Cryptanalysis on a PlayStation 3 cluster

by Joppe Bos

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

- Why is the PlayStation 3 (PS3) hardware of any interest?
- How should we implement our algorithms on the PS3?
- Existing and new video game clusters.
- Projects and results obtained on the PS3s at LACAL.

The PlayStation 3

Facts about the PS3:

- The third video game console by Sony Computer Entertainment
- Released in

Japan	11 November 2006
North America	17 November 2006
Europe	23 March 2007

- As of 30 June 2008 worldwide 14.41 million units sold

Hardware

- PS3 disc drive is an all-in-one type: 2× Blu-ray, 8× DVD and 24× CD
- Hard disk size $\in \{20, 40, 60, 80\}$ GB. This month the 160 GB version will be released
- $\{2, 4\}$ USB 2.0 ports (depending on version)
- A graphics processing unit manufactured by Nvidia
  - Based on the NVIDIA G70 architecture.
  - Makes use of 256 MB GDDR3 RAM clocked at 700 MHz
  - Unavailable to the programmer
- 3.2 GHz Cell Broadband Engine (Cell) microprocessor architecture jointly developed by Sony, Toshiba, and IBM

Cell architecture, overview

The Cell consists of the following components

- external input and output structures
- one “Power Processor Element” (PPE)
- eight Synergistic Processing Elements (SPEs)
  six SPEs available to the user
- the Element Interconnect Bus (EIB)
  a specialized high-bandwidth circular data bus

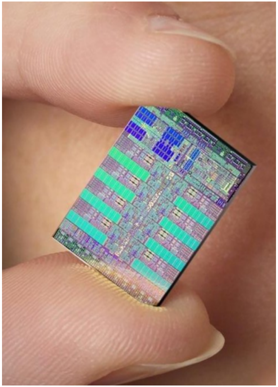

PS3 architecture, the PPE

- 64-bit PowerPC architecture core, can run in 32- and 64-bit mode
- 128-bit AltiVec/VMX SIMD unit
- dual-threaded processor
- 32 KB instruction- and a 32 KB data Level 1 cache
- 512 KB Level 2 cache
- $\sim$ 214 out of 256 MB of memory available to the guest OS
- instruct the workhorses (SPEs) what to do

PS3 architecture, the SPEs

- Synergistic Processing Unit (SPU)
    - Access to an 128 × 128-bit wide register file
    - SIMD architecture

PS3 architecture, the SPEs

- Synergistic Processing Unit (SPU)
  - Access to an 128 × 128-bit wide register file
  - SIMD architecture
- 256 KB of fast local memory (Local Store)

PS3 architecture, the SPEs

- Synergistic Processing Unit (SPU)
  - Access to an 128 × 128-bit wide register file
  - SIMD architecture
- 256 KB of fast local memory (Local Store)
- Memory Flow Controller (MFC)
  - Direct Memory Access (DMA) controller
  - Handles synchronization operations to the other SPUs and the PPU
  - DMA transfers are independent of the SPU program execution

Element Interconnect Bus

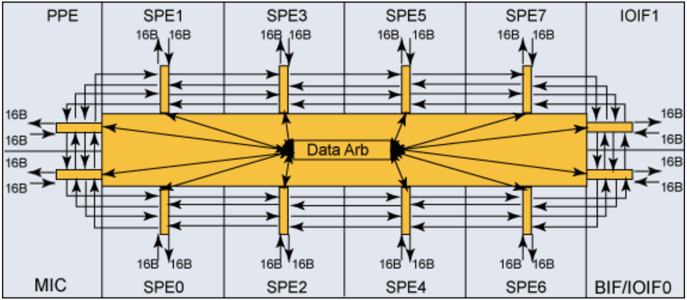

- 12 participants
- circular ring comprised of four 16 Byte-wide unidirectional channels
- peak instantaneous EIB bandwidth:
  $(4 \times 3) \times 16 \; / \; 2 = 96$ Byte per processor cycle (307.2 GB/s)

