Multiprocessor Architectures for Signal Processing Applications

Peter Pirsch

Institute of Microelectronic Circuits and Systems

Universität Hannover

Outline

- Mapping Signal Processing Applications to Multiprocessor Architectures
  - Application Example
  - From Task-specific Processing To Multiprocessing Platforms
    - FFT algorithms: HiPAR-DSP
    - Video compression: Macroblock Processor
    - Stream processing: RISC Core
- Reconfigurable Computing

Introduction

- Problems related to image sequence processing
  - Real-time processing
    - Continuous sequence of data
  - Signal processing requirements
    - Computational rate, data access, storage requirements
  - Implementation requirements
    - Size and power consumption
  - General-purpose processors do not offer needed performance
- Solution
  - Specialized signal processing architectures
  - Multiprocessor on chip

High Performance Signal Processing

- Concurrent processing of several operations necessary
- Parallelism on different levels
  - Independent processing of data blocks
    - ⇨Data-level concurrency
  - Parallel execution of independent instructions
    - ⇨Instruction-level concurrency
  - Independent parts of processing scheme (tasks)
    - ⇨Task-level concurrency

High Performance Signal Processing Example

- Onboard Radar Processor (Synthetic Aperture Radar)
- Compression of data rate
- Modulation for transmission

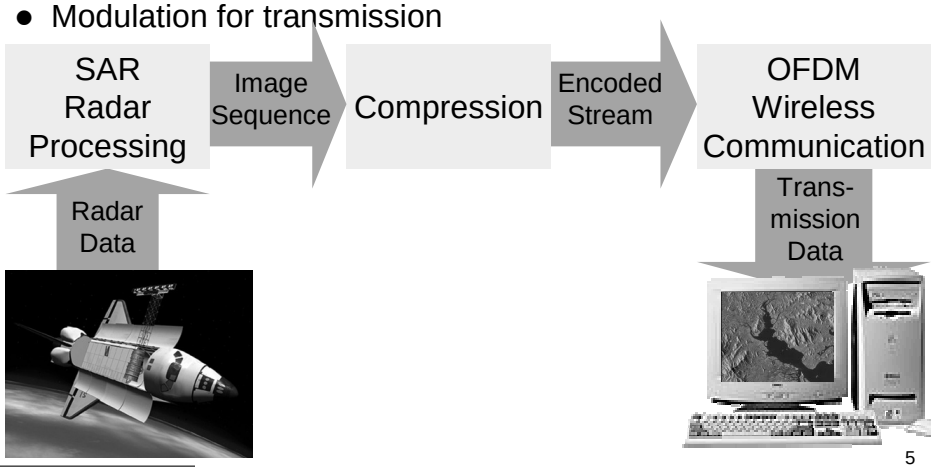

Multiprocessor Approach

- Application driven derivation of task-specific processors
  - Isolation of similar data and processing types in application
  - Mapping onto independent, programmable processors
  - Optimize each processor in terms of
    - Implemented type of parallel processing
    - Instruction set
    - Load-/Store capabilities

  ➔Utilization of parallelization resources on data level and instruction level
- Multiprocessor integration with task-specific processors on a single chip

  ➔Concurrent processing and data access on task level

Derivation of Task-specific Processors

- SAR application
  - Algorithm for image generation
- Modulation for transmission
  - Multicarrier technique
- Video compression
  - MPEG algorithm

Synthetic Aperture Radar

- Emitted impulse is reflected by objects in footprint
- Echo is recorded or processed on-board

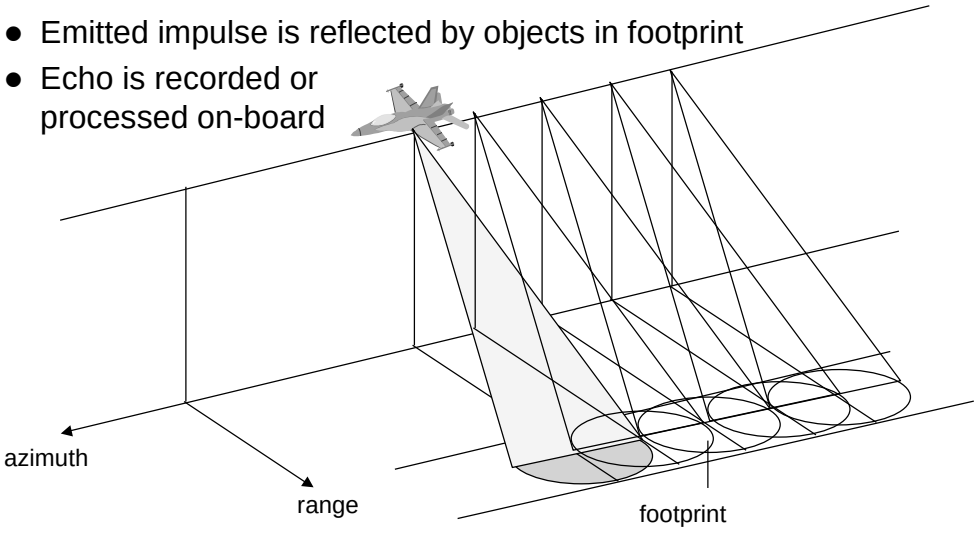

SAR Processing

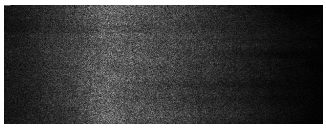

SAR Rawdata

- FFT
- Rangecompression
- IFFT
- 2D-FFT
- Azimuthprocessing
- 2D-IFFT

- SAR processing in frequency domain
- Rangecompression: motion compensation
- Azimuthprocessing: geometric correction
- 6 Fast Fourier Transformations necessary

