

# Synthetic Programming on the Cell BE

---

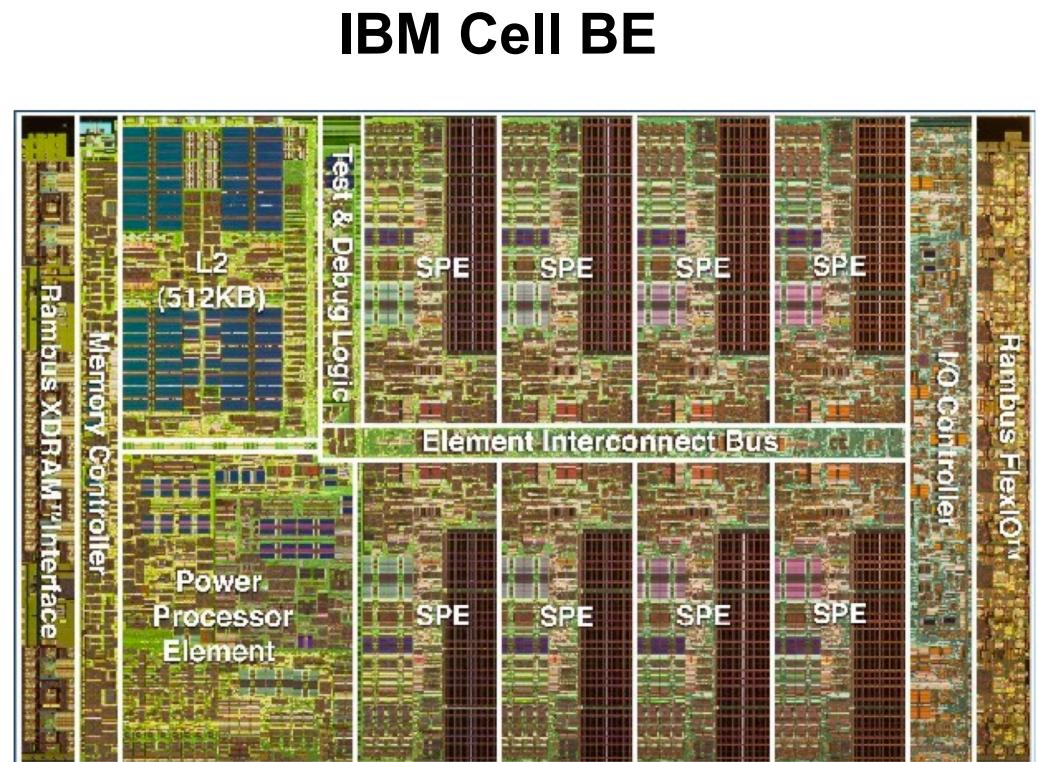
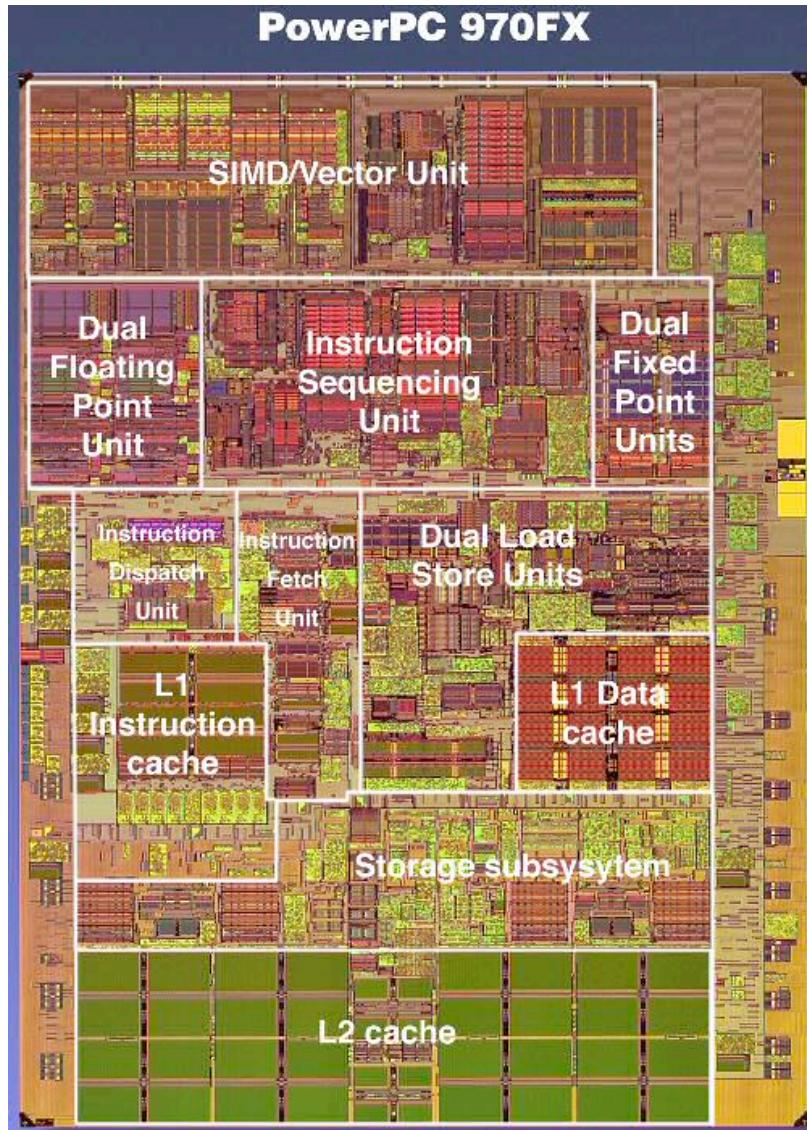
Christopher Mueller, Benjamin Martin and  
Andrew Lumsdaine

