



And now for
something
completely
different..

CorePy: High-Productivity Cell/B.E. Programming


Chris Mueller, Ben Martin, and Andrew Lumsdaine
Indiana University

June 19, 2007

Georgia Tech, Sony/Toshiba/IBM Workshop on Software and
Applications for the Cell/B.E. processor



Cell/B.E.



For *all* your multi-core and SIMD programming needs...



*** Featuring ***

- ⇒ *Advanced “make” build system!*
- ⇒ *Cutting edge “gdb” debugger!*
- ⇒ *Unparalleled C standard library!*
- ⇒ *Works with **any** text editor!*

*Auto-parallelizing, auto-simdizing, optimizing compiler not yet available.
For maximum SIMD performance, use of assembly may be required.
Void where prohibited, prohibited where void.

For *all* your multi-core and SIMD programming needs...



*** Featuring ***

- ⇒ *Advanced “make” build system!*
- ⇒ *Cutting edge “gdb” debugger!*
- ⇒ *Unparalleled C standard library!*
- ⇒ *Works with **any** text editor!*

Is there an alternative?

*Auto-parallelizing, auto-simdizing, optimizing compiler not yet available.
For maximum SIMD performance, use of assembly may be required.
Void where prohibited, prohibited where void.

The Approach

Take a modern programming technique...



(Python + C + agile development)



The Approach

Take a modern programming technique...

...provide direct access to the hardware...



(somewhere between machine code and assembly)



pervasive technology labs
AT INDIANA UNIVERSITY

The Approach

Take a modern programming technique...