Image data

OFDM Modulation

- Orthogonal Frequency Multiplexing
  - Multicarrier transmission technique for wireless data communication
  - Implemented in standards: ADSL, ETSI, DAB, DVB-T, W-LAN
  - Orthogonal modulated signals on multiple modulated carriers

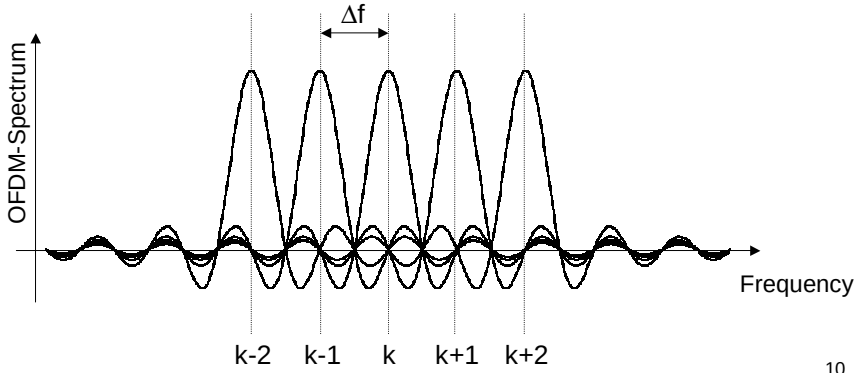

OFDM-System

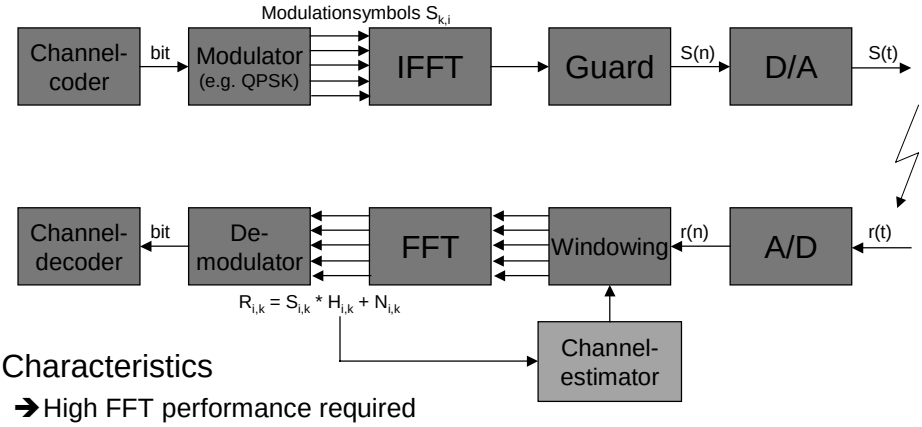

- Characteristics
  - ➔ High FFT performance required
  - ➔ Flexible adaptation to channel dependent guard intervalls
- ➔ FFT is essential operation in SAR- and OFDM-Applications

Butterfly FFT

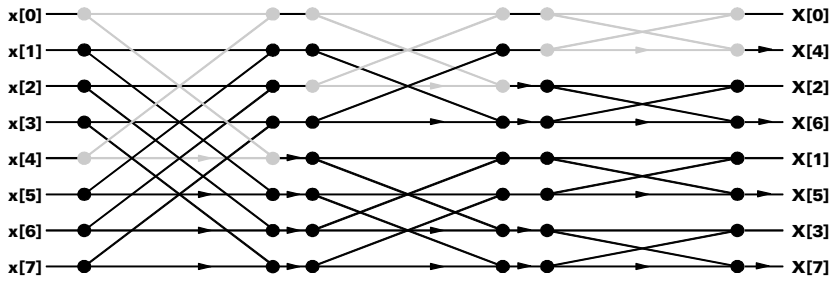

Processing on PU#1

- Parallel butterfly calculation
  - Data Level Concurrency offers high potential for performance gain
    - ⇨ SIMD Controlling to maximize concurrency on limited silicon area
  - Parallelization requires data exchange between processing units (PU)
    - ⇨ **Matrix Memory:** Shared memory with parallel data access