October 25, 2006

UT/IBM Cell Workshop

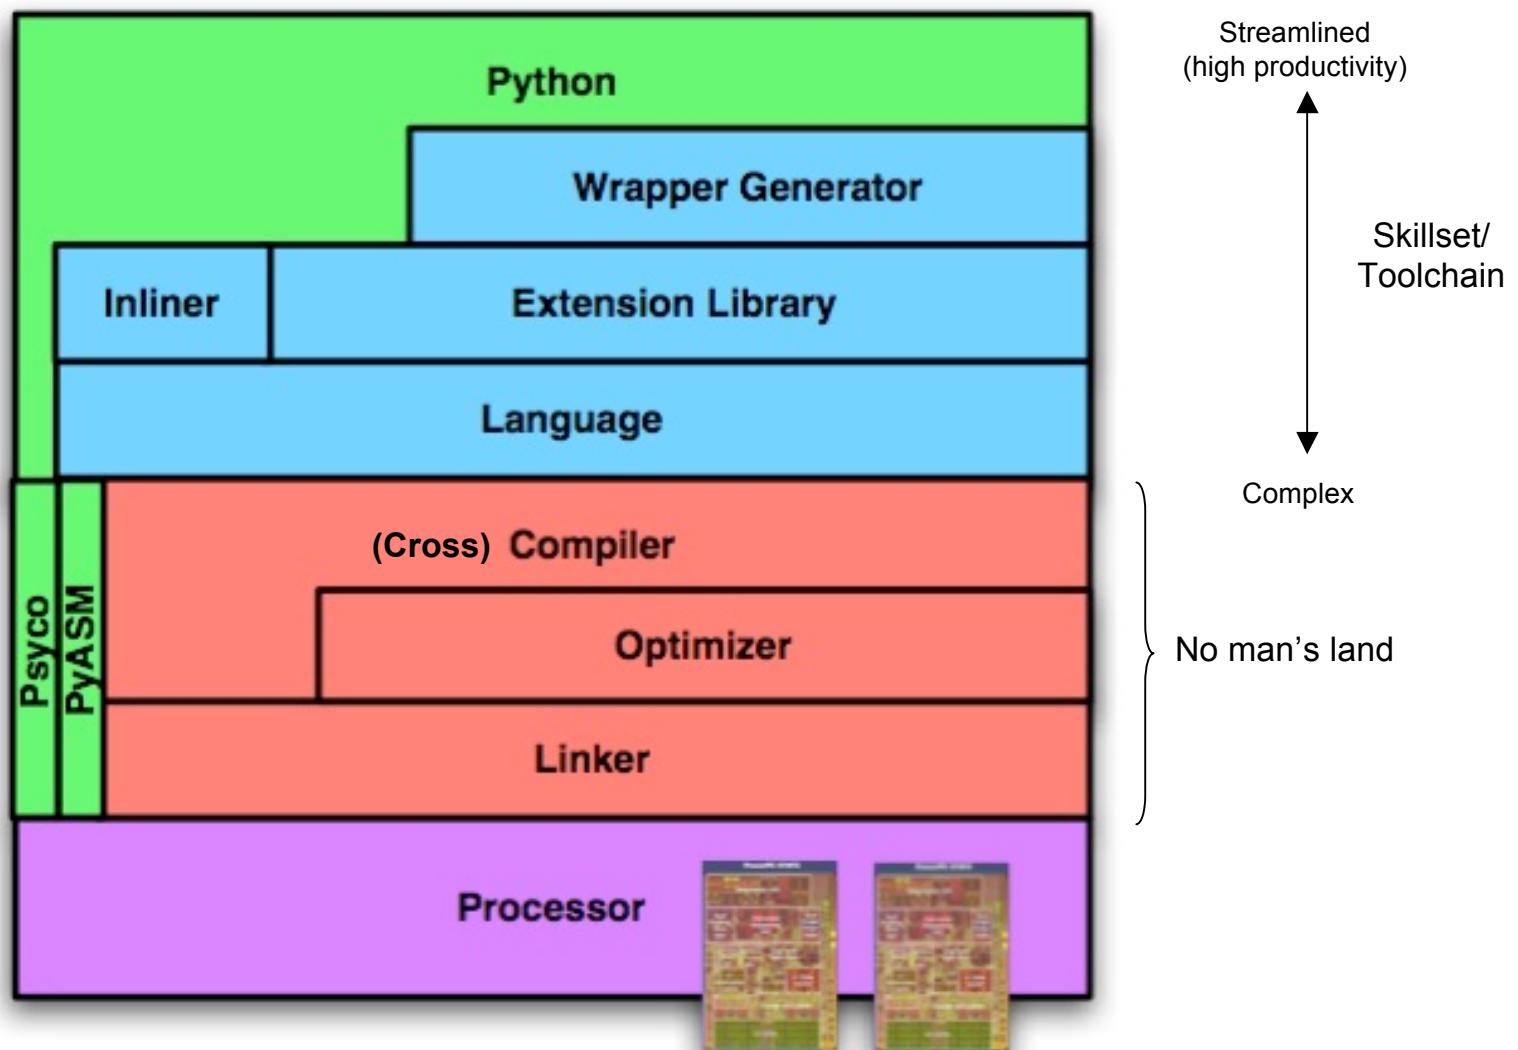


# Modern PowerPC Processors

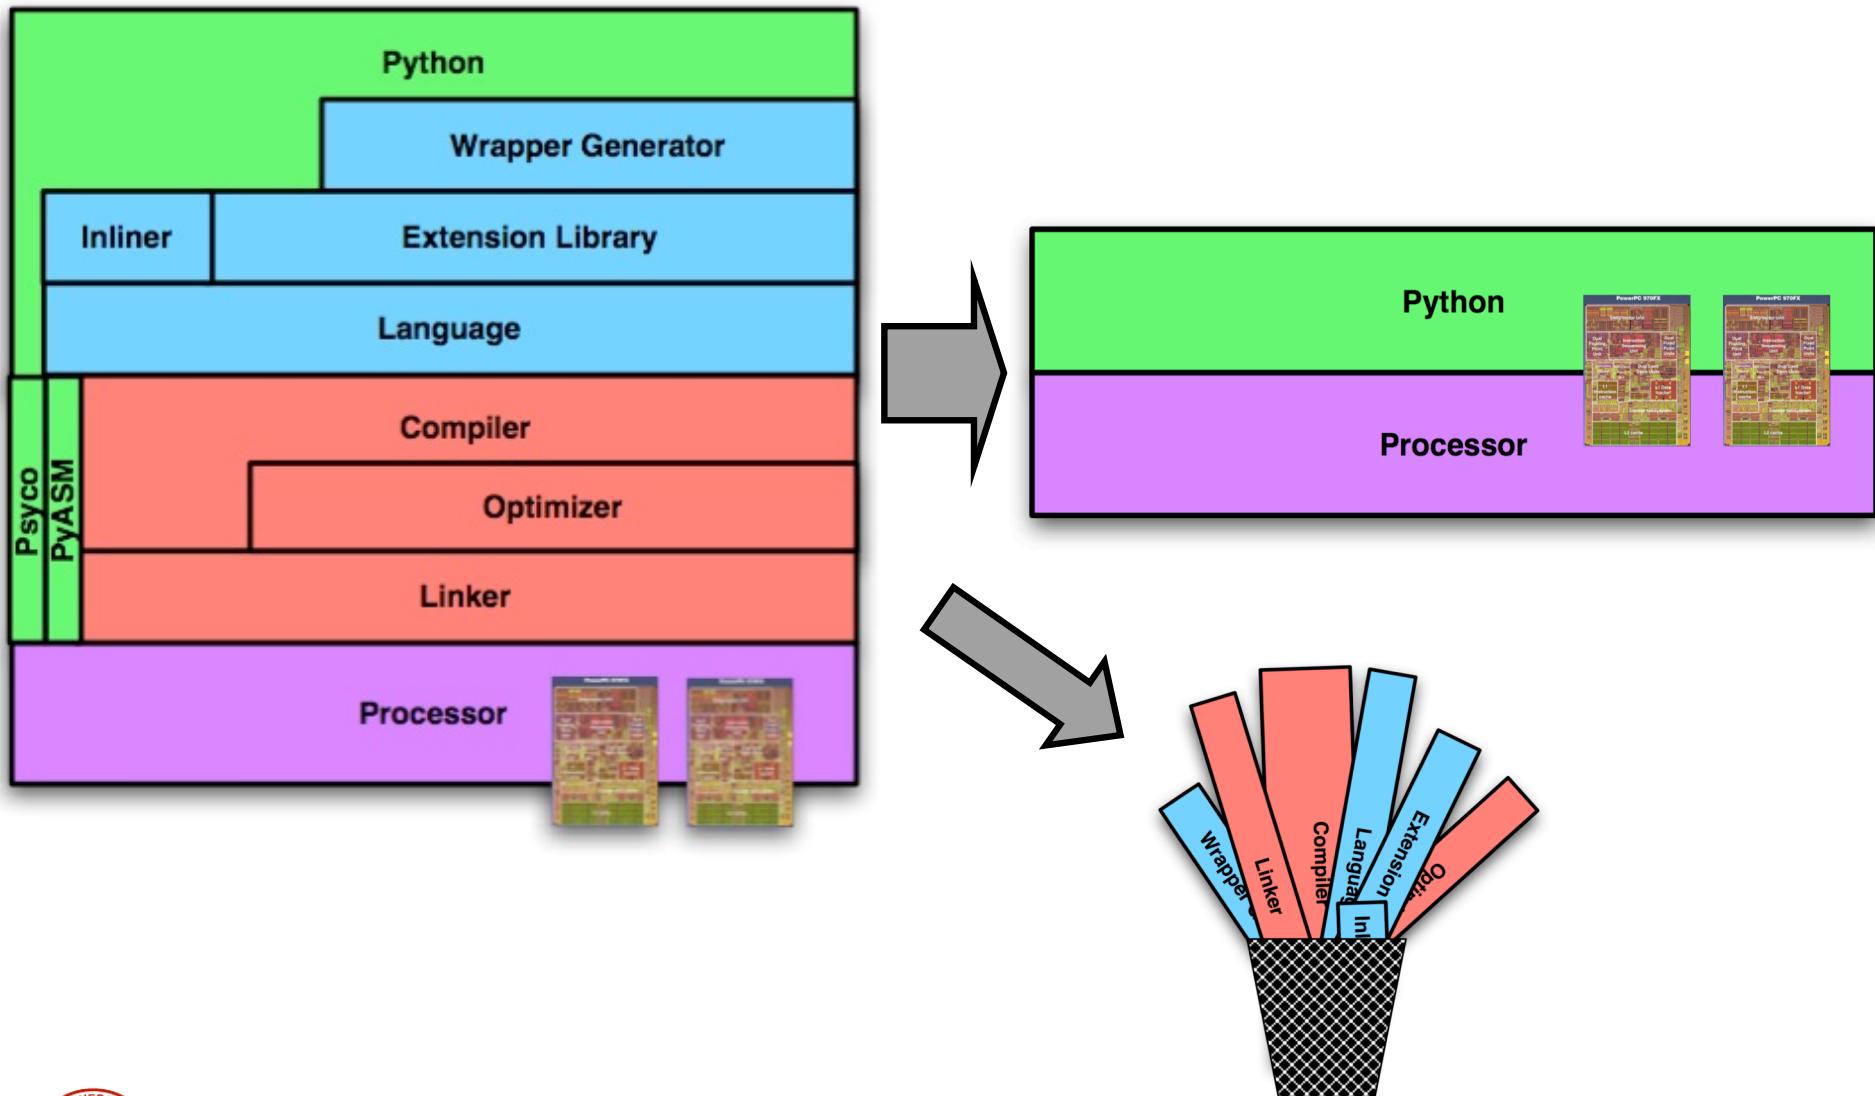


# Programming Modern Processors

## A Python HPC Stack

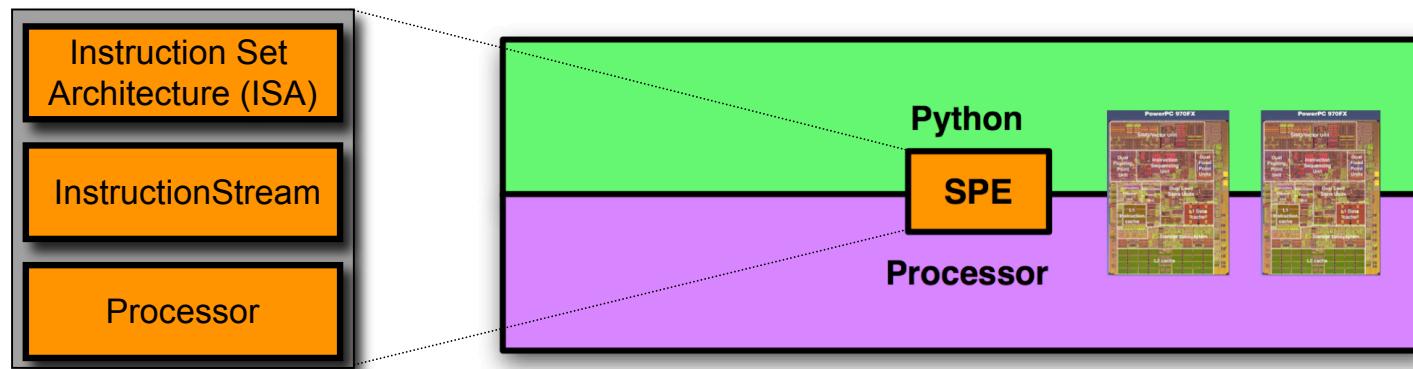


# What if..?



# Synthetic Programming Environment

*The Synthetic Programming Environment (SPE) is a user-level library for synthesizing high-performance computational kernels at runtime.*



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

1. `c = InstructionStream()`  
2. `c.add(ppc.addi(gp_return, 0, 31))`  
3. `c.add(ppc.addi(gp_return, gp_return, 11))`