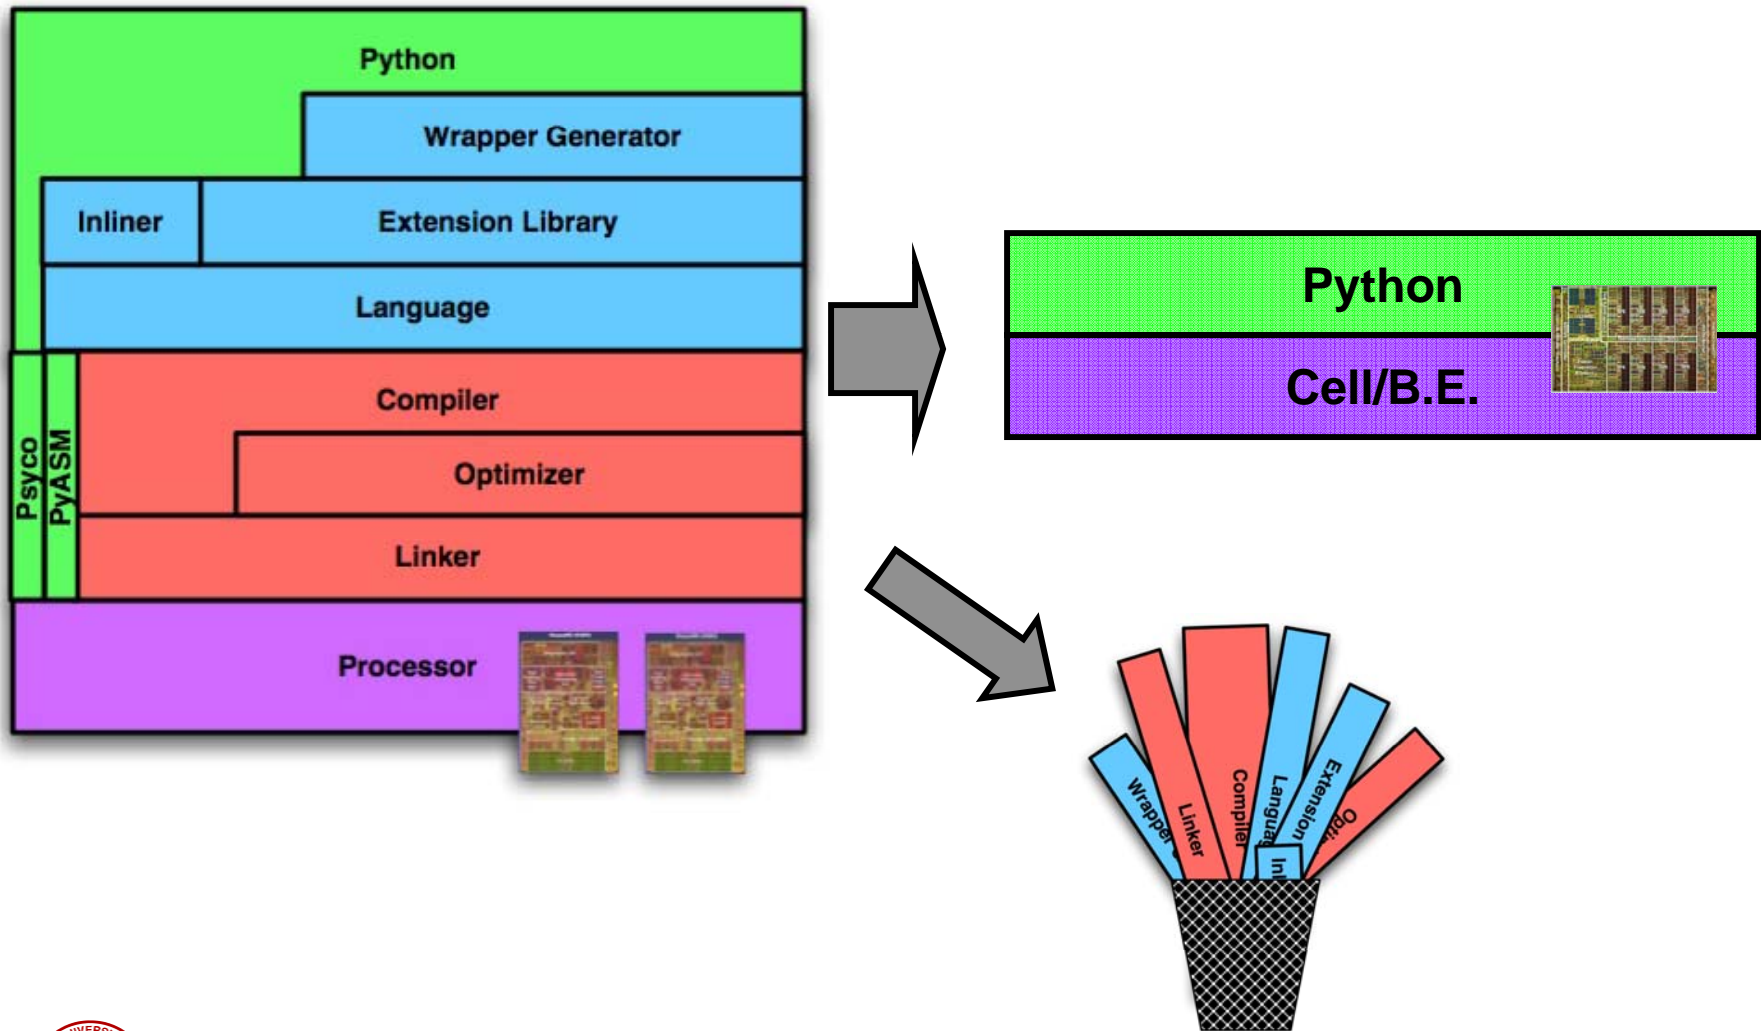
...provide direct access to the hardware...

... and let programmers explore the SIMD and multi-core design spaces.



(power to the people?)

The Approach, Illustrated




Huh?!? Python for Cell Programming?

- Scripting languages increase developer productivity
 - Dynamic type systems
 - Large standard library collections (~240 w/Python)
 - Concise syntax
 - “Edit and execute”
 - Limit reliance on external tools
- Scripting languages are already used in HPC
 - “Glue” code to tie applications together
 - “Kleenex” applications
 - Native libraries for performance



CorePy

CorePy is a library for creating and executing PowerPC, VMX, and SPU programs from Python.