Matrix Memory Access

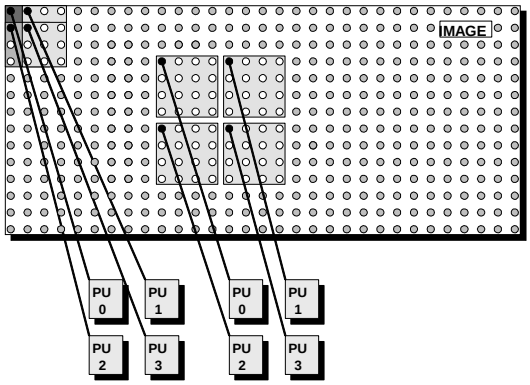

- Concurrent data access of processing units (PU) to Matrix Memory
  - Typical access patterns for filter / transform algorithms:
    Matrices with/without distances between memory elements
  - ➨ Conflicts caused by parallel access can be solved / avoided

Matrix Memory Architecture

- Appropriate distribution of 2x2 data to 3x3 RAM modules
  - Conflict-free access on matrices with distances 0 and 2^x
  - Only 15-20% area overhead for address calculation and shuffle net

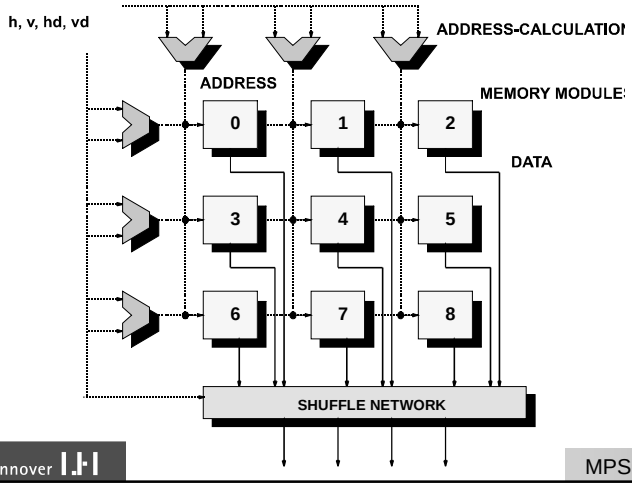

HiPAR-DSP Architecture Approach

- Enhanced data parallelism (SIMD)
  - Single unit controls scalable number of 4 / 16 data paths
  - Additional subword parallelism
- Moderate instruction parallelism
  - VLIW with 3 operations
- Matrix Memory for data exchange
  - Concurrent high bandwidth data access
  - Programmable in high-level language