4. `p = Processor()`  
5. `r = p.execute(c)`  
6. `print r`  
7. `--> 42`



# Synthetic Programming

---

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

## Terminology:

- **Synthetic programs:** complete instruction streams
- **Synthetic components:** meta-programming functions and objects that synthesize partial instruction sequences
- **Synthetic \*:** modifier to denote synthesized code, e.g. *synthetic loop*



# Example: add/reduce

## Synthetic:

```
1. # r_* are processor registers, e.g. r_sum = 3
2. c = InstructionStream()
3. n = 10000000
4. a = array(range(n))

5. # Load a pointer to the array
6. c.load_word(r_addr, addr(a))

7. # Set the counter
8. c.load_word(r_temp, n)
9. c.add(ppc.mtctr(r_temp))

10. # Zero the sum
11. c.add(ppc.addi(r_sum, 0, 0))

12. # Set a loop label and load the next value
13. start = c.add(ppc.lwz(r_current, r_addr, 0))

14. # Update the sum
15. c.add(ppc.add(r_sum, r_sum, r_current))

16. # Increment the pointer
17. c.add(ppc.addi(r_addr, r_addr, 4))

18. # Decrement the counter and loop
19. next = code.size() + 1
20. c.add(ppc.bdnz(-((next - start) * 4)))

21. result = proc.execute(c)
```

## Numeric:

```
1. n = 10000000
2. a = Numeric.arange(n)
3. sum = Numeric.add.reduce(a)
```

## Python:

```
1. n = 10000000
2. a = array(range(n))
3. sum = 0

4. for i in a:
5.     s += i
```

| Test      | Time (s) |
|-----------|----------|
| Synthetic | 0.019    |
| Numeric   | 0.054    |
| Python    | 3.64     |

2.5 MHz G5, 3 GB, timed code gen and reduce operations only



# Execution - PPU

---

*Synthetic code is executed on the PPU using a thin native interface. ABI compliance is managed in Python.*

## Parameters and return values:

1. `# Pass two parameters`
2. `v = Params()`
3. `v.p1 = 31`
4. `v.p2 = 11`
5. `p.execute(c, params=v)`
  