- ❑ Execute arbitrary SPE programs from Python
- ❑ Talk to SPE programs using libspe wrappers
- ❑ Create *new* SPE or PPE programs directly from Python



CorePy Components

Execution Management

InstructionStream

Container for SPE/PPE programs

Processor

Execution management and `libspe` functions

Memory

Memory alignment and transfer utilities

Code Synthesis

Instruction Set
Architecture (ISA)

PowerPC, VMX, and SPU instruction interface

Variables

Library-defined “primitive” types

Iterators

Loop generation and optimization




Example: Population Count

Population count (popc) counts the number of '1' bits in a bit vector. It is used in many data clustering and machine-learning algorithms to compare feature vectors.

1. `v = {0b0101010111101101, 0b1100110000000001,`
2. `0b0000001000100001, 0b1111111111000001}`
3. `c = popc(v)`
4. `print c`
5. `-> 29`

Structure



Bit-vector Representation

00101101101100100001001111...

Alkaloid bit Rule-of-five bits

Metrics

$$D_{Rand} = \frac{c + d}{a + b + c + d}$$

$$D_{Jaccard} = \frac{c}{a + b + c}$$

x:101011101

y:110011010

cbadccaba

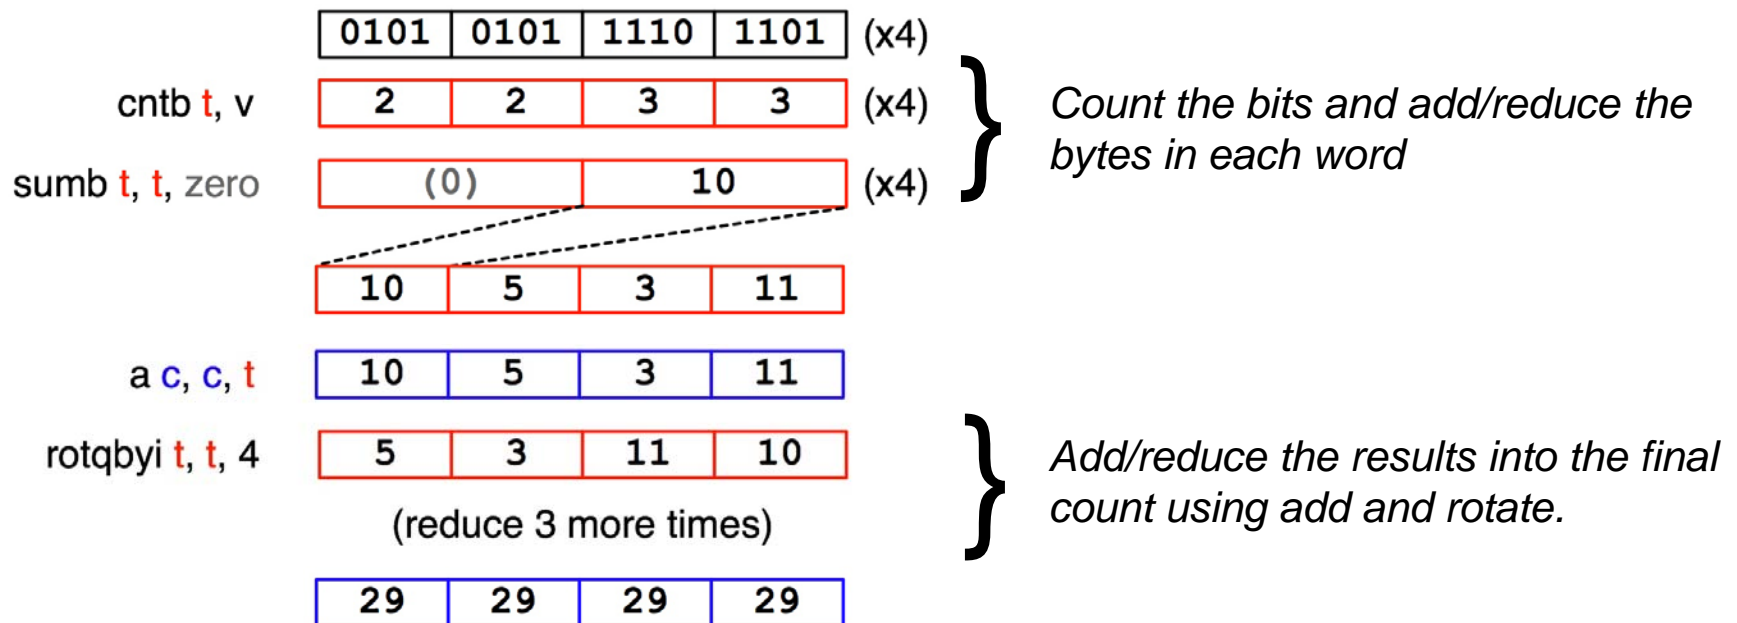
$D_{Rand} = .444, D_{Jaccard} = .375$



An SPE Population Count

128-bit bit vector population count using SPE instructions.

