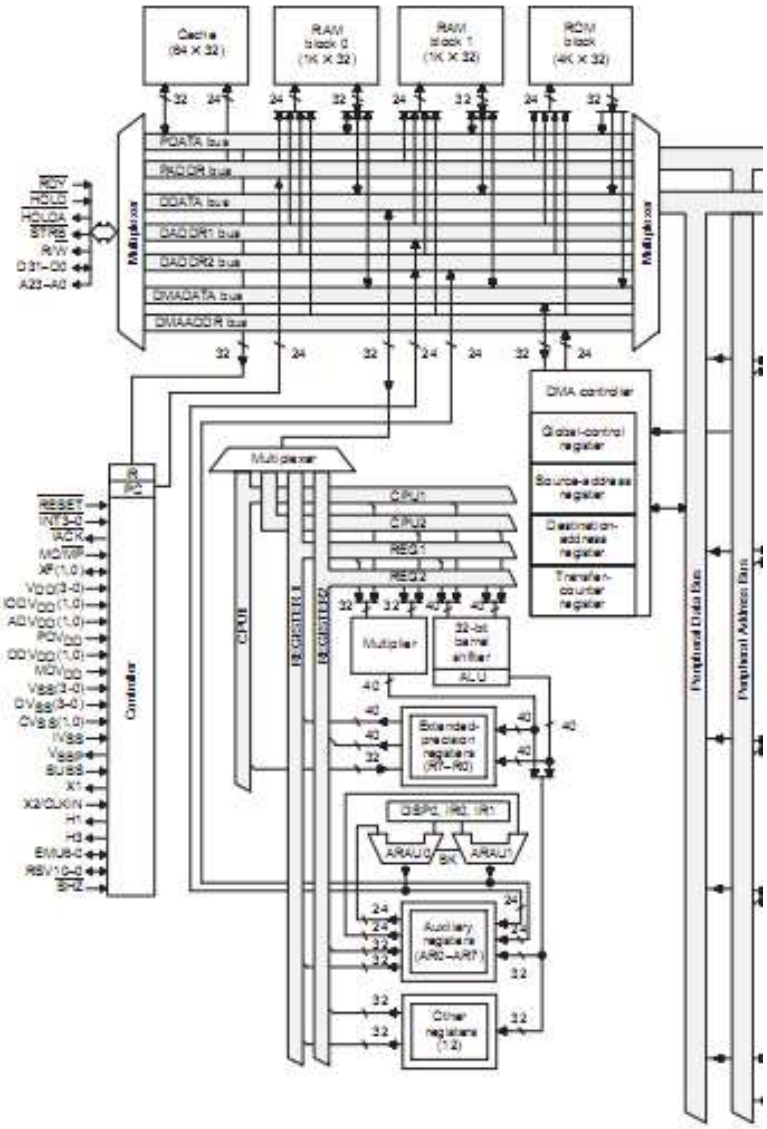

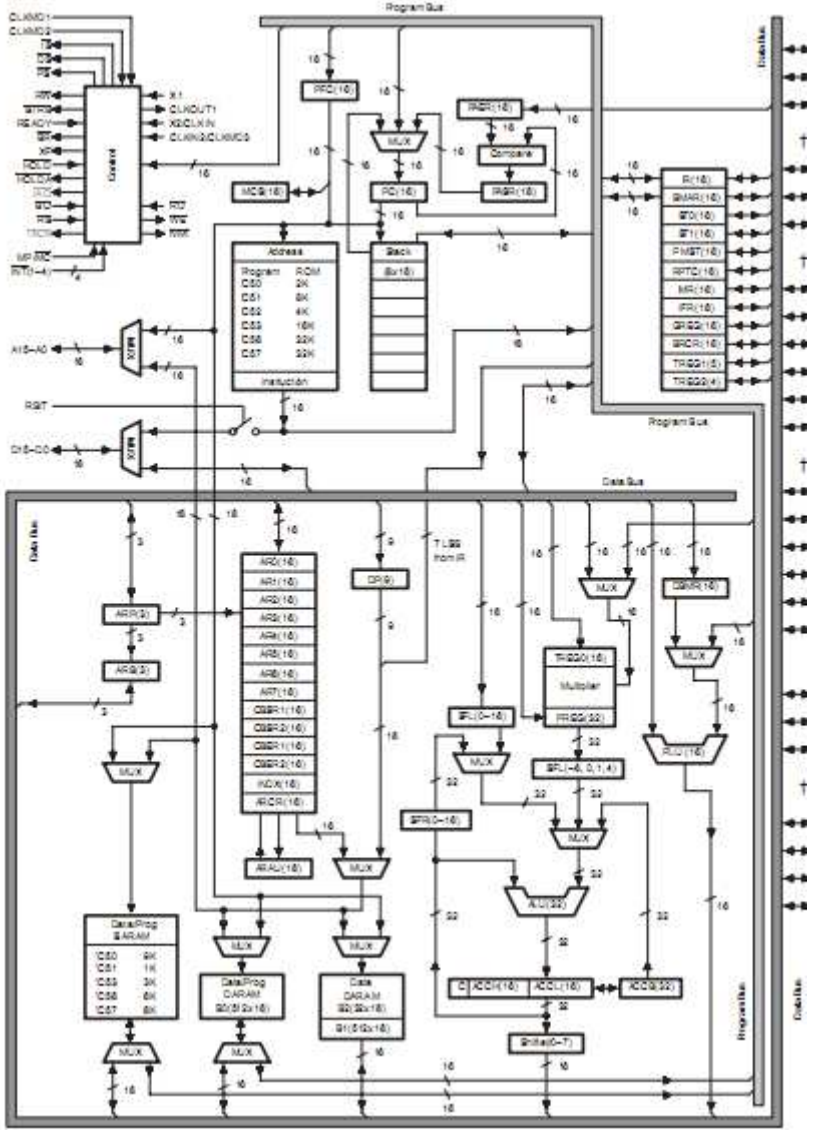

Early 90’s: The great divide

Emerging DSP markets and enablers

- Multimedia, image encoding
  - MP3: 1993, MPEG-2: 1996, MPEG-4: 1998
- Mobile communications, base stations
  - First GSM pilot network in 1991
- Digital control, enabled by lower costs
  - 1994: TMS320C2x, ADI 210x: $10 (1k units)
- Increased transistor budget
  - TMS32010 (1982): 57,000 transistors
  - TMS320C50 (1989): 1M transistors
    - Cray-1A CPU (1976-1979) was 2.5M
- Lower power consumption
  - ‘C50 was 1.8x lower power (typ. @ 5V) than the TMS32010 while featuring a 4x instruction rate.

The great divide

- Pervasive DSPs
  - Entry-level, variety of on-chip peripherals, general-purpose capabilities eliminating MCU, low cost...
- Mobile DSPs
  - Mobile communications standards setting stringent requirements (interrupt response time, bit-exactness...), power consumption (energy/function, not power per MHz), digital consumer applications.
- High-performance DSPs
  - Multiple on-chip execution units, multiprocessing, VLIW, superscalar, large on-chip caches...

Other major DSP trends

- Efficient C programming, hybrid MCU/DSP capabilities.
- Emergence of IP vendors
  - ARM, ARC (now Synopsys), DSP Group (now CEVA), Tensilica...
- Core-based design: cDSP, C2xLP, Carmel...
- Specialized hard-wired blocks and accelerators
  - MPEG-2/4...
  - Hard disk drive read channel
  - ADSL, xDSL
  - Channel equalization (CDMA...)
  - etc.
- Software “ecosystem”, development tools

A look at software for TI DSPs

- 1987: “Digital Signal Processing Applications with the TMS320 Family” textbook
- 1990: TI C compiler and source debugger, DSP Starter Kit (DSP, EPROM, audio AD/DA)
- 1991: TI sponsors the first Educators’ Conference for DSP educators and researchers
- 1993: TMS320 Software Cooperative
- 1995: On-Line DSP Lab, WWW DSP hotline
- 1997: TI acquires Spectron Microsystems (SPOX [stdio, DSP math pkg], BIOSuite) and GO DSP (Code Composer Studio)
- 1998: Real-Time Data Exchange (RTDX)
- 1999: eXpressDSP (CCS, DSP/BIOS, APIs)
- MATLAB/Simulink, SPW, VisSim/Embedded Controls Developer

Some more DSP algorithms – CELP

- Codebook Excited Linear Predictive coding
- Analysis-by-synthesis
- Codebook contains random white noise sequences
- Use of a perceptual filter so coding error noise falls below speech spectral envelope.

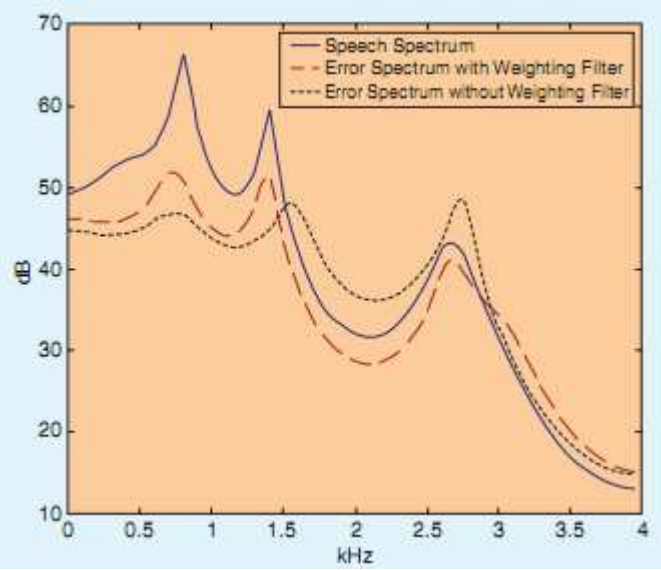

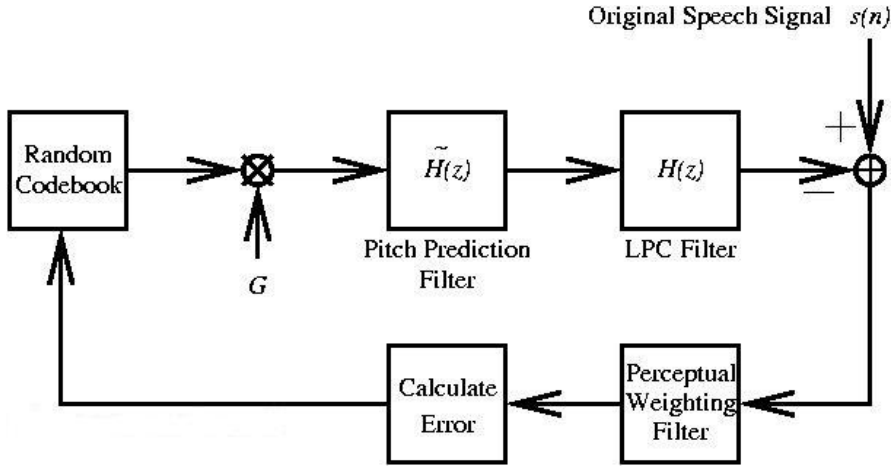

Image sources: Jerry D. Gibson, Speech coding methods, standards, and applications (2005), www.data-compression.com

CELP variants

- Required to reduce computational complexity.
- VSELP (Vector Sum):
  - Code vectors are linear combinations of a few basis vectors.
  - Americas D-AMPS IS-641 (1990, 14 MIPS), RealAudio 1 (1995)...
- ACELP (Algebraic):
  - Large codebook containing very few pulses of ±1 amplitude.
  - GSM EFR (ETSI, 1995, 14 MIPS), AMR (3GPP, 1998 , 14 MIPS)...
- G.729 (1995, 12-20 MIPS):
  - Adaptive codebook to encode past residuals (step 1) then fixed codebook.