Limitations

- Branching
  - No “smart” dynamic branch prediction
  - Instead “prepare-to-branch” instructions to redirect instruction prefetch to branch targets

Limitations

- Branching
  - No “smart” dynamic branch prediction
  - Instead “prepare-to-branch” instructions to redirect instruction prefetch to branch targets
- Memory
  - The binary **and** all the needed memory should fit in the LS
  - Or perform manual DMA requests to the main memory (max. 214 MB)

Limitations

- Branching
  - No “smart” dynamic branch prediction
  - Instead “prepare-to-branch” instructions to redirect instruction prefetch to branch targets
- Memory
  - The binary **and** all the needed memory should fit in the LS
  - Or perform manual DMA requests to the main memory (max. 214 MB)
- Instruction set limitations
  - 16 bit multiplier

SPU registers

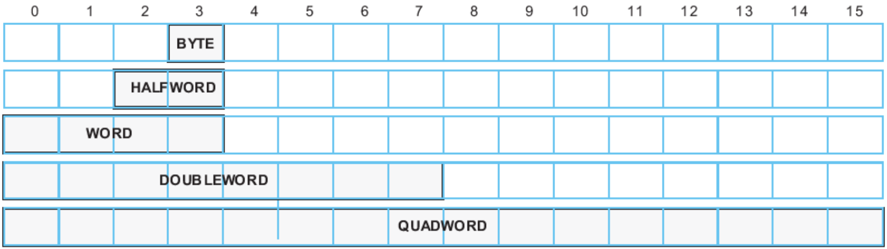

- Byte: 16 × 8-bit SIMD
- Half-word: 8 × 16-bit SIMD
- Word: 4 × 32-bit SIMD

SPU registers

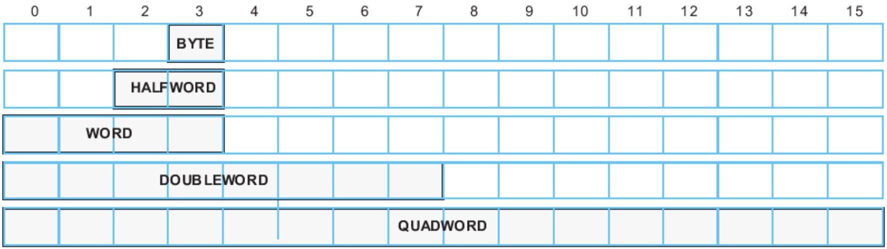

- Byte: 16 × 8-bit SIMD
- Half-word: 8 × 16-bit SIMD
- Word: 4 × 32-bit SIMD

Theoretical performance of $16 \times 3.2 \cdot 10^9 = 51.2$ billion 8-bit integer operations per second.

Special SPU instructions

All distinct binary operations $f : \{0,1\}^2 \rightarrow \{0,1\}$ are present.

shuffle bytes
or across
average of two vectors
select bits
carry/borrow generate
multiply and add
element-wise absolute difference
add/sub extended
count leading zeros
count ones in bytes
gather lsb
sum bytes
multiply and subtract

Special SPU instructions

All distinct binary operations $f : \{0,1\}^2 \to \{0,1\}$ are present.

- shuffle bytes
- or across
- average of two vectors
- select bits
- carry/borrow generate
- multiply and add
- element-wise absolute difference
- add/sub extended
- count leading zeros
- count ones in bytes
- gather lsb
- sum bytes
- multiply and subtract

shufb Concatenate two input registers to form a 32-byte lookup table Each byte in the third register selects either a constant value (0x00/0x80/0xFF) or a location in the lookup table → 16 table lookups per cycle