cntb: count bits in bytes
sumb: sum bytes into halfwords
a: add
rotqbyi: rotate quadword by bytes, immediate



spu_popc: Population Count in C

```
1.  #include "corepy.h"

2.  int main(unsigned long long id) {
3.      // Load the vector from input parameter 3
4.      vector unsigned int x = get_vector_param_3();
5.      vector unsigned int count = (vector unsigned int){0,0,0,0};
6.      vector unsigned int result = (vector unsigned int){0,0,0,0};
7.      int i = 0;

8.      // Initialize CorePy run time
9.      spu_ready();

10.     count = (vector unsigned int)spu_cntb((vector unsigned char)x);
11.     count = (vector unsigned int)spu_sumb((vector unsigned char)count,
12.                                           (vector unsigned char)0);

13.     for(i; i < 4; i++) {
14.         result = result + x;
15.         x = si_rotqbyi(x, 4);
16.     }
17.
18.     // Report the result using the out mailbox
19.     spu_write_out_mbox(spu_extract(result, 0));

20.     return SPU_SUCCESS;
21. }
```



Calling spu_popc Using CorePy

The Processor object executes SPE programs and retrieves return values.

```
1. import corepy.arch.spu.platform as env
2. # Load the popc SPE program
3. code = env.NativeInstructionStream("spu_popc")
4. proc = env.Processor()

5. # Set the input parameters
6. params = env.spu_exec.ExecParams()
7. params.v3 = [0xDEAD, 0xBEEF, 0x1337, 0xCAFE]

8. # Execute the popc program with the parameters
9. count = proc.execute(code, 'mbox', params)
```



Processor and InstructionStream

InstructionStream manages executable instruction sequences

```
# Use NativeInstructionStream for compiled SPE programs
c = env.NativeInstructionStream("spu_popc")
```

Processor manages execution and return values.

Return values:

1. `# result in gp_return`
2. `r = p.execute(c)`
3. `# result in fp_return`
4. `f = p.execute(c, 'fp')`
5. `# result in mbox (spe)`
6. `r = p.execute(c, 'mbox')`

Asynchronous execution:

1. `t1 = p.execute(c1, mode='async')`
2. `t2 = p.execute(c2, mode='async')`
3. `p.suspend(t2)`
4. `p.resume(t2)`
5. `p.stop(t1)`
6. `p.join(t2)`



CorePy and Synthetic Programming

Synthetic programming is a meta-programming technique for synthesizing instruction sequences at run-time from high-level languages.

A Simple Example:

$$r = ((0 + 31) + 11)$$

```
1. c = InstructionStream()
2. ppc.set_active_code(c)

3. ppc.addi(gp_return, 0, 31)
4. ppc.addi(gp_return, gp_return, 11)

5. p = Processor()
6. r = p.execute(c)
7.
7.
7. print r
-
```



Synthetic Population Count

```
1.  def syn_popc (code) :
2.      spu.set_active_code (code)

3.      # Reserve four variable registers
4.      x, count, result, temp = code.acquire_registers (4)
5.
6.      # 'Load' the input vector x from register 5
7.      spu.ai (x, 5, 0)

8.      # Zero count and result
9.      spu.xor (count, count, count)
10.     spu.xor (result, result, result)
11.
12.     # Inline the popc and reduce operations
13.     spu.cntb (temp, x)
14.     spu.sumb (temp, temp, 0)
15.     spu.a (count, count, temp)

16.     for i in range (4) :
17.         spu.a (result, x, result)
18.         spu.rotqbyi (x, x, 4)

19.     # Send the result to the caller
20.     spu.wrch (result, dma.SPU_WrOutMbox)

21.     code.release_registers (x, count, result, temp)
```