Architecture: HiPAR-DSP

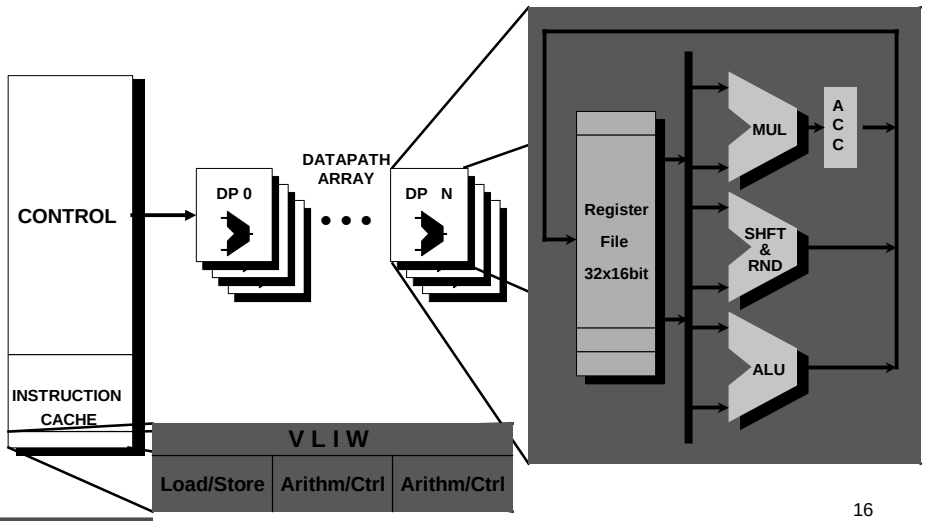

Architecture: HiPAR-DSP

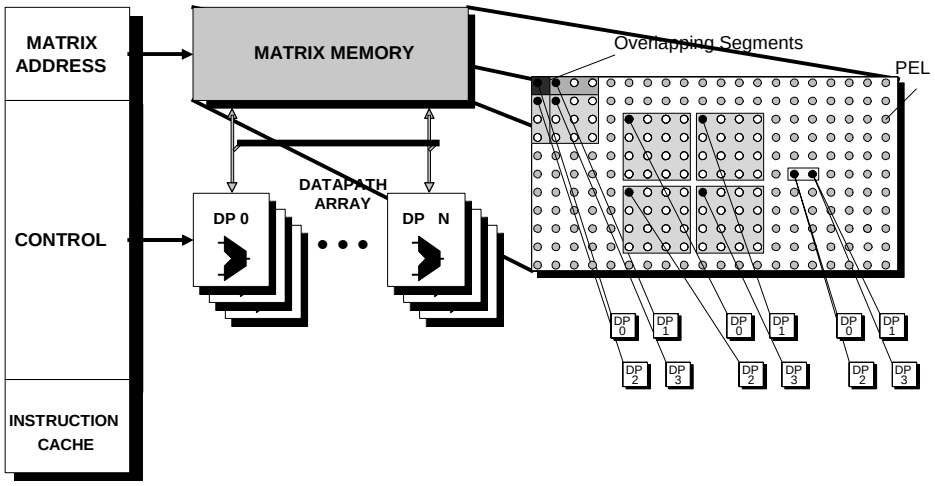

Architecture: HiPAR-DSP

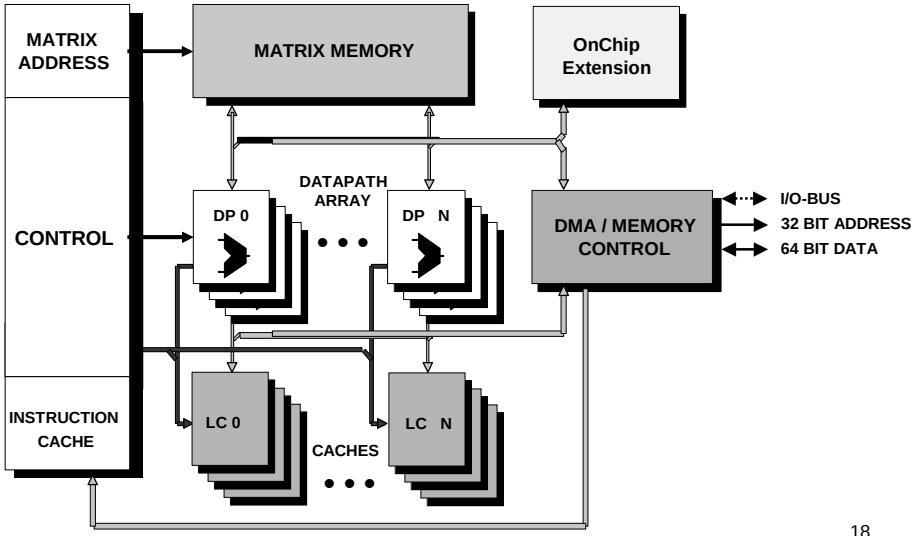

Video Compression Based on MPEG

- Characteristics
  - Motion compensated prediction
  - DCT on prediction error
  - Variable length code words
- High performance requirements
  - Motion estimation
    - Block Matching
  - Hardware extensions for higher throughput
    - Specialized instructions in VLIW processor

Motion Estimation : Block Matching (BM)

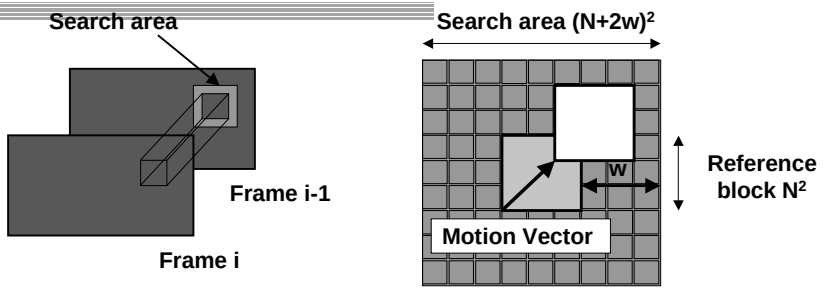

- Comparison of NxN reference blocks in MxM search window
- Distance measure

$$D(m,n) \equiv \sum_{i\equiv 0}^{N-1} \sum_{j\equiv 0}^{N-1} \left| x(i,j) - y(i+m, j+n) \right|$$

- Search for best candidate block in search window

BM: Software-Implementation with Sum of Absolute Differences Instruction

- Vector instruction: Sum-of-absolute differences V_SAD
- $S = \sum(|OpA-OpB|)$
- Example: 32-Bit

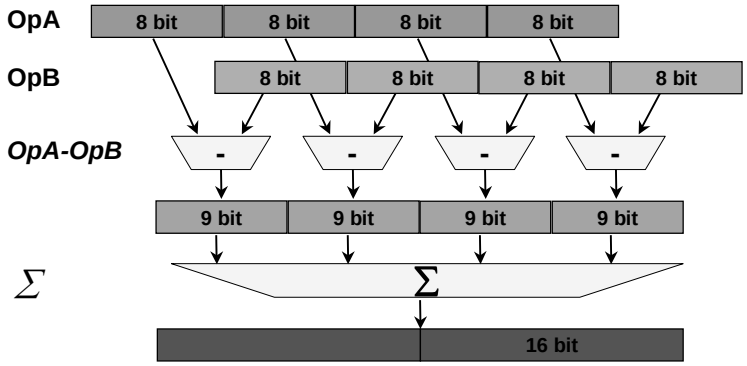

Motion Estimation for MPEG 4

- Increase of compression rate
  - Sub-pel accuracy
  - MPEG 4: up to 1/8 pel resolution
  - Interpolation filter necessary