SPU pipelines and latencies

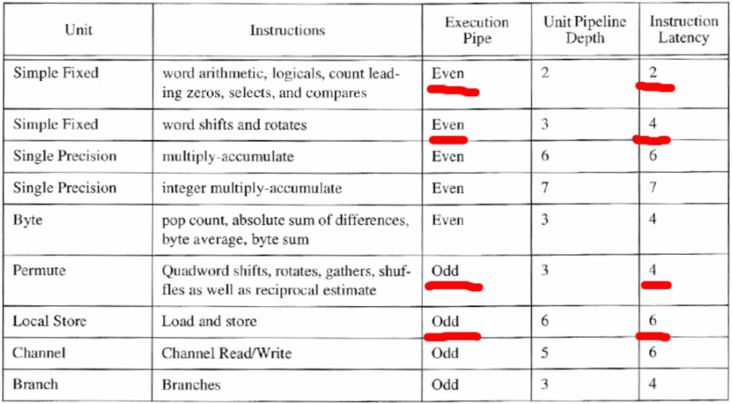

Unit	Instructions	Execution Pipe	Unit Pipeline Depth	Instruction Latency
Simple Fixed	word arithmetic, logicals, count leading zeros, selects, and compares	Even	2	2
Simple Fixed	word shifts and rotates	Even	3	4
Single Precision	multiply-accumulate	Even	6	6
Single Precision	integer multiply-accumulate	Even	7	7
Byte	pop count, absolute sum of differences, byte average, byte sum	Even	3	4
Permute	Quadword shifts, rotates, gathers, shuffles as well as reciprocal estimate	Odd	3	4
Local Store	Load and store	Odd	6	6
Channel	Channel Read/Write	Odd	5	6
Branch	Branches	Odd	3	4

One odd and one even instruction can be dispatched per clock cycle.
Challenge to the programmer (or compiler).

Cluster of game console

- Using the compute power of video game consoles is not new
- 65-node PS2 cluster build by the National Center for Supercomputing Applications and the University of Illinois in 2003

Cluster of game console

- Using the compute power of video game consoles is not new
- 65-node PS2 cluster build by the National Center for Supercomputing Applications and the University of Illinois in 2003
- Other uses, besides gaming and computing, include grilling:

Small clusters

Academic clusters

An 8 PS3-cluster at the North Carolina State University

An 16 PS3-cluster "Gravity Grid" at the University of Massachusetts

Small clusters

Academic clusters

An 8 PS3-cluster at the North Carolina State University

An 16 PS3-cluster “Gravity Grid” at the University of Massachusetts

Commercial clusters

Pre-installed PS3 from Terra Soft solutions:
8 Node PS3 Cluster $17,650 (≈ $2,200 per PS3)
32 Node PS3 Cluster $42,250 (≈ $1,300 per PS3)
(current PS3 price ≈ $400)

Warhawk mayhem

Ranked-Dedicated servers for the PS3 games called Warhawk mayhem

Warhawk mayhem

Ranked-Dedicated servers for the PS3 games called Warhawk mayhem

U.S. Air Force wants to buy **300** PS3s

LACAL setup

- Physically in the cluster room: 186 PS3s
- 6 × 4 PS3s in the PlayLaB (attached to the cluster)
- 9 PS3 scattered over our offices for programming purposes
- ⇒ 219 PS3s in total.

LACAL setup

- Physically in the cluster room: 186 PS3s
- 6 × 4 PS3s in the PlayLaB (attached to the cluster)
- 9 PS3 scattered over our offices for programming purposes
- ⇒ 219 PS3s in total.

How do we put these machines to work?

Finding MD5 multi-collisions

Performed by: Marc Stevens, Arjen Lenstra, Benne de Weger.

- Summer 2007:
  Single chosen-prefixes MD5 collision after half year on BOINC network (no PS3s used)
- Fall 2007:
  Previous attack in 3 hours on single PS3
  (with 30-fold MD5 speed-up on PS3 over desktop)
- Proof of concept example:
  12 PDF turned into a MD5 multi-collision: “Predicting the winner of the 2008 US Presidential Elections using a Sony PlayStation 3”

Multi-Stream Hashing on the PlayStation 3
Joppe Bos, Nathalie Casati and Dag Arne Osvik
PARA 2008: State-of-the-Art in Scientific and Parallel Computing

Idea: Using the SIMD-organization of the SPUs to hash multiple streams and hide latencies.