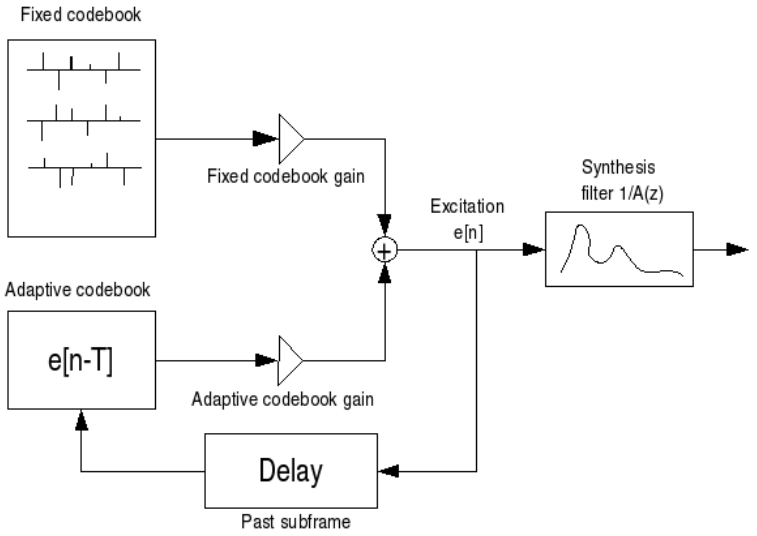

Image source: Wikipedia article on CELP

MPEG-1/2 audio layer 3 (MP3, 1993)

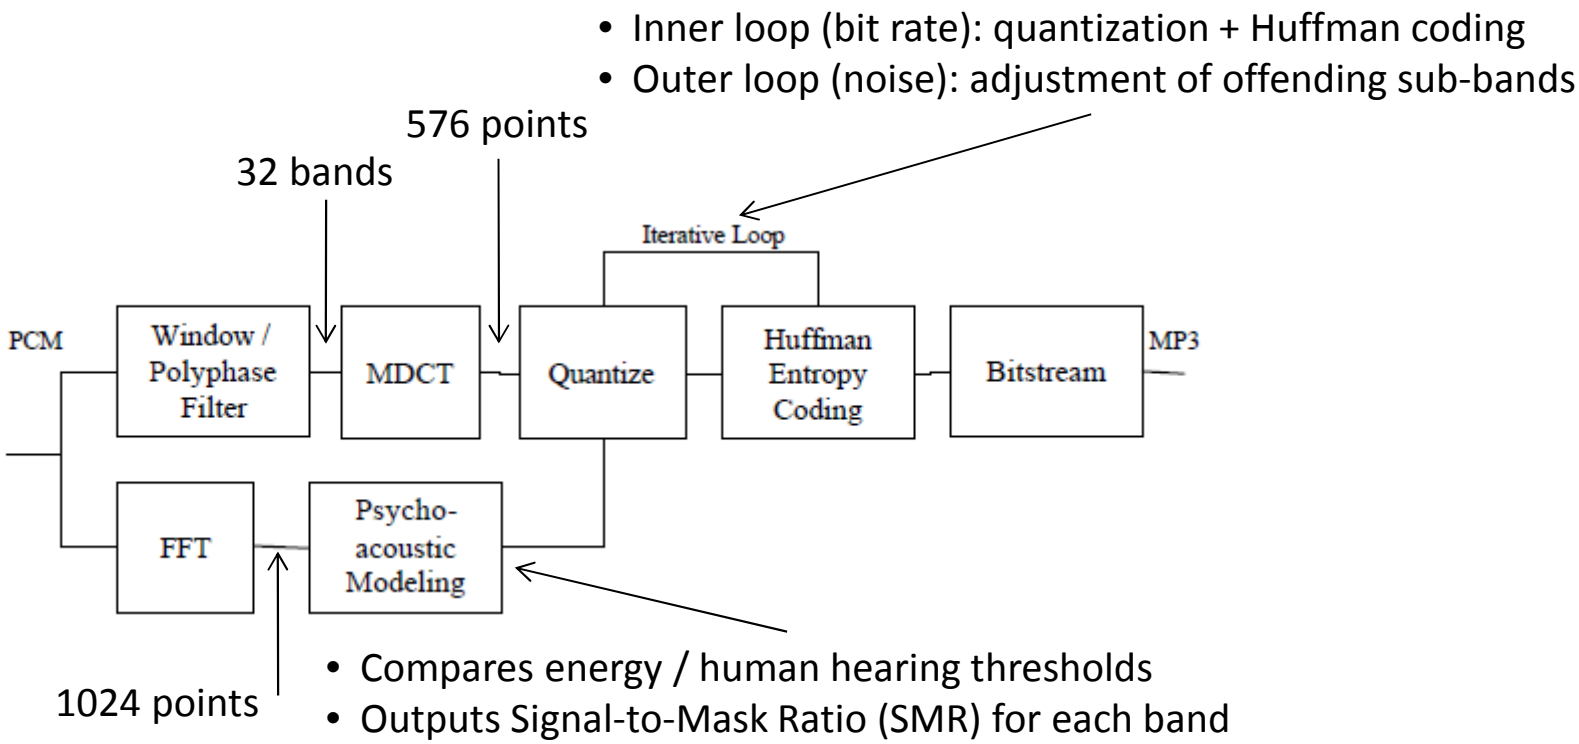

- 23 MIPS encode, 12 MIPS decode (‘C55x, 128 kbps)

Figures of SPIRIT DSP implementation – www.spiritdsp.com

Still-image compression

- Transform
  - 'C55x, CIF†, 30 fps:
    ~40 MHz
    ~10 MHz (HWA)
- Quantization
  - Reduces number of bits per coeff.
- Run-length coding
- Huffman coding

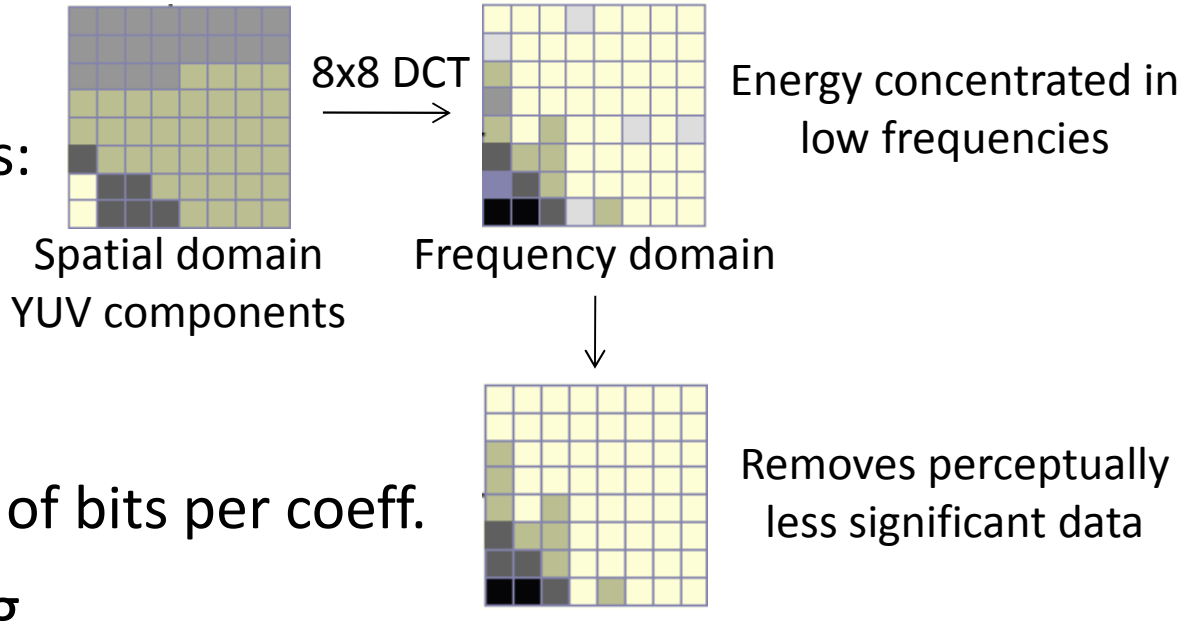

- Still-image encoding alone achieves ~10:1 compression.

Full tutorial available at http://www.bdti.com/articles

†CIF: Common Intermediate Format, 352 x 288 pixels

Motion estimation

- Intra- and predictive-coded frames.
- Image divided in 16x16 pixels macroblocks.
- Process:
  - Search ref. frame for a 16x16 region matching macroblock.
  - Encode motion vector.
  - Encode difference between predicted and actual macroblocks.
- Cannot perform exhaustive search.
  Strategy of motion estimation is a key differentiator.
- Sum-of-absolute differences.
  - 2/3 of encoder MIPS spent there.

Full tutorial available at http://www.bdti.com/articles

Image post-processing

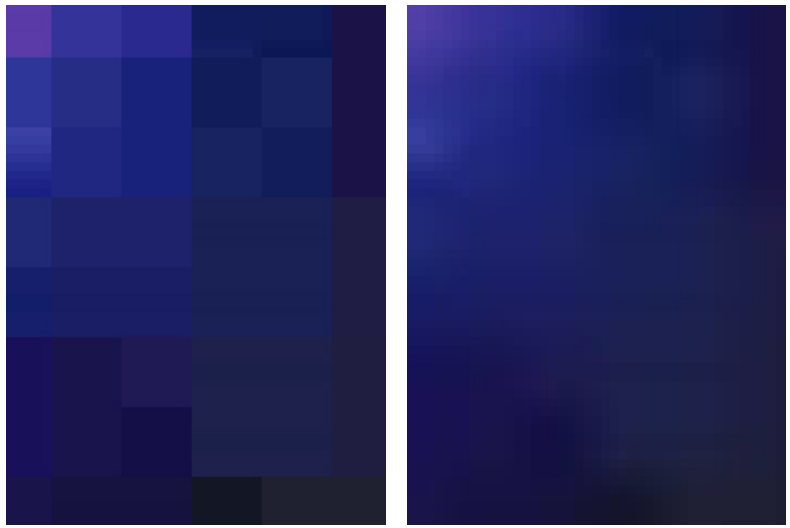

- Removing artifacts:
  - De-blocking (see images)
  - De-ringing: remove distortions near edges of image features.
- Color-space conversion:
  - YUV → RGB
- Very intensive since performed at pixel level.
  - e.g., YUV → RGB takes ~36 MIPS ‘C55x (CIF, 30 fps)