6. `# result in gp_return`
7. `r = p.execute(c)`
  
8. `# result in fp_return`
9. `f = p.execute(c, 'fp')`

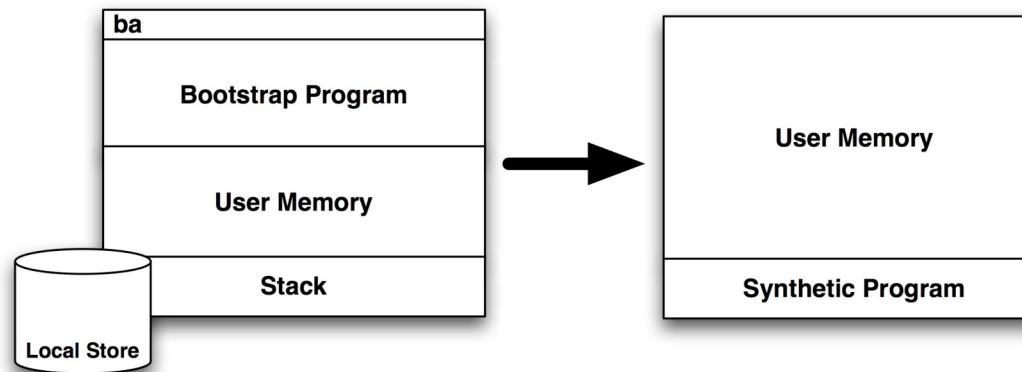
## Asynchronous execution:

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)`



# Execution - SPU

*Synthetic code is generated on the PPU and executed on the SPU using a native interface and an SPU bootstrap program.*



**Single SPU, blocking execution:**

1. `# Execute and stop`
2. `c.add(spu.stop(0x7))`
  
3. `# result from stop`
4. `r = p.execute(c)`
5. `--> r = 0x7`



**8 SPUs, non-blocking execution:**

1. `ids = p.execute(c,`
2. `mode = 'async',`
3. `n_spus = 8)`
  
4. `for id in ids:`
5. `p.join(id)`

# Synthetic Variables

***Synthetic variables*** encapsulate a register, backing store, and valid operations for a processor data type.

## Scalar variables (PPU):

1. `a = var(0)`
2. `b = var(1.2)`
3. `c = var(4, reg=gp_return)`
4. `b.store()`
5. `... execute code ...`
  
6. `print b.value`
7. `--> 1.2`

## Vector variables (PPU/SPU):

1. `a = vector(1)`
2. `b = vector([1,2,3,4])`
3. `Altivec.vadd(a.reg,`  
   `a.reg,`  
   `b.reg)`
  
6. `... execute code ...`
7. `--> a = [2,3,4,5]`



# Synthetic Expressions

***Synthetic expressions*** use synthetic variables and Python operators to generate instruction sequences for arbitrary expressions.

## Scalar expressions (PPU):

```
1. a = var(11)
2. b = var(31)
3. c = var(0, reg=gp_return)

4. c = (a + b) * 10
5. --> c = 420
```

## Vector expressions (PPU/SPU):

```
1. a = vector([2,3,4,5])
2. b = vector([3,3,3,3])
3. c = vector(0)

4. c = vmin(a, b) * b + 10
5. --> c = [16, 19, 19, 19]

6. d = selb((a < b), a, b)
7. --> d = [2,3,3,3]
```