Algorithm	Gb / sec / PS3	Gb / sec / Core2Quad (*)
MD5	88.17	64
SHA-1	43.60	34.8
SHA-256	18.70	13.5

(*) Upper-bound by carefully counting instructions
Hashing 10^5 150 KB messages with the assembly version.

Finished student projects related to ECM at LACAL

- Sylvain Pelissier and Aniruddha Bhargava
  First attempt to port GMP to the SPU
    - code size versus performance

Finished student projects related to ECM at LACAL

- Sylvain Pelissier and Aniruddha Bhargava
  First attempt to port GMP to the SPU
  - code size versus performance
- Thomas Kunz: GMP-ECM on the PS3
  - Non-trivial, code size problems
  - Replace low-level building blocks

Finished student projects related to ECM at LACAL

- Sylvain Pelissier and Aniruddha Bhargava
  First attempt to port GMP to the SPU
  - code size versus performance
- Thomas Kunz: GMP-ECM on the PS3
  - Non-trivial, code size problems
  - Replace low-level building blocks
- Donato Verardi: MPM-ECM based on GMP-ECM
  - Fast! But many improvements are still possible

Finished student projects related to ECM at LACAL

- Sylvain Pelissier and Aniruddha Bhargava
  First attempt to port GMP to the SPU
  - code size versus performance
- Thomas Kunz: GMP-ECM on the PS3
  - Non-trivial, code size problems
  - Replace low-level building blocks
- Donato Verardi: MPM-ECM based on GMP-ECM
  - Fast! But many improvements are still possible
- Stage 1 only
- Limitations: input number must be < 2048 bits

Time in seconds to run 12 curves on different input length with different B1-values.

B1-value	Donato	Thomas	PENTIUM-D
512-bit input			
250000	26	30	**22**
1000000	108	**68**	89
3000000	322	341	**274**
768-bit input			
250000	37	**34**	44
1000000	150	**138**	179
3000000	448	**414**	543
1024-bit input			
250000	**47**	50	72
1000000	**189**	200	300
3000000	**567**	601	877

Pollard rho for finding ECDL

Work in progress:
Pollard rho on the PS3 by Joppe Bos and Marcelo Kaihara

Motivation

Branch-free SIMD Pollard rho to calculate the elliptic curve discrete logarithms (over prime fields)
Currently runs on SPU only;
An implementation which offloads work to the PPE is in progress

Pollard rho for finding ECDL

Work in progress:
Pollard rho on the PS3 by Joppe Bos and Marcelo Kaihara

Motivation

Branch-free SIMD Pollard rho to calculate the elliptic curve discrete logarithms (over prime fields)
Currently runs on SPU only;
An implementation which offloads work to the PPE is in progress

Current speed:
ECCP-109: $1.5 \cdot 10^7$ iterations per second per PS3
$\Rightarrow$ less than 4 months on a PS3 cluster with 200 nodes.
ECCP-131: 10^7 iterations per second per PS3
$\Rightarrow$ 800 years on a PS3 cluster with 200 nodes.

New projects

- PS3s attract {bachelor, master} students!
- This semester:
  - Implementation of ECM stage 2 on the SPE.
  - Creating a set of script to handle all the ECM jobs on the cluster.
  - “Monster RSA”; RSA encryption/decryption with 15k modulus
  - Efficient arithmetic using the residue number system (RNS)

Conclusions

- The PS3 hardware (i.e. Cell) is very interesting
  - Some limitations: memory, 16 bit multiplier
  - Think SIMD, avoid branching, exploit the dual-pipeline and use the rich instruction set
- The cluster attracts many students
  → lots of new PS3 project are on their way!

Conclusions

- The PS3 hardware (i.e. Cell) is very interesting
  - Some limitations: memory, 16 bit multiplier
  - Think SIMD, avoid branching, exploit the dual-pipeline and use the rich instruction set
- The cluster attracts many students
  → lots of new PS3 project are on their way!
- In the future: PS4 (rumors say 2012)? More main memory? More SPEs?