Control-oriented peripherals

- Pulse Width Modulation
- Dead-band control
- Quadrature Encoder

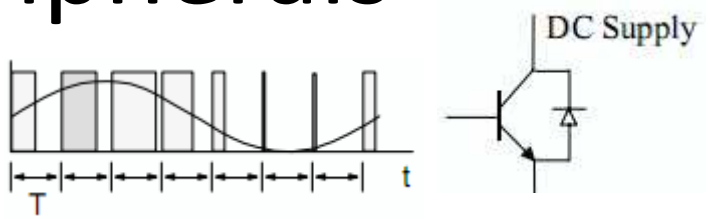

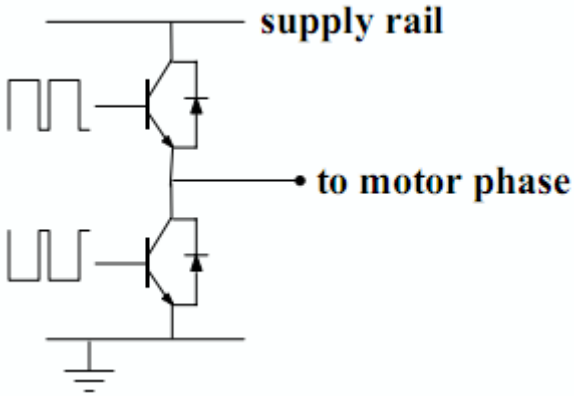

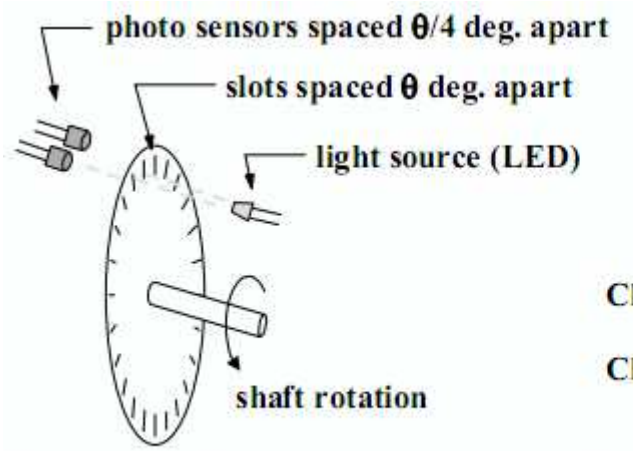

Incremental Optical Encoder

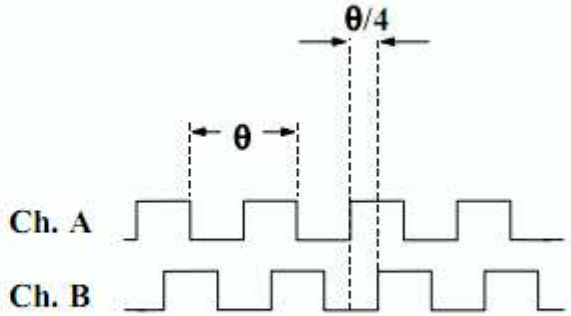

Quadrature Output from Photo Sensors

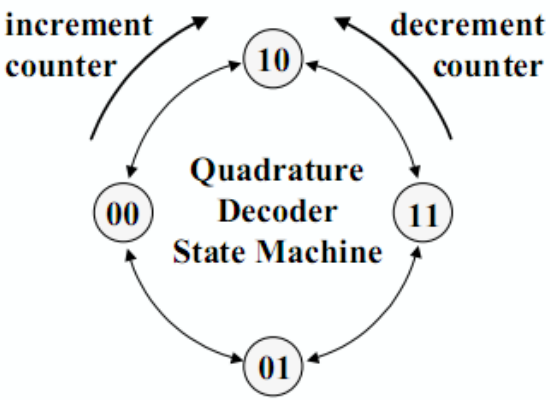

Images source: C28x DSP Design Workshop, Texas Instruments, 2004

Benchmarking DSP performance

- BDTI DSP Kernel benchmarks
  - FIR, LMS, IIR, FFT, vector dot product/add/max, Viterbi, control, bit unpack
- BDTI Communications benchmark (OFDM)
  - IQ, FIR, FFT, slicer (FFT to QAM constellation), Viterbi
- BDTI Video Encoder, Decoder and Kernel benchmarks
  - Deblocking, DCT, motion compensation and estimation, image resize
- Also at BDTI:
  - Comparison high-level synthesis on FPGA and RTL or ‘C64x+

Digital Signal Processors 90’s-00’s

Pervasive DSPs

- 1988: TMS320C14
  - UART, timers, capture inputs, compare outputs (PWM), EPROM
  - Positioned towards motor control and automotive (ABS...)
- 1991: $5 for a TMS320C1x DSP (1 kunits)
- ‘C2xx: C2xLP core-based, JTAG emulation
  - 1996: 6 products; flash; UART, timers, GPIO
  - 1997: ’F240 aimed at motor control

TI hybrid DSP/MCUs

- 1998: ‘C27x
  - Code efficiency: better density, better compiler target
    - 15-20% advantage on ARM7 in mass storage benchmarks
  - 4M 16-bit words address space, byte addressability, stack-based addressing
  - Improved interrupt and context switch responses
  - Memory-to-memory and register-to-register operations
  - Real-time emulation: DMA, real-time data exchange...
- 2001: ‘C28x
  - ‘C2xx code compatible, 32-bit arithmetic, dual 16x16 MAC
  - A/D converter, CAN, SPI, SCI, I^2C, LIN, McBSP...
  - High-resolution PWM (65 ps), quadrature encoder (QEP)...
  - “Virtual floating point” (IQ Math), 1-to-1 C-to-assembly ratio

TI hybrid DSP/MCU current lineup

- Not a DSP anymore!
  ‘C2000 appears in the TI MCU product tree.
- Floating-point products since 2008, 50-300 MHz
- 80+ code-compatible products
- Entry price: $2.00 (1,000 units)
  - 40 MHz, 16/6 KB flash/RAM, 8 PWM, 2 MHz 12-bit ADC
  - Similar price as ARM Cortex M3 microcontrollers (same features)
- Applications:
  - Digital motor control, automotive (hybrid, power steering, x-by-wire...), renewable energy, lighting, power line comms, precision sensing and control...
- Competition: Analog 2199x, Freescale 5685x, boosted MCUs (Atmel, Infineon, Microchip dsPIC)

TI Mobile DSPs

	'C5x (1989)	'C54x (1995)	'C55x (2000)
ALU	No guard bits 1 acc+buf. Pipelined MAC	8 guard bits 2 accumulators Dual 16-bit ALU 32-bit operands Separate MAC	4 accumulators Dual MAC Additional 16-bit ALU
Specialized func.	Parallel logic unit	Compare-Select-Store Exponent encoder	More orthogonal instr. set
Addressing	1 address gen. 2+1 data read+write 64 kw program	2 address gen. 2+1 data read+write Up to 8 Mw program Parallel load/store Conditional store	3 address gen. 3+2 data read+write 16 MB program/data More circular addr.
Program control	Repeat block (1 level) Delayed instr., XC	(same)	Repeat block (3 levels) 64-byte instr. queue Speculative fetching
Power modes	Idle1 (CPU stopped) Idle2 (CPU+periph)	Idle3 (CPU+periph+PLL)	User-configurable idle domains
Other			Up to 40% denser code

Mobile DSPs timeline

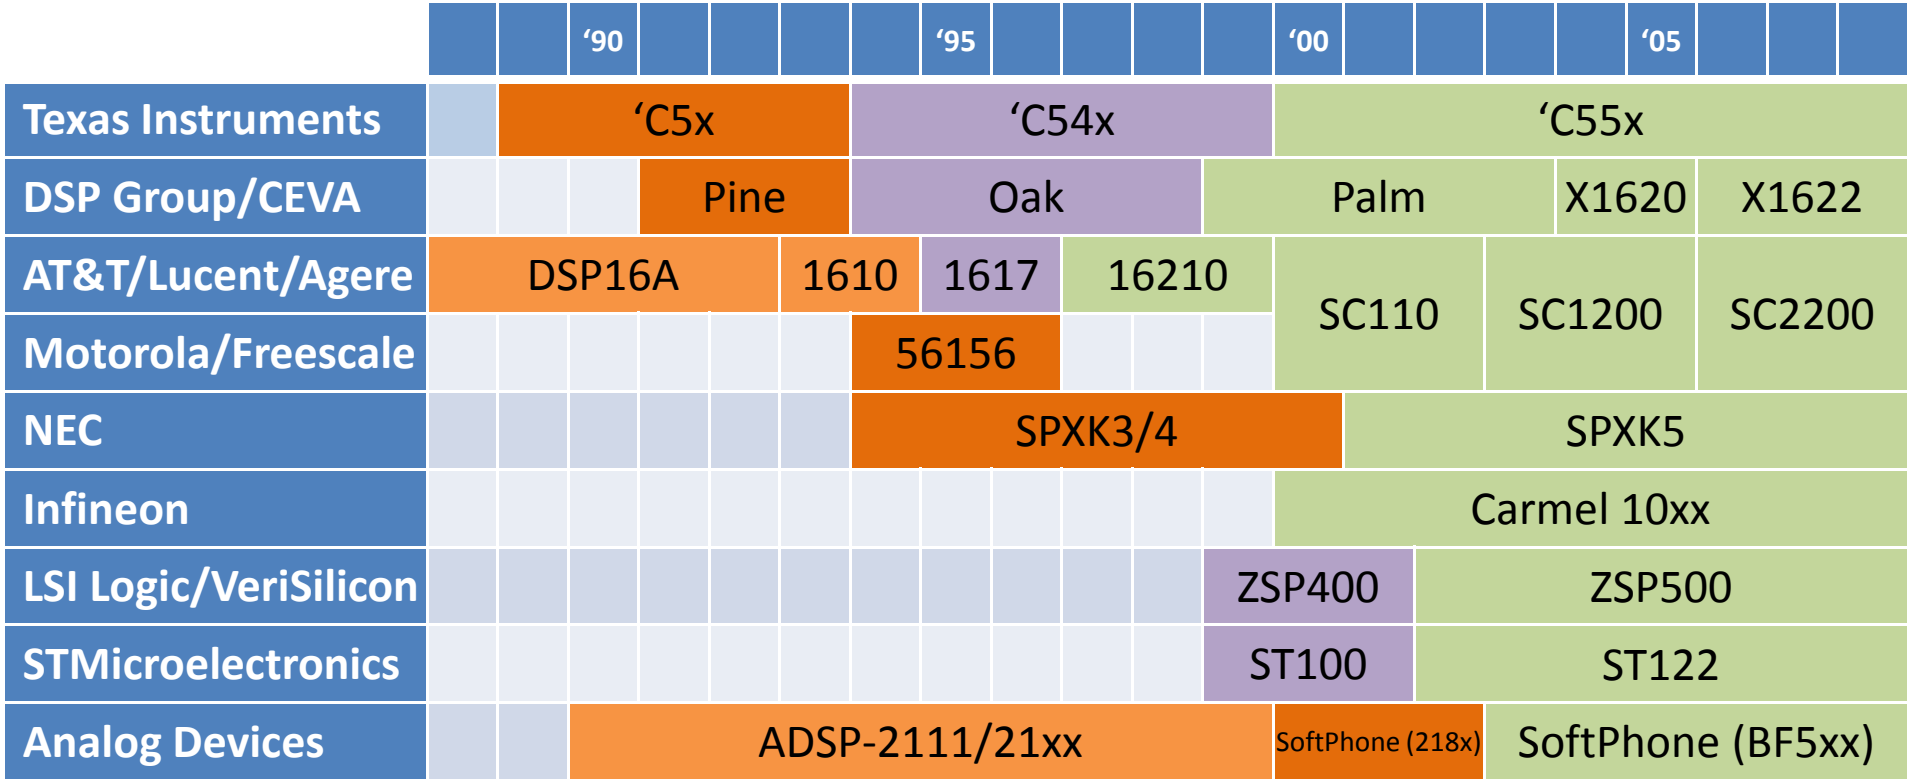

- Today:
  - Qualcomm: in-house DSP
  - Intel: CEVA
  - MediaTek: CEVA, but acquired Coresonic in 2012
  - Japan: Tensilica

High performance DSPs

- 1987: TMS320C30 is 45 kgates plus memory.
- Massively Parallel Processing (MPP) starts in the 80's.
  - e.g. nCUBE 10 (1985, 1024 processors).
- Some processors are designed with high-speed links.
  - Inmos Transputer T212 (1984, 16-bit fixed point)
  - TI TMS320C40 (1990, 32-bit fp)
  - ADI ADSP-2106x "SHARC" (1994, 32-bit fp)
- Early '90s: silicon budget starts allowing multiple execution units.
  - TMS320C60 "Juggernaut": single-cycle complex multiply.
    - Project redirected to the TMS320C6000.