# Synthetic Iterators

---

***Synthetic iterators*** synthesize code to manage loops on both the PPU and SPU.

```
1. # Basic Iteration (PPU/SPU)
2. a = var(c, 0)

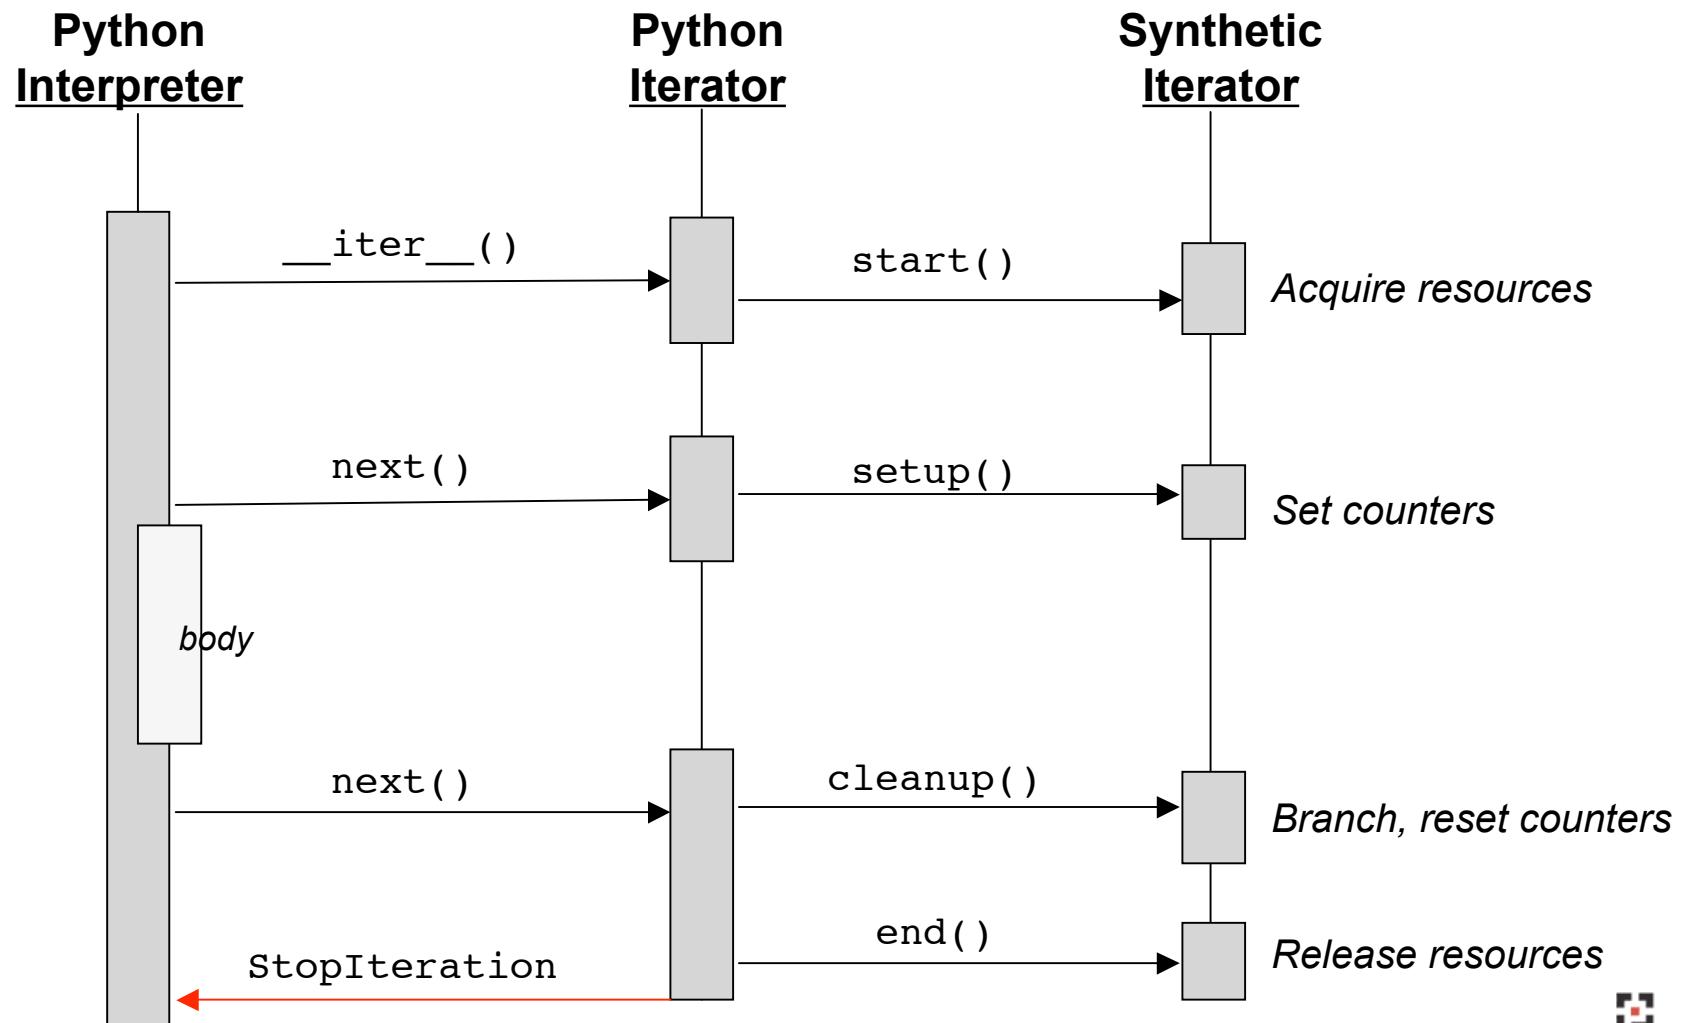
3. for i in syn_iter(c, 5):
4.     for j in syn_iter(c, 5):
5.         a.v = a + i + j

6. proc.execute(c)
7. --> a = 100
```