BM: Sub-Pel Interpolation

- Interpolation filter
  - 4 Cases: FF/FH/HF/HH

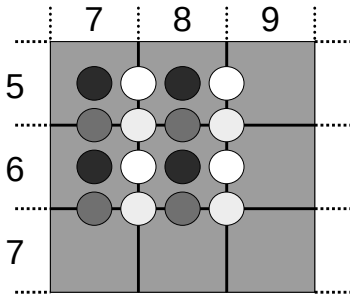

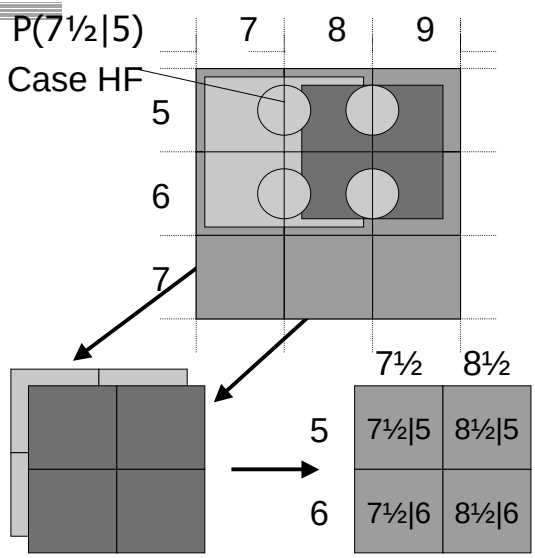

Software Interpolation with Average Instruction

- Vector average instruction: V_AVGRC
- *p = (OpA+OpB+rndctrl) shr 1*
- 4 Parallel interpolations / clock-cycle (32-Bit)
- Example: 32-Bit

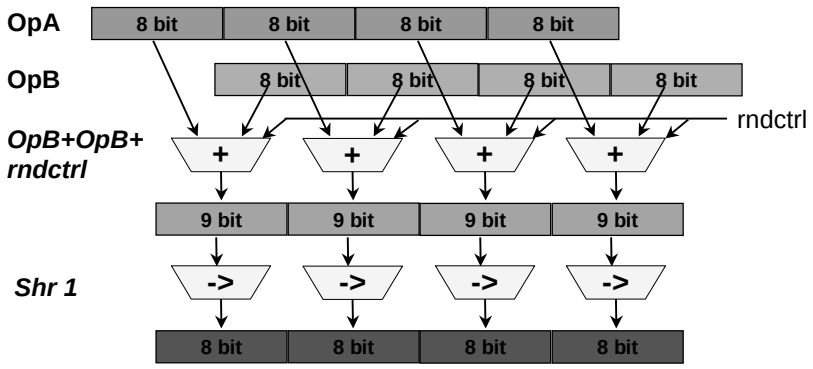

Macroblock Processor (MP)

- VLIW core: 2-issue, 64-bit (splittable)
- Vector unit (64 bit) Scalar unit (32 bit)
- Local memories: Instruction (8kb) Data (2kb)
- Peak performance: 1.7 GOPS @ 108 MHz
  - Data, instruction level parallelism
  - Specialized instructions
  - Parallel data transfer

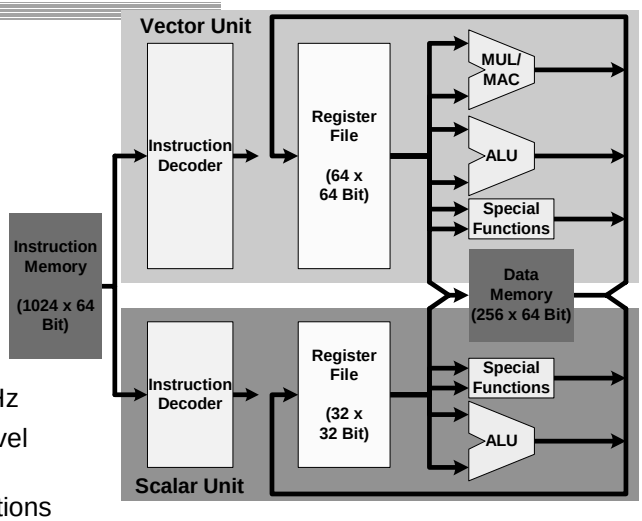

Specialized Instructions for Macroblock Processing

- Quad 16 → 32-Bit MAC with saturation (128-Bit result)
- Shift with round to 0 / ∞ unsigned, signed
  - ➔Transform, filter (QMC, deblocking)
- Average value with rounding control
  - ➔Sub-pel motion compensation
- Addition of absolute value (ABSADD)
- Controlled Addition/Subtraction (ADDSUB)
  - ➔Dequantization
- Permute instruction (perfect data shuffle with 2 regs)
  - ➔Motion compensation, deblocking
- Branch on vector status registers
  - ➔Deblocking mode selection