High performance DSPs timeline

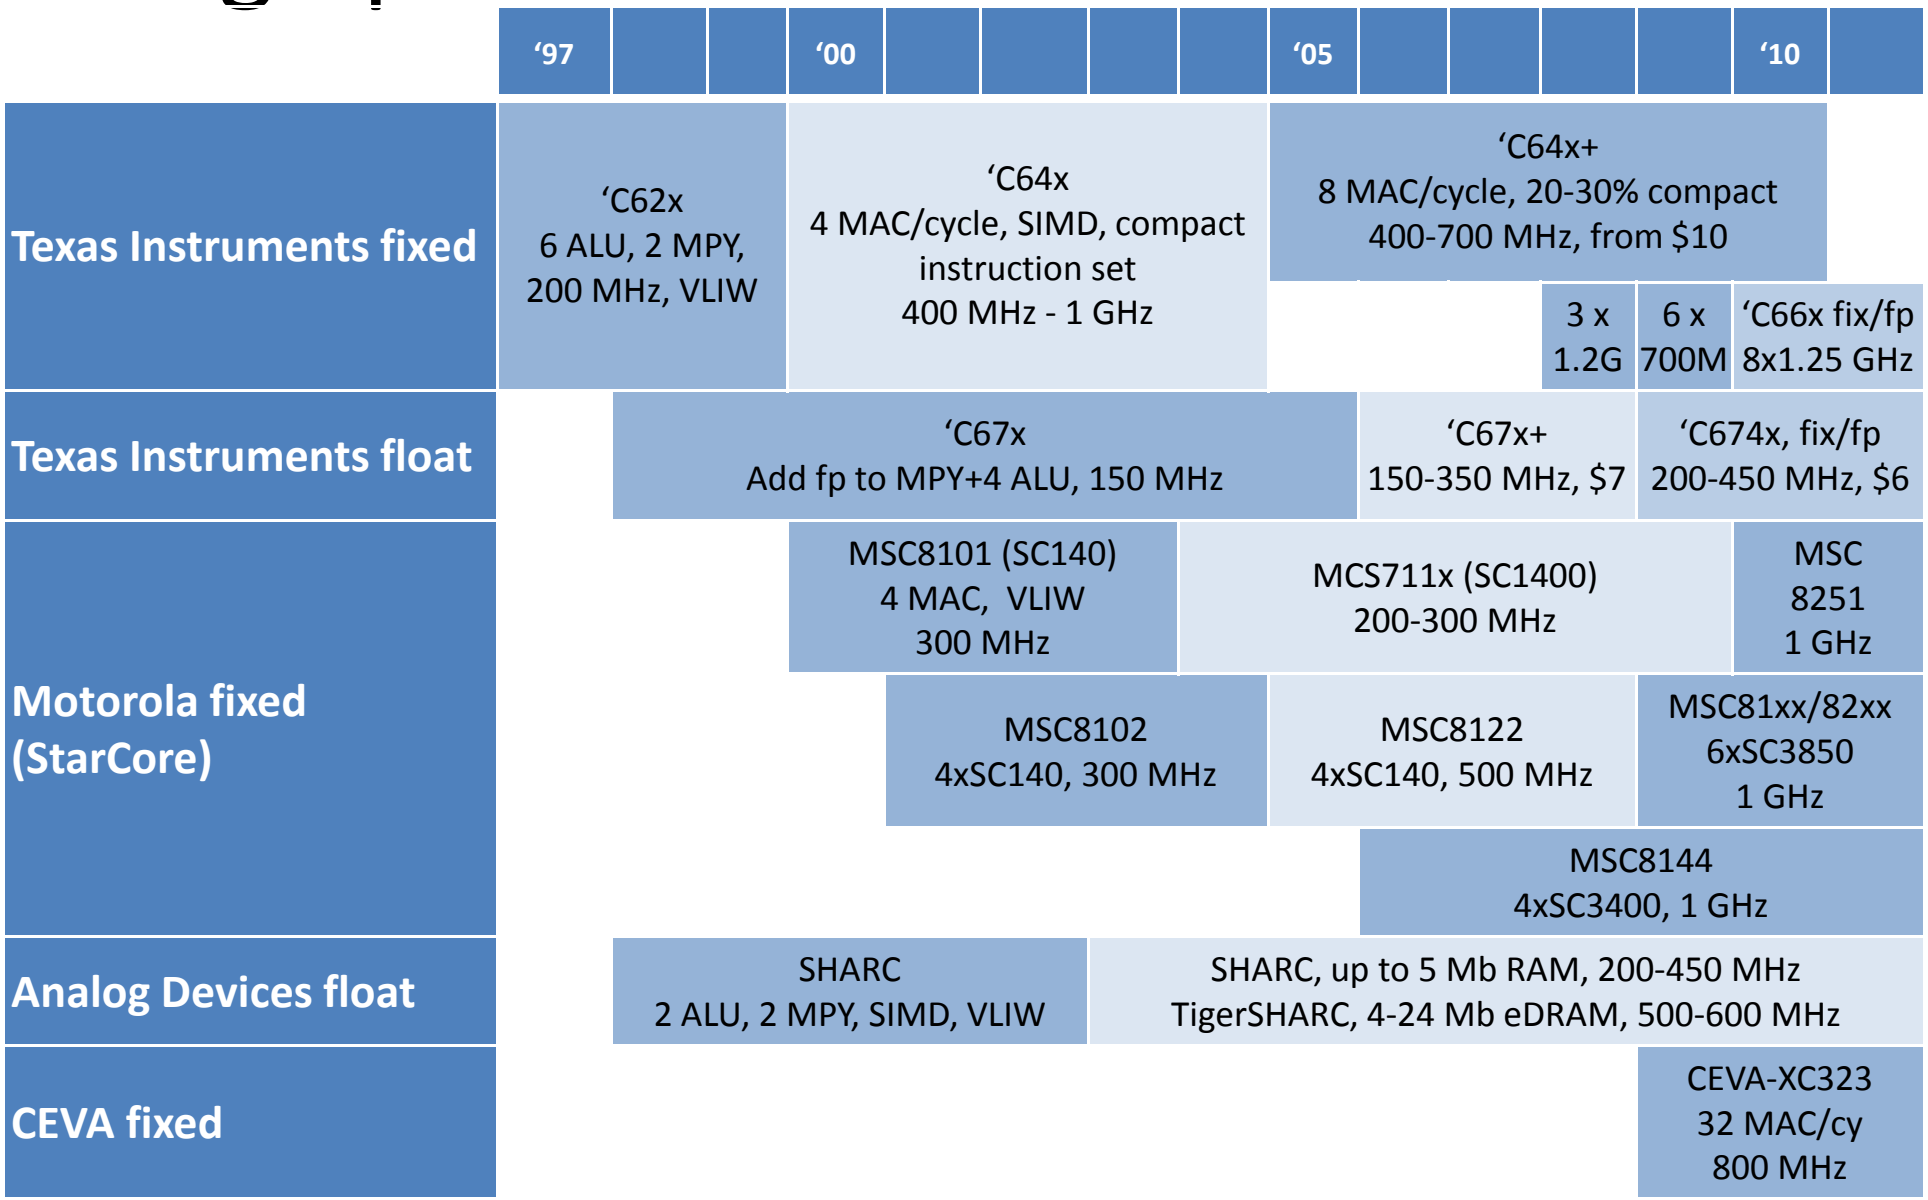

Specialized DSPs from TI

- 2001, TMS320DA250, digital audio
  - ‘C55x @ 120 MHz, USB, MS/MMC/SD, LCD, I^2C, SDRAM...
- 2001, TMS320DSC21/24, digital still camera
  - ‘C54x @ 250 MHz + ARM7TDMI, SDRAM, NTSC/PAL, 2 Mpix...
- 2001, TMS320IP5472, IP telephony (‘C54x + ARM7TDMI)
- 2003, TMS320DM310, digital media (‘C54x + ARM925)
- 2003, TMS320DM642, DaVinci video/imaging, ‘C64x
- 2008, TMS320TCI6487, wireless infrastructure, 3x’C64x+
- Currently: DaVinci video/imaging, OMAP apps processors, KeyStone imaging/vision/high performance computing

Configurable DSPs – ARC

- 1999: ARC 3, 2 MAC/cycle, 40-bit registers, saturating arithmetic, XY mem, circular and bit-reverse addressing, zero-overhead looping, 40 MHz.
- 2002: ARCtangent-A5, ARCompact instruction set, targets VoIP, wireless baseband, digital imaging and audio.
- 2003: ARC600, 200 MHz, vertical applications (audio...).
- 2004: ARM700, 300 MHz, MMU, dynamic branch prediction...
- 2005: ARM710D, 128-bit SIMD for video algorithms, 533 MHz.
- 2007: VRaptor media architecture. SIMD, coprocessors (ME, entropy enc/dec...).
- 2009: Virage Logic acquires ARC
- 2010: Synopsys acquires Virage Logic
- 2011: ConnX for baseband.
- 2012: HiFi 3 for high-end audio and voice processing.
- 700 millions units shipped annually.

Configurable DSPs – Tensilica

- 2000: Vectra DSP extension to Xtensa III
  - 200 MHz, 40-bit registers, 4 MAC/cycle, SIMD
  - FFT 1.8x faster than TI ‘C55x, Viterbi butterfly in 2 cycles (TI: 4)
- 2002: Xtensa V / Vectra scores 1.6x higher than ‘C62x EEMBC Telecom benchmark (both *optimized*).
- 2003: HiFi audio package (software codecs, new instr.)
- 2004: Xtensa/Vectra LX
  - FLIX (VLIW). 1.8x BDTImark of ‘C64x (8 MAC configuration, Viterbi + bit unpack instr.)
- 2005: video decoder package
  - H.264 D1 decoding 30 fps in 10.5 mm^2 (130 nm).
- 2006: Diamond series of pre-configured processors
  - 545CK DSP (8 MACs) displays twice the ‘C64x+ BDTImark2000 score at same MHz.
- 2009: ConnX Baseband Engine (Diamond 545CK, 16 MACs, new instructions)
  - Tensilica positions its products as customizable dataplane processors
- 800 millions units shipped annually. Claim license revenue bigger than CEVA’s.
- 2013: announced they shipped 2 billions cores. Cadence acquires Tensilica.

ARM attempts at DSP – Piccolo

- 1997. An ARM7 coprocessor adding $^2/_3$ area.
- In ARM's words,
  about the 'C52 performance
- 16x16 MAC, 48 bits accums., dual 16-bit instructions, saturation, hardware loop
- Issues:
  - Register file starvation
  - No ARM/Piccolo signaling
    (need for data, interrupts...)