Calling `syn_popc`

InstructionStream can be used to collect arbitrary instruction sequences.

```
1. import corepy.arch.spu.platform as env
2. # Load the popc SPE program
3. code = env.InstructionStream()
4. proc = env.Processor()

5. # Create the SPE program
6. syn_popc(code)

7. # Set the input parameters
8. params = env.spu_exec.ExecParams()
9. params.v3 = [0xDEAD, 0xBEEF, 0x1337, 0xCAFE]

10. # Execute the popc program with the parameters
11. count = proc.execute(code, 'mbox', params)
```



InstructionStream, Revisited

InstructionStream manages a sequence of instructions, allocates registers, collects object references, and ensures OS ABI (application binary interface) compliance.

Register Allocation:

```
1. ra = c.acquire_register()
2. fa = c.acquire_register('fp')
3.   ... use the registers ...
4. c.release_register(ra)
5. c.release_register(fa, 'fp')
```

Reference Counting:

```
1. data = Numeric.ones(100)
2. c.add_storage(data)
3.   ... execute the code ...
4. c.release_storage()
```

ABI Compliance:



Prologue:
Save registers
Set vector flags

... user instructions ...

Epilogue:
Restore registers



ISAs

ISA module exposes each machine instruction as a Python function that generates the proper binary coded machine instruction.

add_x Instruction Example

PPC Assembly:

32-bit Binary Layout:

			Opcode	Field	Field	Field	Field	Ext. Opcode	Field
add	D, A, B	(OE = 0, Rc = 0)	31	D	A	B	OE	266	Rc
add.	D, A, B	(OE = 0, Rc = 1)							
addo	D, A, B	(OE = 1, Rc = 0)							
addo.	D, A, B	(OE = 1, Rc = 1)							
			0	5 6	10 11	15 16	20 21 22		30 31

Rc = 1 sets condition register CR0, OE = 1 overflow register is set

1. `ppc.addx(rd, ra, rb)` # asm: add D, A, B
2. `ppc.addx(rd, ra, rb, Rc=1)` # asm: add. D, A, B
3. `ppc.addx(rd, ra, rb, OE=1)` # asm: addo D, A, B
4. `ppc.addx(rd, ra, rb, Rc=1, OE=1)` # asm: addo. D, A, B



Variables

CorePy **Variables** encapsulate a register, backing store, and valid operations for a user-defined data type.

Scalar example:

```
1. a = SignedWord(11)
2. b = SignedWord(31)
3. c = SignedWord(0, reg=gp_return)
4. c.v = (a + b) * 10
5. --> c = 420
```

Vector example:

```
1. a = VecWord([2, 3, 4, 5])
2. b = VecWord([3, 3, 3, 3])
3. c = VecWord(0)
4. c.v = vmin(a, b) * b + 10
5. --> c = [16, 19, 19, 19]
```



Synthetic Population Count, Take 2

```
1.  def syn_popc_var(code):
2.      spu.set_active_code(code)

3.      x = Word(0)
4.      count = Word(0)
5.      result = Word(0)
6.
7.      # 'Load' the input vector x from register 5
8.      x.v = spu.ai.ex(5, 0)

9.      # Inline the popc and reduce operations
10.     count.v = spu.sumb.ex(spu.cntb.ex(x), 0)

11.     for i in range(4):
12.         result.v = result + x
13.         x.v = spu.rotqbyi.ex(x, 4)

14.     # Send the result to the caller
15.     spu.wrch(result, dma.SPU_WrOutMbox)

16.     code.release_registers(x, count, result)
```



Iterators

Iterators enable user-defined loop semantics.

```
1. # Basic Iteration
2. a = SignedWord(c, 0)

3. for i in syn_iter(c, 5):
4.     for j in syn_iter(c, 5, mode = 'ctr'):
5.         a.v = a + 1

6. proc.execute(c)
7. --> a = 25
```