Bitstream Processing

- Application
  - Audio/Video stream generation and separation
- Characteristics in MPEG encoding
  - Multiplexing of different parts of bitstream
  - Run-Length coding of DCT coefficients
  - Variable length coding of coded DCT coefficients (using Huffman table)

Variable Length Coding

- Short codewords for frequent occurring symbols
- Long codewords for rare occurring symbols

Codewords	Symbols
0	S0
10	S1
110	S2
111	S3

Variable Length Decoding

- Parsing bitstream
  - Variable number of bits form a codeword
  - Extraction of codewords
  - Replacing codeword with symbol

➔ Large amount of bit oriented operations

- Problem:
  - Standard operations use byte/word widths (8/16 bits)
- Solution:
  - Implementation of specialized bit-operations

Instruction-Set Extensions

⇨ Specialized instructions (1 cycle) for frequently used functions

- ShowBits

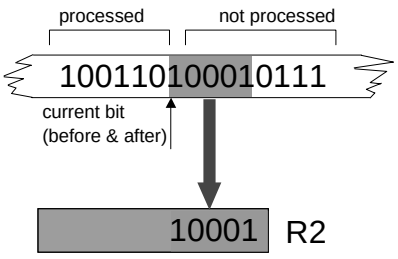

- GetBits

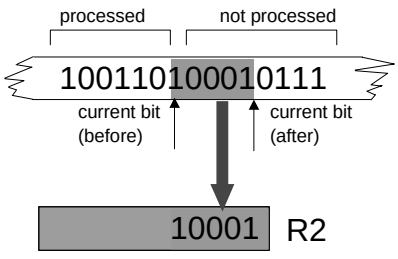

RISC Core Architecture Approach

- Scalar operation flow
  - ➔Usage of standard RISC with concurrent access on instructions and data
- Enhancements to instruction set
  - ➔Inclusion of bit oriented operations

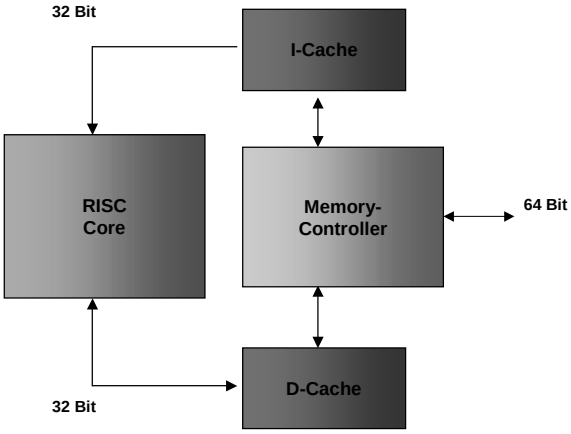

Multiprocessing

- Multiprocessing application featuring several tasks
  - Preprocessing (e.g., filtering)
  - Image/Video compression
  - Channel coding (e.g., OFDM modulation)
- Implementation goals
  - Meeting real-time requirements for entire application
  - Compact realization / low power constraints

HiBRID-SoC Multiprocessor Approach

- Integration of 3 processor cores on a single chip
  - HiPAR-16 Core
    - SAR processing
    - Filter and transform operations (e.g., FFT)
    - Modulation (e.g., OFDM)
  - Macroblock Processor
    - VLIW architecture adapted to video coding applications (MPEG-4 Advanced Simple Profile featuring Global Motion Compensation)
  - RISC-Core
    - Controlling tasks
    - VLC, VLD
    - Bitstream processing

HiBRID-SoC Multiprocessor

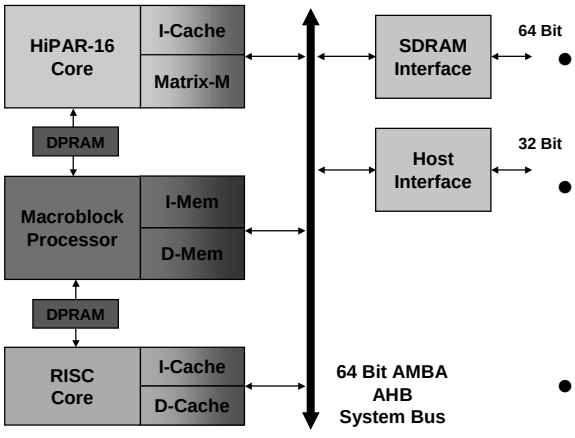

Communication Platform

- Caches and DMA control allow autonomous data transfers concurrently to processing
- External communication via SDRAM interface
  - 64-Bit AMBA AHB system bus
    - ➔Flexible use of DSPs as building blocks
- Internal communication of cores via Dual-Port-RAM
    - ➔Low controlling overhead

Outline

- Mapping Signal Processing Applications to Multiprocessor Architectures
  - Application Example
  - From Task-specific Processing To Multiprocessing Platforms
    - FFT algorithms: HiPAR-DSP
    - Video compression: Macroblock Processor
    - Stream processing: RISC Core
- Reconfigurable Computing