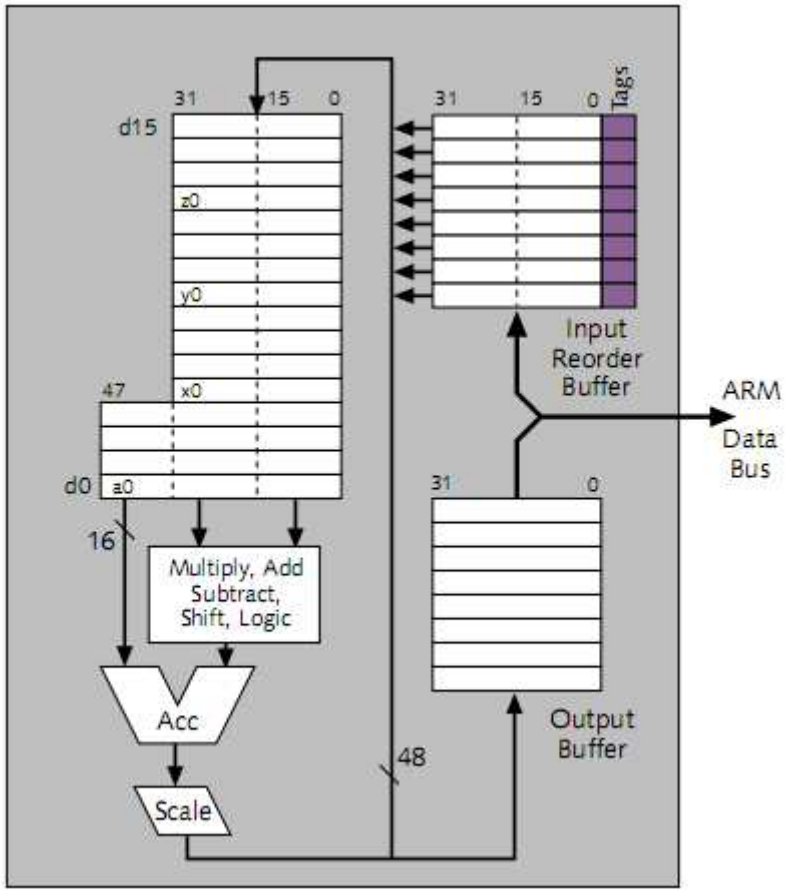

ARM attempts at DSP – v5TE

- 1999. 32x16 MAC, saturation. 30% area adder to v4T.
- Speedup vs. v4T: FFT 20-40%, G.723 codec almost 2x.
- No modulo or bit-reverse addressing, no hardware loop, no guard bits.
- ARM946/966 (1999), ARM926 (2001), ARM968 (2004)
- v5TE maintained in v6 and v7

MIPS	v5TE	‘C54x	‘C55x
Adaptive Echo Cancellation (64 ms)	35.9	12.4	8.6
G.729AB codec	40.3	12.2	10.3

Performance figures source: www.adaptivedigital.com

ARM attempts at DSP – OptimoDE

- 2003, acquisition of Adelante Technologies
- VLIW with Custom Functional Units and configurable micro-architecture

1. C code evaluated on pre-defined configurations
2. Define CFUs
3. Develop microcode using retargetable C compiler

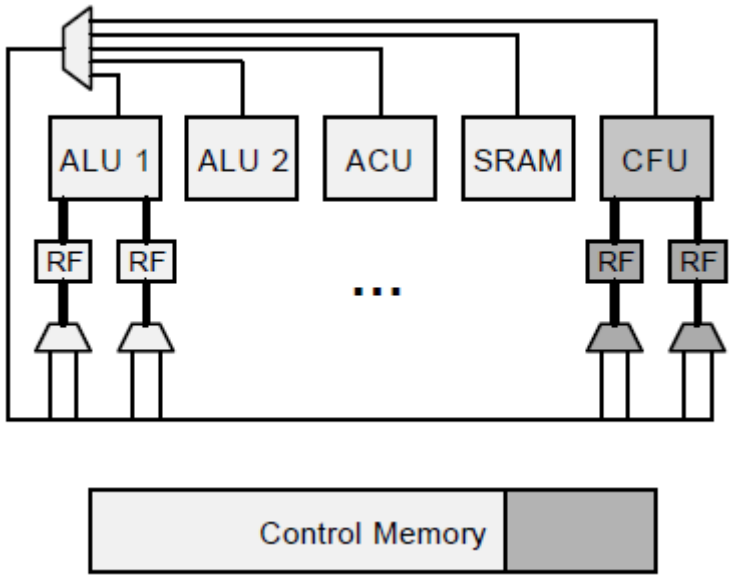

ARM attempts at DSP – OptimoDE

- Limited success: Thomson (video), LG (HDTV), Broadcom (networking, wireless), Phonak (hearing aids), Toshiba (portable).
- Competition: ARC, Tensilica. Patent portfolio, development environment.
- ARM Leuven office closed in 2009.

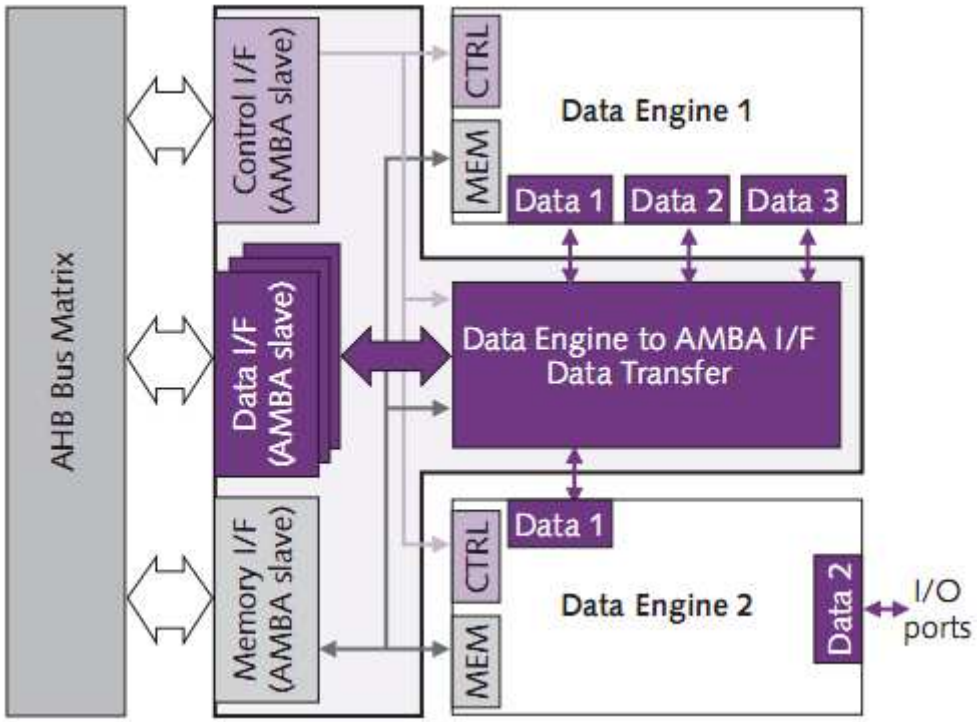

ARM attempts at DSP – NEON

- 2005, introduced on Cortex A8 (ARMv7-A).
- MMX-like (DSP-capable), not DSP-like
- 128-bit fixed/floating-point SIMD
- 32 64-bit registers
- 8/16/32/64-bit integers
- 8/16-bit polynomials with 1-bit coefficients
- Single- and double-precision f.p., IEEE and fast
- Competed with Intel Wireless MMX
  - PXA27x, 2004, 312-624 MHz, 64-bit int. SIMD (16 registers)

ARM attempts at DSP – NEON

- Performance increase *vs.* ARMv5TE:
  - MPEG4 decoding: x4.5
  - GSM-AMR: x3
  - MP3 decoding: x2.5
  - FFT: x4
- Comes at a cost:
  - Cortex A9 core: 600 kgates
  - Cortex A9 NEON: 500 kgates

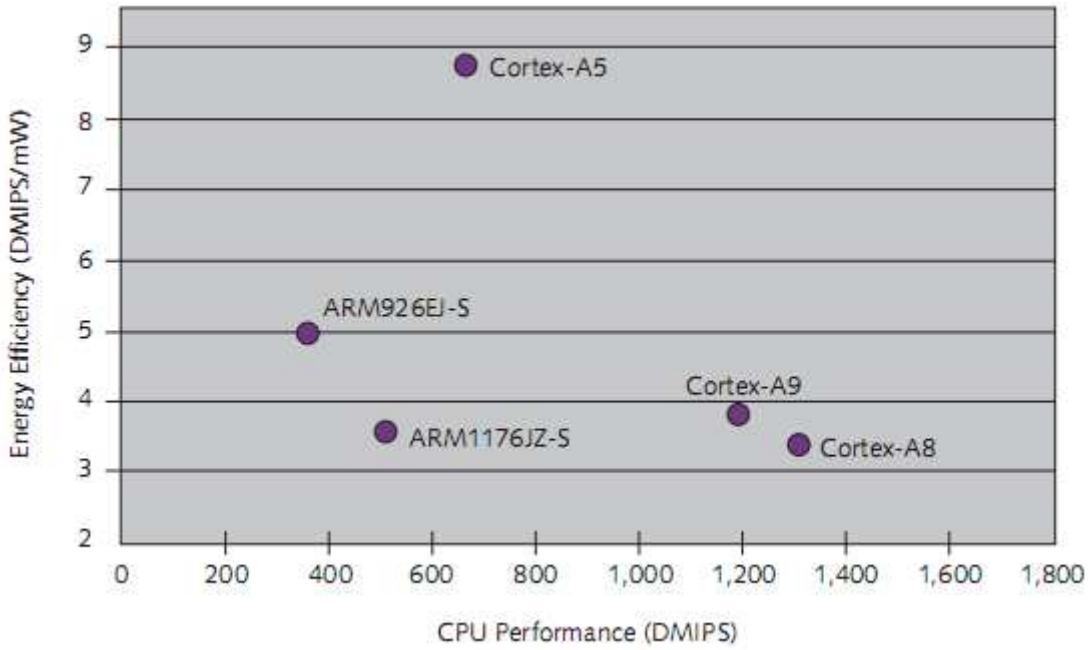

- Cortex A5 introduced in 2009 as a replacement to ARM1176.

ARM attempts at DSP – Cortex M4

- Introduced in 2010 as a Digital Signal Controller.
- Adds single-cycle *à la* v5TE instructions to Cortex M3.
  - 32-bit MAC or dual 16-bit MAC
- Optional FPU, MPU†, NVIC†, WIC†.
- But no bit-reverse or circular addressing, no hardware loops, no parallel load/store and ALU ops.
- Typically 150 MHz (targets flash memory) but capable of 300 MHz (65 nm LP).

†MPU: Memory Protection Unit; NVIC: Nested Vectored Interrupt Controller; WIC: Wake-up Interrupt Controller