Iterator Examples

```
1. # Array iteration
2. for x in var_iter(c, a): sum.v = sum + x
3. for x in vec_iter(c, a): sum.v = sum + x

4. # Data stream merge
5. for x,y,z,r in zip_iter(c, X,Y,Z,R):
6.     r.v = vmadd(x,y,z)

7. # Auto-parallelization
8. for x in parallel(vec_iter(c, a)): body(x)
9. t1 = proc.execute(c, mode='async', params=[0,2,0])
10. t2 = proc.execute(c, mode='async', params=[1,2,0])

11. # SPE main memory/local stream stream
12. strm = stream_buffer(c, data, buffer size, buffer addr,
13.                     double = True, save = True)
14. for buffer in strm:
15.     for x in vec_iter(c, buffer): x.v = x * x
```



CorePy Population Count, Take 3

```
1. def stream_popc(code):
2.     # Use the popc and reduce from the previous example
3.     popc = syn_popc_var()

4.     x = var.Word(0)
5.     count = var.Word(0)
6.     total = var.Word(0)

7.     # Create a streaming iterator with buffer size < stream size
8.     stream = stream_buffer(code, stream_addr, stream_size * 4,
9.                             buffer_size, lsa)

10.    # Create an array descriptor for the LS buffer
11.    ls_data = memory_desc('I', lsa, buffer_size / 4)

12.    # Count the population in the stream
13.    for buffer in stream:
14.        for x in spu_vec_iter(code, ls_data, addr_reg = buffer):
15.            popc.popc(count, x)

16.    popc.reduce_word(total, count)

17.    # Send the result to the caller
18.    spu.wrch(total, dma.SPU_WrOutMbox)

19.    return
```



Light-weight SPU Debugger

CorePy, combined with the unique execution model of the Cell SPUs, makes it possible to implement a simple, interactive SPU debugger entirely in Python.

Example:

```
# Create the synthetic program
code = InstructionStream()
spu.set_active_code(code)
```

```
spu.ai(127, 0, 1)
spu.ai(126, 0, 2)
spu.ai(125, 0, 3)
spu.brnz(125, 2)
spu.ai(124, 0, 4)
spu.ai(123, 0, 5)
```

```
# Execute it with a debug processor
proc = DebugProcessor()
r = proc.execute(code)

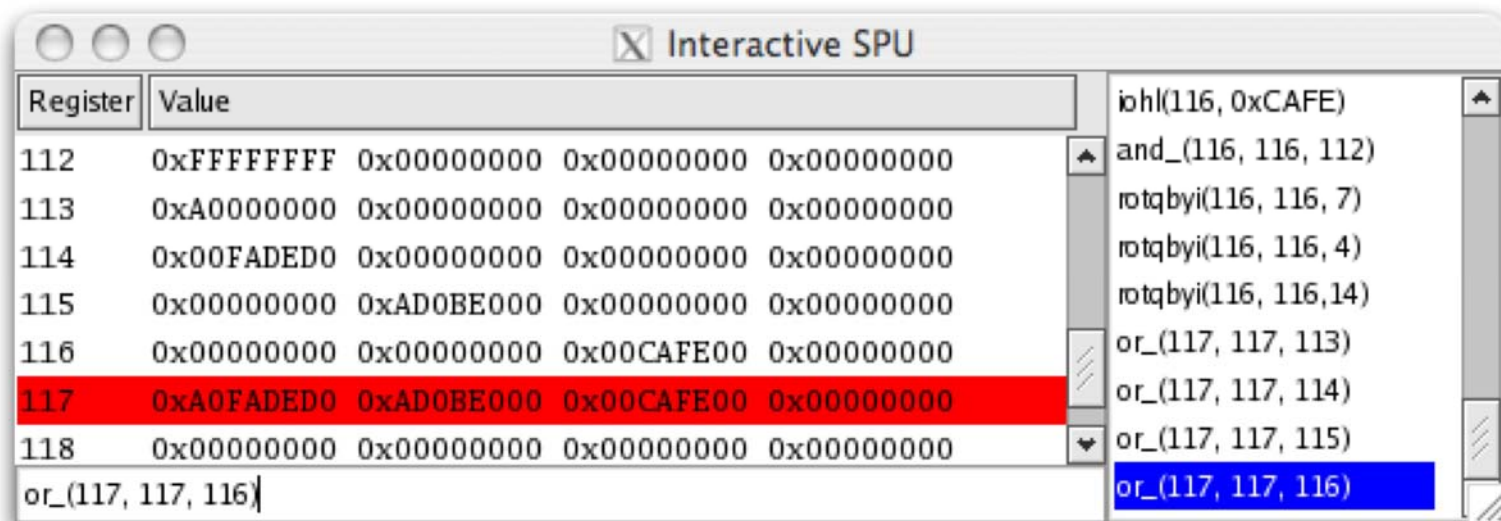
while r is not None:
    regs = proc.get_regs()
    print [reg[0] for reg in regs[123:]]
    r = proc.nexti()
```

```
# Output
--> [1, 0, 0, 0, 0]
--> [1, 2, 0, 0, 0]
--> [1, 2, 3, 0, 0]
--> [1, 2, 3, 0, 5]
```