Reconfigurable Computing (RC)

- Goal
  - Adaptation of hardware configuration to algorithms at run time
  - Direct mapping onto suitable processing elements allows optimum computing density (power, area, performance)
  - Closing the gap between HW and SW design
  - Capability of in-field upgrades
  - Time-to-market shortening ⇒ reduces product life-cycle cost
- Reconfigurable computing system
  - Reconfigurable basic cells, reconfigurable network
  - Combination of reconfigurable resources, processor cores and dedicated arithmetic HW as R-SoC (reconfigurable SoC) to higher performance on data-intensive tasks

Efficiency of Reconfigurable Computing

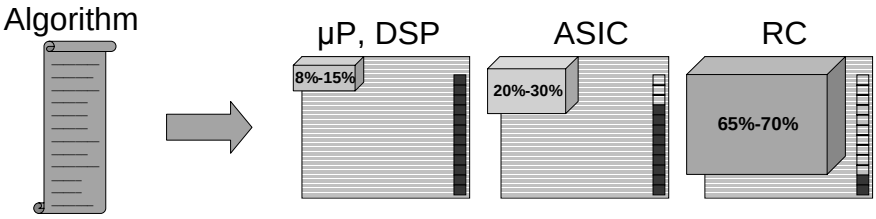

- Future of general purpose and application-specific processors
  - Increasing clock speed, gate number and performance
  - Decreasing efficient use of transistors, e.g. for parallel computing applications
- High computational and power efficiency
  - Much work accomplished within a given timeframe against the amount of power consumed
  - Measurement not in terms of MHz, MIPS or number of gates but in efficient use of transistors

Current Virtex-II FPGA Platform

- µP
  - PowerPC
  - D/I – Caches
  - Controllers
  - Interfaces
- DSP
  - Distributed RAM (Filter coefficients)
  - 18x18 Multipliers
  - 600 billion MACs/second
- Connectivity
  - 3.125 Gbps serial
  - 100+ Gb Bandwidth

DSP vs. Xilinx FPGA benchmark

Function	Industry's fastest DSP Processor Core	Xilinx Virtex-E -8	Xilinx Virtex-II
8x8 Multiply-and-Accumulate (MAC)	8.8 billion MACs/s	128 billion MACs/s	600 billion MACs/s
FIR filter 256-tap linear phase 16-bit data/coeff.	17 Mega samples per second (MSPS) 1.1 GHz	160 MSPS 160 MHz	180 MSPS 180 MHz
FFT 1024 point 16-bit data	7.7 µsec 800 MHz	41 µsec 100 MHz	1 µsec 140 MHz

Different granularities for future RC

- Coarse-grained components
  - General Purpose Processor Core
- Medium-grained components
  - DSP functionality (dedicated hardware: ALU, MAC, MULT...)
  - Synchronous Block-RAM
  - Intracomponent connection network
- Fine-grained components
  - Reconfigurable processing elements
  - Reconfigurable interconnection network

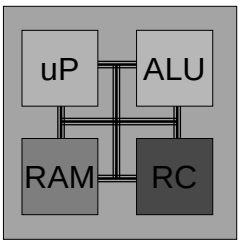

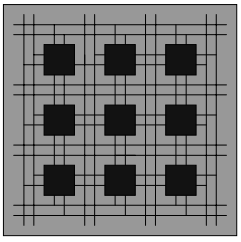

Conclusion

- Implementation efficiency:
  Adaptation of processor architectures to processing schemes
- High performance requirements:
  Exploitation of all existing parallelization levels
- System implementations:
  Several task specific processors on single chip together with communication network
- Outlook:
  Programmable processors vs. reconfigurable computing still open question

Acknowledgements

- Mladen Berekovic
- Willm Hinrichs
- Javier Martín Langerwerf
- Carsten Reuter
- Christian Simon-Klar

References

Processor Architectures

[Pirs95] P. Pirsch, J. Kneip, K. Rönner, "Parallelization Resources of Image Processing Algorithms and their Mapping on a Programmable Parallel Videosignal Processor," *Proc. of the International Symposium on Circuits and Systems 1995*, pp. I-562–565, Seattle, 1995

[Pirs99] P. Pirsch, "Architectures for Multimedia Signal Processing," *Proc. IEEE Workshop on Signal Processing Systems*, pp. 1–12, Nov. 1999.

[Pirs02] P. Pirsch, A. Freimann, C. Klar, J. P. Wittenburg, "Processor Architectures for Multimedia Applications," *Embedded Processor Design Challenges: Systems, Architectures, Modeling, and Simulation - SAMOS*, Springer Verlag, Heidelberg, pp. 188–206, Feb. 2002

Synthetic Aperture Radar SAR

[Curl91] John C. Curlander, Robert N. McDonough, *Synthetic Aperture Radar*, John Wiley and Sons, Inc., 1991

[Soum99] Mehrdad Soumekh, *Synthetic Aperture Radar Signal Processing with MATLAB algorithms*, John Wiley and Sons, Inc., 1999

[Fitc88] J. Patrick Fitch, *Synthetic Aperture Radar*, Springer, 1988

[Kloo01] H. Kloos, L. Friebe, J. P. Wittenburg, W. Hinrichs, H. Lieske and P. Pirsch, "HiPAR-DSP 16, A new DSP for Onboard Real-Time SAR Systems," *Proc. of 15th Aerosense Conference on Phototonic and Quantum Technologies for Aerospace and Application III*, July 2001.