Performance comparison

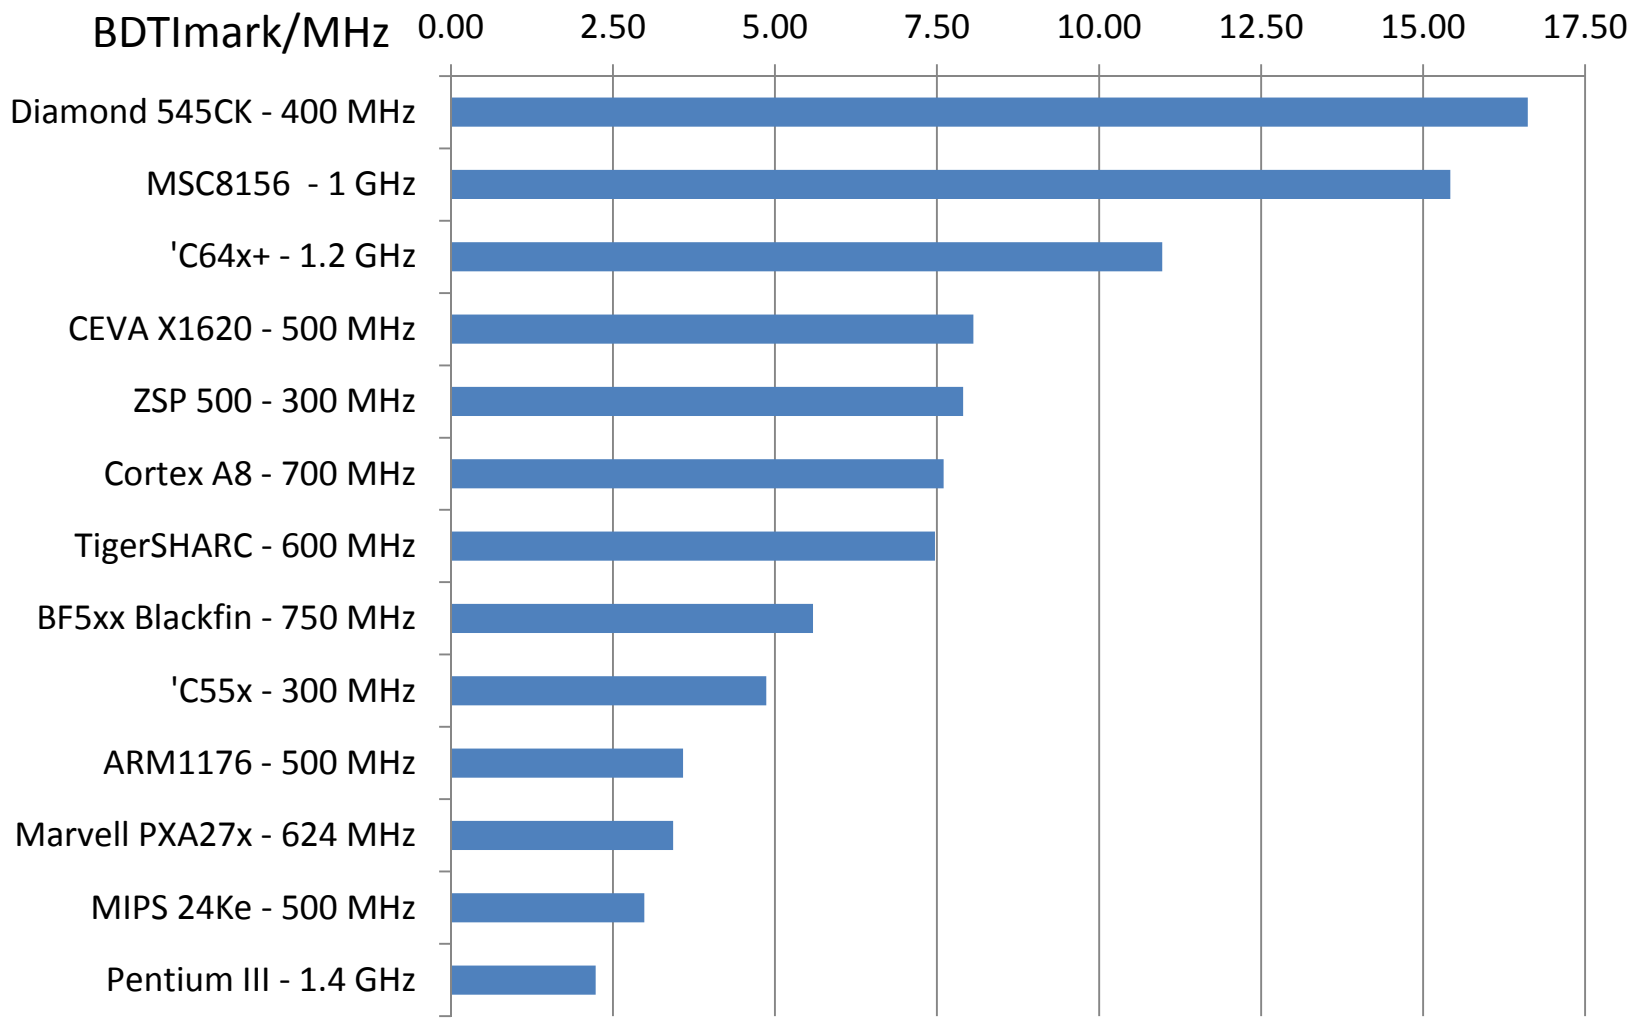

Source: www.bdti.com

AnySP – The best mobile DSP?

- Data-level parallelism analysis for mobile signal processing algorithms:

Algorithm	SIMD workload (%)	Scalar workload (%)	Overhead workload (%)	SIMD width (elements)	Amount of thread-level parallelism
FFT / inverse FFT	75	5	20	1024	Low
Space-time block coding(STBC)	81	5	14	4	High
Low-density parity-check (LDPC)	49	18	33	96	Low
Deblocking filter	72	13	15	8	Medium
Intraprediction	85	5	10	16	Medium
Inverse transform	80	5	15	8	High
Motion compensation	75	5	10	8	High

Source: AnySP: Anytime Anywhere Anyway Signal Processing,
Proc. of the International Symposium on Computer Architecture, June 2009
http://www.public.asu.edu/~chaitali/papers.html

AnySP – Workload analysis

- Multiple SIMD widths, substantial scalar and overhead loads
  - Avoid fixed-width SIMD, improve scalar/address generation/ data shuffling performance.
- Register values lifetimes
  - Bypass register file whenever possible; split register file into a small and a large region to optimize power consumption.
- Instruction pair frequency
  - Fuse most frequent instruction pairs (loses unneeded interim result).
- Data reordering patterns
  - All studied algorithms have a predefined set of swizzle patterns (<10).

AnySP – Architecture

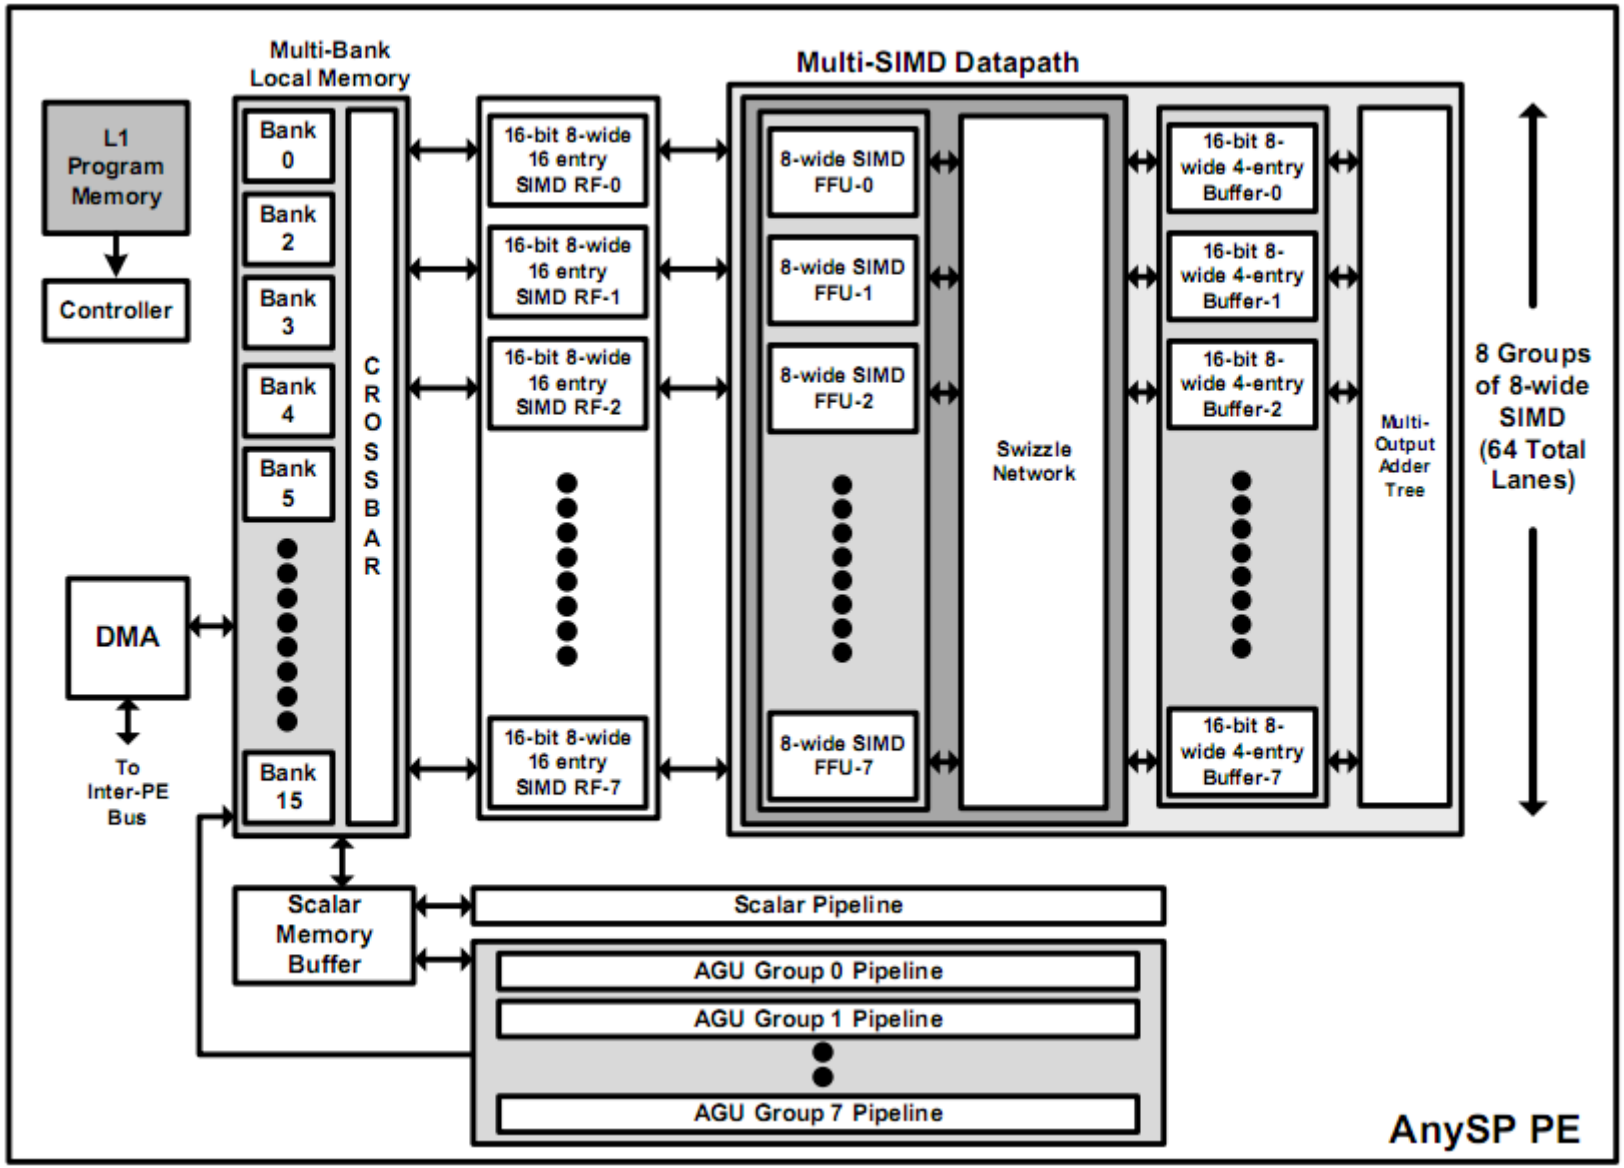

AnySP – Results

- 90 nm core area: 25.2 mm² (est. 6.85 mm² in 45 nm)
- 100 Mbps high mobility 4G wireless:
  - 90 nm, 1 V, 300 MHz: 1.3 W
  - (est.) 45 nm, 0.8 V, 300 MHz: 850 mW – 1000 Mops/mW!
- High quality H.264 4CIF 30 fps decoding: 60 mW (90 nm)