Interactive SPU

*The **Interactive SPU** module provides a graphical or command line interface to a running SPU.*



Other Applications

Chemical fingerprint comparison

Rat brain neural analysis

Cell BLAST prototype

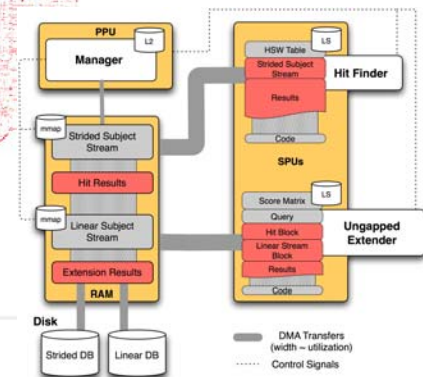
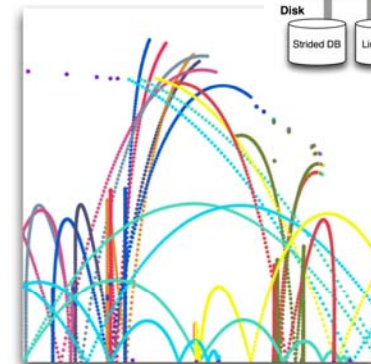

SPU big num library

PS3 framebuffer support

VMX particle system

Scalable PPC dgemm


















Interactive SPE



Register	Value	
112	0xFFFFFFFF 0x00000000 0x00000000 0x00000000	lshi(116, 0xCAFE)
113	0xA0000000 0x00000000 0x00000000 0x00000000	and_(116, 116, 112)
114	0x00FADED0 0x00000000 0x00000000 0x00000000	rotqbyi(116, 116, 7)
115	0x00000000 0xAD0BE000 0x00000000 0x00000000	rotqbyi(116, 116, 4)
116	0x00000000 0x00000000 0x00CAFE00 0x00000000	rotqbyi(116, 116, 14)
117	0xADFADED0 0xAD0BE000 0x00CAFE00 0x00000000	or_(117, 117, 113)
118	0x00000000 0x00000000 0x00000000 0x00000000	or_(117, 117, 114)
		or_(117, 117, 115)
		or_(117, 117, 116)



CorePy Status

	<u>SPE</u>	<u>PPE</u>	<u>VMX</u>
Processor			()
InstructionStream			()
NativeInstStream			
Variables			
Iterators			
Debug Print			()
Interactive Debug			
Operating System	CorePy Linux	Linux OS X	Linux OS X

 - Beta  - Alpha

- Availability
 - Source distribution
 - Evaluation license
- Tested Platforms
 - IBM Cell Blades
 - PS3
 - Apple G4/G5 Macs
- Caveat
 - libspe 1.x



Conclusion

CorePy provides a framework for rapid development of high-performance applications on the Cell/B.E.

CorePy Benefits:

- Control C/C++/etc SPU programs from Python
- Develop assembly-level code using an intuitive API
- Design and debug new SPU SIMD algorithms interactively



Thank You!

Special Thanks To:

IBM Cell Ecosystem Team, especially:
Hema Reddy, Gordon Ellison, Jennifer Turner

The Lilly Foundation

<http://www.corepy.org>
chemuell@cs.indiana.edu

“CorePy makes assembly fun again!”
-Alex Breuer