Orthogonal Frequency Division Multiplexing OFDM

[Baha99] Ahmad R. S. Bahai, Burton R. Saltzberg, *Multi-carrier digital communications : theory and applications of OFDM*, Kluwer Academic/Plenum, 1999

[Nee00] Richard van Nee, Ramjee Prasad, *OFDM wireless multimedia communications*, Artech House, 2000

[Rohl01] H. Rohling, R. Grünheid and D. Galda, "OFDM Air Interface for the 4th Generation of Mobile Communication Systems," *Proc. of the 6th International OFDM-Workshop*, pp. 1–28, Hamburg, Germany, Sept. 2001.

MPEG

[MPEG-2] ISO/IEC 13818–2, "Generic coding of moving pictures and associated audio, (MPEG-2), Part 2: Video," Nov. 1993.

[H.263] ITU-T Recommendation Draft H.263, "Video Coding for Low Bitrate Communications," International Telecommunication Union, May 1996.

[Chia98] L. Chiariglione, "Impact of MPEG Standards on Multimedia Industry," *Multimedia Signal Processing, Proc. of the IEEE*, Vol. 86, No. 6, pp. 1222–1227, June 1998.

[ISO98] ISO/IEC JTC11/SC29/WG11 N2323, "Overview of the MPEG-4 Standard," July 1998.

[MPEG-4] ISO/IEC JTC11/SC29/WG11 W2502, "ISO/IEC 14496-2. Final Draft international standard. Part 2: Visual," Atlantic-City, Oct. 1998

[Hask98] B. Haskell, P. G. Howard, Y. A. Lecun, A. Puri, J. Ostermann, M. R. Civanlar, L. Rabiner, L. Bottou and P. Haffner, "Image and Video Coding - Emerging Standards and beyond," *IEEE Trans. on Circuits and Systems for Video Technology*, Vol. 8 No. 7, pp. 878–891, Nov. 1998.

[Knei98] J. Kneip, S. Bauer, J. Volmer, B. Schmale, P. Kuhn, M. ReiSSmann, "The MPEG-4 Video Coding Standard - a VLSI point of view," *IEEE International Workshop on Signal Processing Systems SIPS98*, Boston, Oct. 1998.

[Bere99] M. Berekovic, H.-J. Stolberg, M. B. Kulaczewski and P. Pirsch, "Instruction Set Extension for MPEG-4 Video," *Journal of VLSI Signal Processing 23*, pp. 27–49, 1999.

[Bere02] M. Berekovic, H.-J. Stolberg, S. Flügel, S. Moch, M. B. Kulaczewski, L. Friebe, J. Hilgenstock, X. Mao, H. Klussmann, P. Pirsch, "Implementing The MPEG-4 AS Profile on a Multi-Core System on Chip Architecture," *Proccedings of 3rd Workshop and Exhibition on MPEG-4 (WEMP4)*, June 2002

Reconfigurable Computing

[Mirs96] E. Mirsky, A. DeHon, "MATRIX: A Reconfigurable Computing Architecture with Configurable Instruction Distribution and Deployable Resources," *Proceedings of the IEEE Symposium on FPGAs for Custom Computing Machines*, pp. 157–166, Apr. 1996.

[Radu98] B. Radunovic and V. Milutinovic, "A Survey of Reconfigurable Computing Architectures," *8th International Workshop on Field-Programmable Logic and Applications*, pp. 376–385, Tallinn, Estonia, Aug./Sept. 1998.

[Gucc00] S. A. Guccione, "List of FPGA-based Computing Machines," http://www.io.com/∼guccione/HW_list.html, last updated August 21, 2000.

[Tess01] Russell Tessier and Wayne Burleson, "Reconfigurable Computing for Digital Signal Processing: A Survey," *Journal of VLSI Signal Processing*, vol. 28, pp. 7–28, 2001.

Further Readings

[Cool65] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Computation of Complex Fourier Series," *Mathematics of Computation*, Vol. 9, pp. 297–301, Apr. 1965.

[Toma67] R. M. Tomasulo, "An Efficient Algorithm for Exploiting Multiple Arithmetic Units," *IBM Journal of Research and Development*, volume 11 (1), pp. 25-33, IBM, 1967

[Yama94] W. Yamamoto, M. J. Serrano, A. R. Talcott, R. C. Wood, M. Nemirovski, "Performance Estimation of Multistreamed, Superscalar Processors," *Proceedings of the 27th Annual Hawaii International Conference on System Sciences*, Honolulu, 1994.