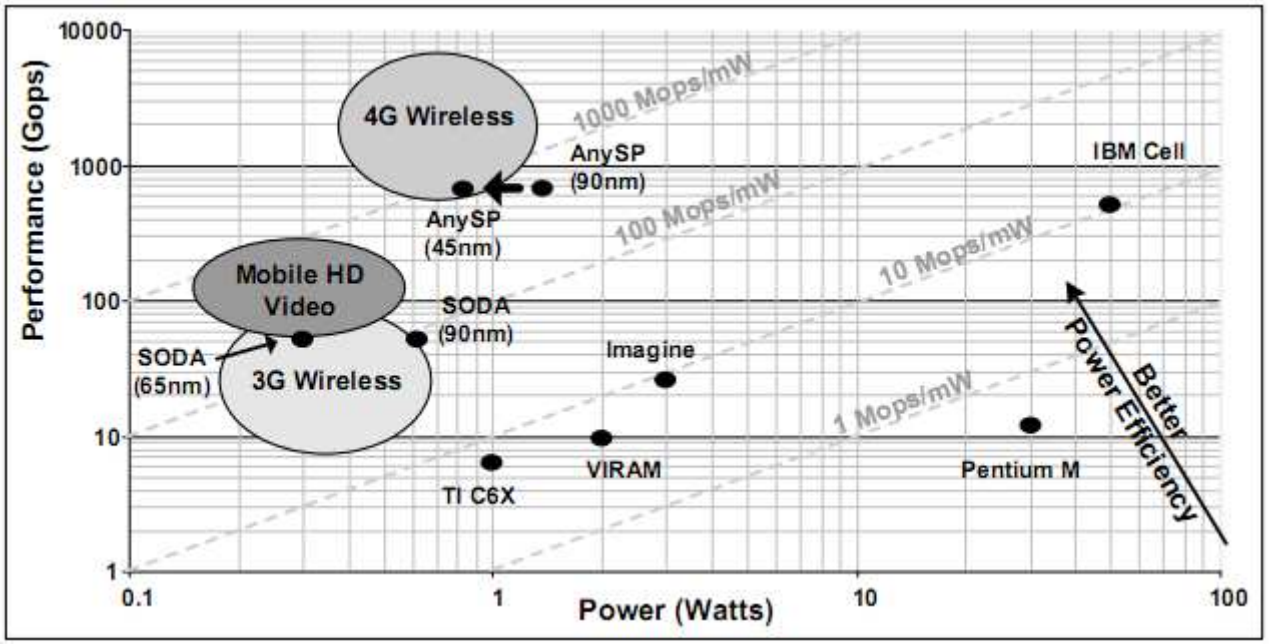

FPGAs in digital signal processing

- 2004: FPGAs appeared in the EDN DSP Directory.
  - Altera DSP Builder interfaces with MATLAB/Simulink; FIR and IIR MegaCores.
  - Xilinx Virtex II: up to 566 18x18 multipliers. Pro version includes PowerPC core. Licensable IPs (Viterbi, Turbo...). Support of MATLAB, Simulink and SPW.
- Markets:
  - Aerospace/defense, broadcast/video/imaging, wireless infrastructure.
- 2005: Xilinx introduces XtremeDSP platform.
  - Defines XtremeDSP slices building blocks (mpy, add/sub...).
  - Low-cost development environment.
- 2011: Xilinx acquires AutoESL.

DSP/FPGA system partitioning

- Altera's view of an OFDMA base station:

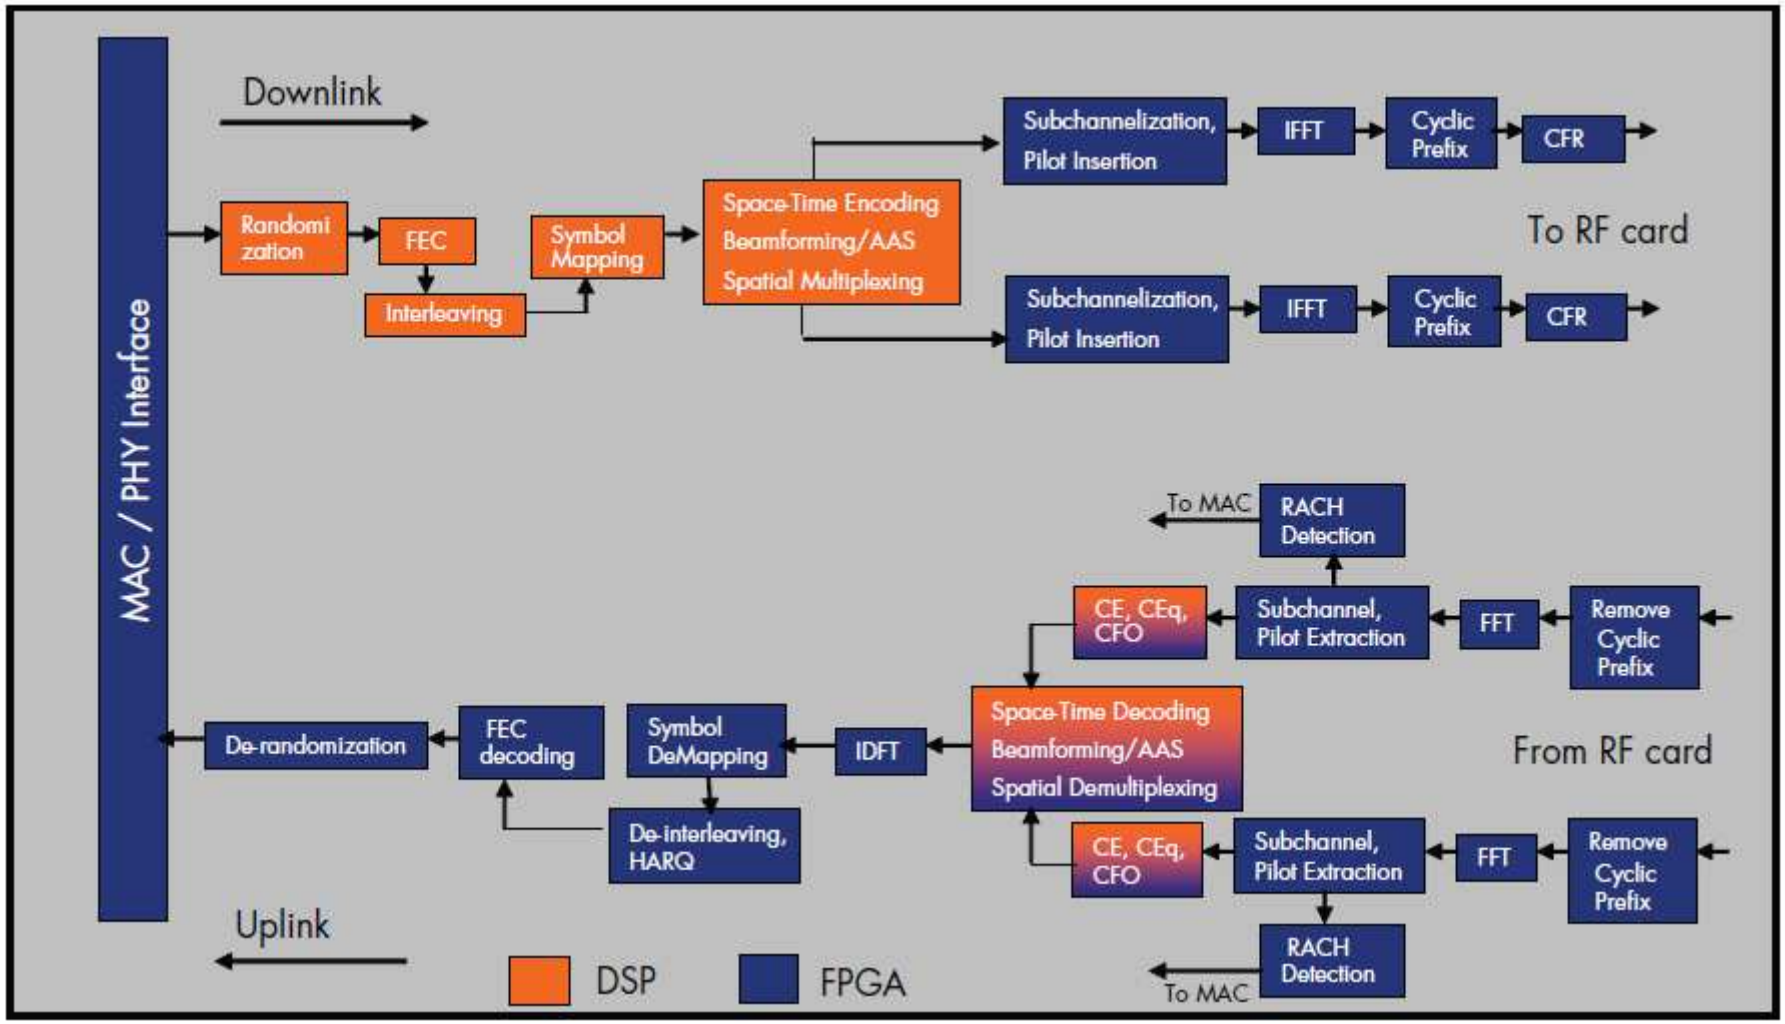

Source: Altera white paper WP-01043-1.0, October 2007

The DSP-centric view (MSC8156)

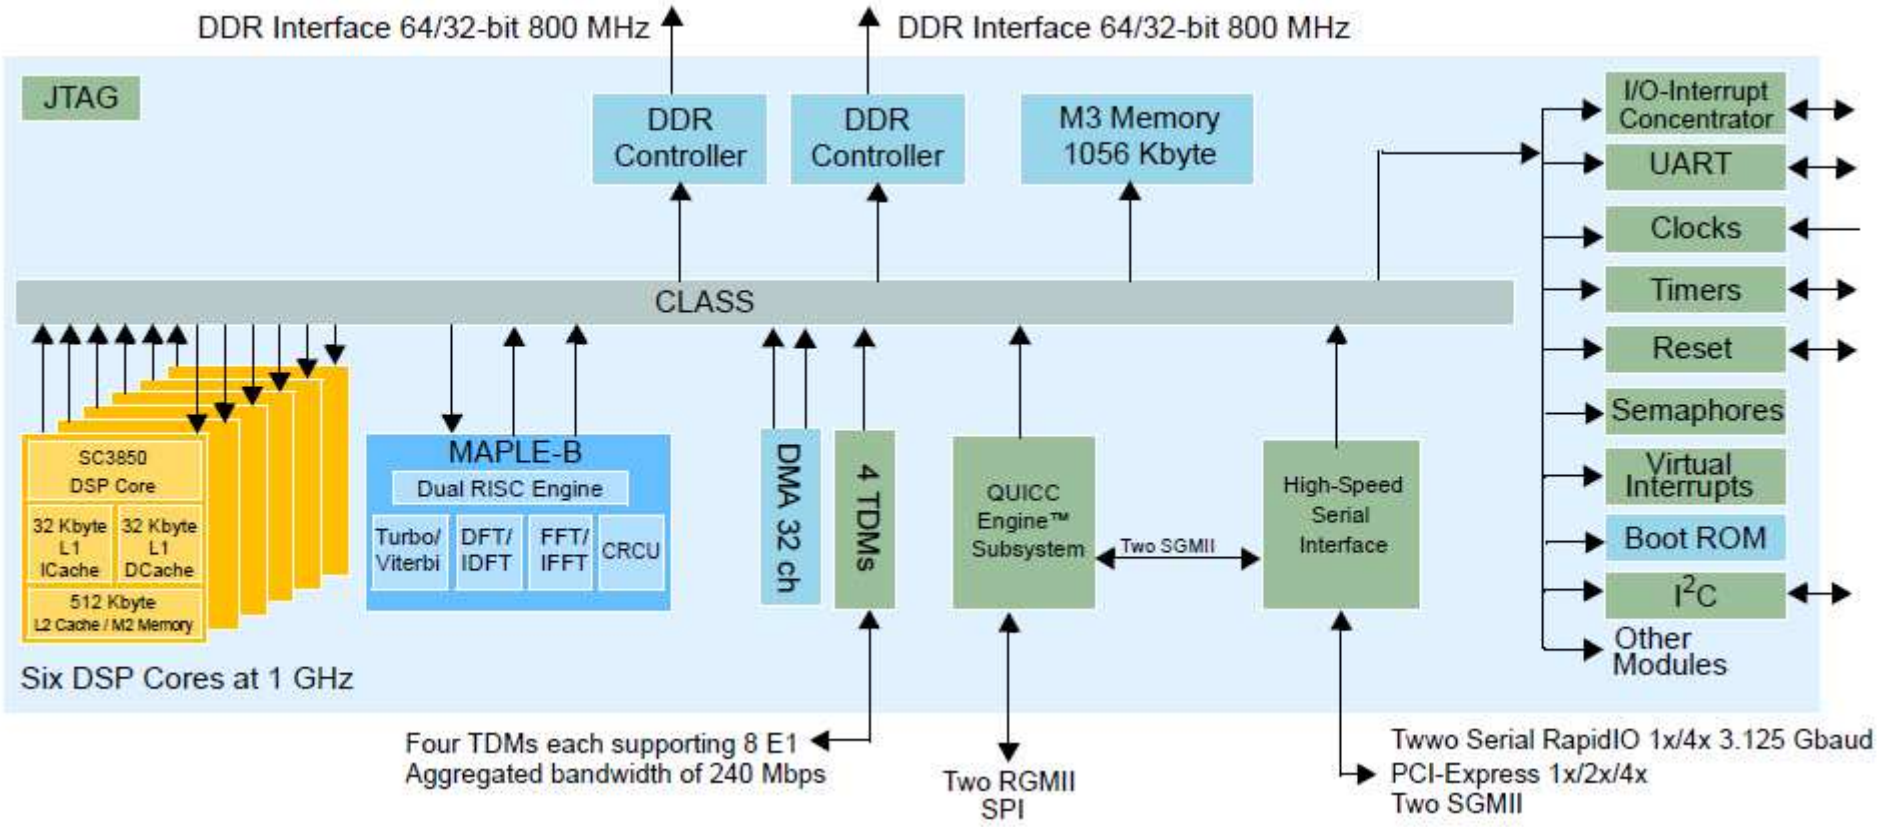

Source: MSC8156 Product Brief, Broadband Wireless Access DSP, Freescale, March 2010

Comparison of DSP implementations

Implementation type	Area	Power	Flexibility	Sw. Dev.	Risk	Remote Upgrade
FPGA	- -	- -	+ +	-	+ +	+ +
Parallel Homogeneous	-	0	0	- -	0	+
Configurable CPU	0	+	-	0	- -	0
CPU + accelerators	+	+ +	- -	+	-	-
Single speed demon	+ +	-	+	+ +	+	+ +

Table shows relative value of implementation by criteria: from - - (poorest) through 0 (average) to + + (best).
Source: Microprocessor Report, “Mixed Architectures Dominate Consumer Sockets”, 11-Aug-2008

Some failed attempts – Mwave

- Mwave: virtual DSP (hard, soft, native).
- PC-based multimedia (modem, fax, voice, audio, video decompression...). Upgradeability.
  - Was to be the DSP “killer application”.
- Initially a TI/IBM alliance.
  - IBM design marketed by TI (TMS320M500, 1992).
  - TI withdrew in 1994 when IBM decided to sell directly the silicon.
- Mwave DSP: 16-bit data, 32-bit ALU/MPY, 25-33 MHz; ISA, MIDI and audio codec interfaces.
- Dropped by IBM because of performance and compatibility issues. Important support burden compared to off-the-shelf sound or modem cards.

Some failed attempts – TMS320C80

- 1994: Multimedia Video Processor
  - 1 Master Processor : 32-bit RISC, IEEE f.p., I/D caches. 100 Mflops.
  - 4 Parallel Processors : 32-bit DSP, 64-bit opcodes, I-cache, local RAM.
  - Transfer controller (400 MB/s). Video controller (DRAM/VRAM).
  - Dual display. 32-bit address space. Crossbar switch. 50 MHz. 2 BOPS.
- Too complex to program efficiently.
  - STI's Cell processor started shipping in 2006...
- Karl Guttag:
  - Video Display Controller ("sprite"): MSX, ColecoVision, TI-99/4...
  - Video palettes, SDRAM, VRAM.
  - TMS340 Graphics Signal Processor family.
  - TI Fellow 11 years after graduating and having joined TI.
  - Left TI in 1997.

Some failed attempts – MSA @ Intel

- Intel/ADI DSP/MCU architecture co-development since 2000.
  - ADI started marketing the Blackfin in 2001.
  - OthelloOne, SoftFone, TTPCom GPRS stack (ADI, 2002) along with Intel's XScale or NeoMagic's application processors.
  - Intel: Micro Signal Architecture. PXA800F / Manitoba in 2003.

– 0.13 µm
– 312 MHz XScale
– 104 MHz MSA
– 4 MB flash
– 512 KB SRAM

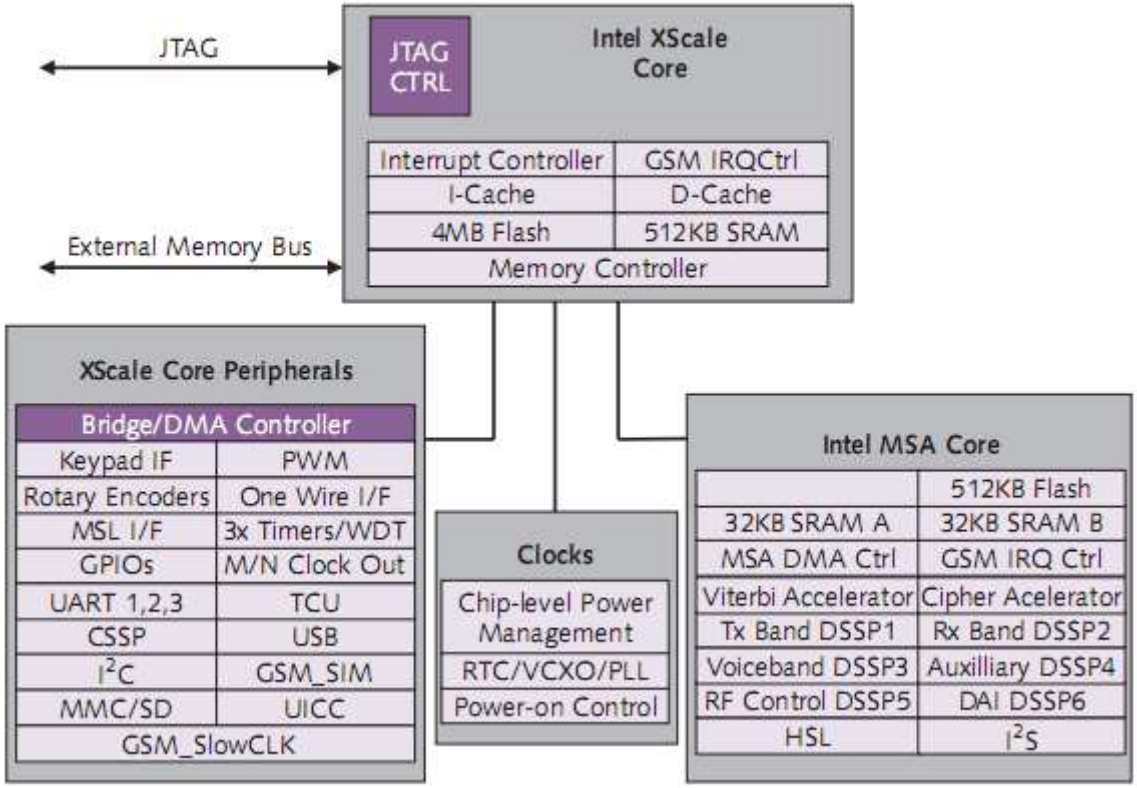

DSP market

- DSP-enabled silicon market is 10x the discrete DSP's.
  - 2008: $27.2B *vs.* $3.0B[†] (total semiconductor: $260B [‡])
- Definition of what a DSP is varies widely.
  - Forward Concepts: $3.0B *vs.* iSuppli $5.8B [‡] (2008).
- Shrinking market for discrete DSPs [‡] :
  - DSPs: -8.2%/year over 2008-2014
  - Total processors: +4.1%
  - Total digital ICs: +4.5%

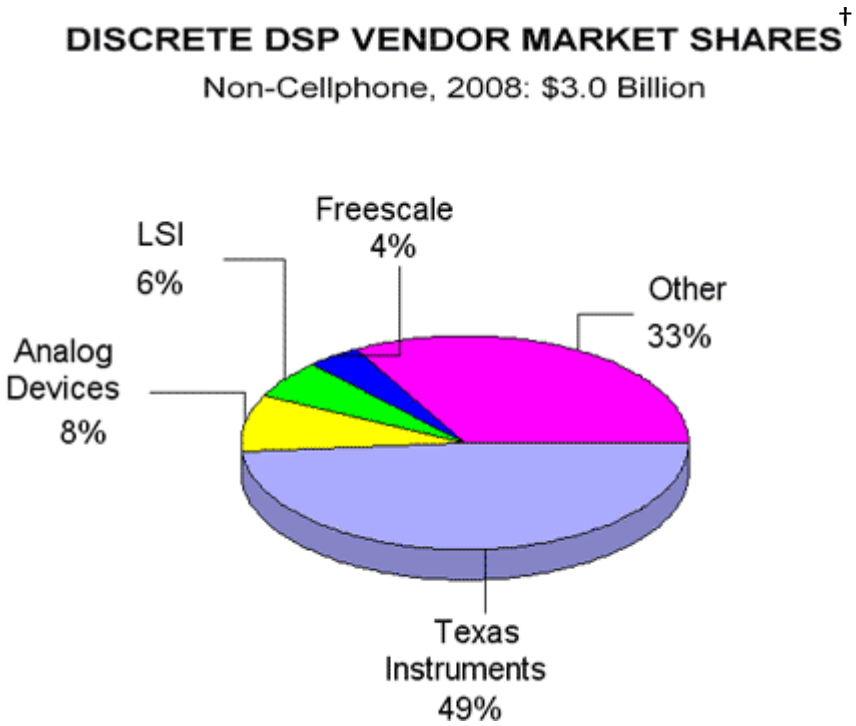

[†]Source: Forward Concepts, May 2009
[‡]Source: iSuppli, May 2010

Conclusion

- DSP technology has enabled many applications.
- Size of this market has attracted many competitors with a broad set of optimized solutions.
- Low-performance, low-cost DSP against MCUs.
- Low-power, high-volume DSP against ASSP with configurable processors or licensable DSP cores.
  - CEVA: Infineon, Broadcom, MediaTek, Spreadtrum, ST-Ericsson...
  - Tensilica: NTT DoCoMo/NEC/Fujitsu/Panasonic.
- High-performance DSP against MCU+FPGA and ASSP with configurable processors or licensable DSP cores or DSP arrays.
  - PicoChip, Tilera...

To probe further...

- The Scientist and Engineer's Guide to Digital Signal Processing
  - www.dspguide.com
- www.data-compression.com
- Forward Concepts, www.fwdconcepts.com (since 1984)
- Berkeley Design Technology, www.bdti.com (since 1991)
  - Hosts the comp.dsp FAQ www.bdti.com/Resources/Comp.DSP.FAQ
- EDN DSP Directory, www.edn.com (yearly)
- ieeexplore.ieee.org
- www.datasheetarchive.com
- This presentation is available on SlideShare.