# How it works...

```
1. for i in syn_iter(c, 20):  
2.     sum.v = sum + i
```



# Array Iterators

**Array iterators** iterate over values in an array, one element at a time or in vector-size steps.

```
1. a = arange(32)
2. # Serial add/reduce (PPU)
3. for i in var_iter(c, a):
4.     sum.v = sum + i
5. # SIMD add/reduce (PPU/SPU)
6. for i in vec_iter(c, a):
7.     sum.v = sum + i
```

```
1. # r_* are processor registers, e.g. r_sum = 3
2. c = InstructionStream()
3. n = 10000000
4. a = array(range(n))

5. # Load a pointer to the array
c.load_word(r_addr, addr(a))

6. # Set the counter
c.load_word(r_temp, n)
c.add(ppc.mtctr(r_temp))

7. # Zero the sum
c.add(ppc.addi(r_sum, 0, 0))

8. # Set a loop label and load the next value
start = c.add(ppc.lwz(r_current, r_addr, 0))

9. # Update the sum
c.add(ppc.add(r_sum, r_sum, r_current))

10. # Increment the pointer
c.add(ppc.addi(r_addr, r_addr, 4))

11. # Decrement the counter and loop
next = code.size() + 1
c.add(ppc.bdnz(-(next - start) * 4))

12. result = proc.execute(c)

13.
14.
15.
16.
17.
18.
19.
20.
21.
```



# Pythonic Iterators

---

***syn\_range*** and ***zip\_iter*** duplicate the functionality of Python's range and zip.

```
1. for i in syn_range(c, 10, 20, 2):  
2.     a.v = a + i  
  
3. --> a = 70
```

```
1. # X,Y,Z,R are vec_iters  
  
2. for x,y,z,r in zip_iter(c, X,Y,Z,R):  
3.     r = vmadd(x,y,z)  
  
4. --> r[i] = x[i]*z[i] + y[i]
```



# High-Performance Iterators

**High-Performance iterators** modify existing iterators for optimized code generation.

```
1. # Unroll the loop 3 times
2. for x,y,z,r in unroll(zip_iter(c, X,Y,Z,R), 3):
3.     r = vmadd(x,y,z)

1. # Divide the data for parallel execution
2. c = ParallelInstructionStream()

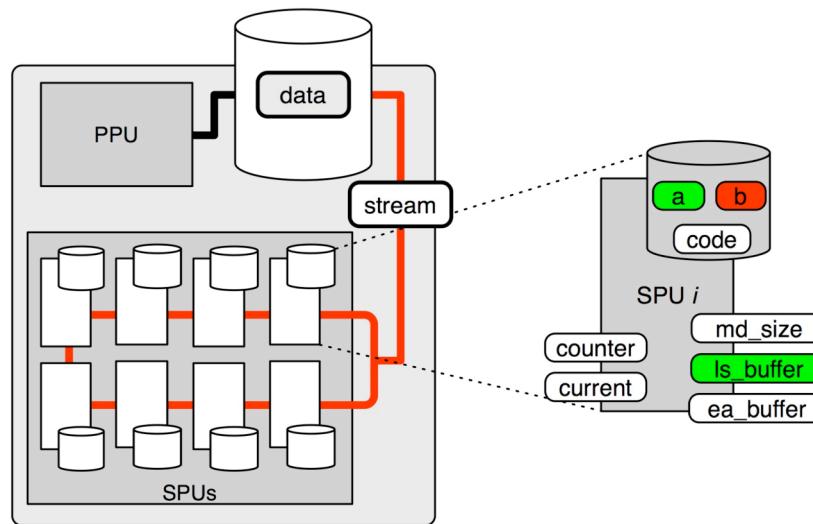
3. for x,y,z,r in parallel(zip_iter(c, X,Y,Z,R)):
4.     r = vmadd(x,y,z)

5. t1 = proc.execute(c, mode='async', params=[0,2,0])
6. t2 = proc.execute(c, mode='async', params=[1,2,0])
```



# Stream Buffer Iterators

***stream\_buffer iterators*** move data between main memory and the SPU local store.



```
1. stream = stream_buffer(c, data, buffer_size, buffer_addr,  
                           double = True, save = True)  
2.  
3. stream = parallel(stream)  
  
4. for buffer in stream:  
5.     for x in vec_iter(c, buffer):  
6.         x.v = x * x  
  
7. proc.execute(c, n_spus = 4)
```



# Communication

---

*All libspe functions are available to the host program.*

```
# Send a message to the PPU
spu_write_out_mbox(code, 0xDEADBEAF)

# Get a message from the PPU
reg = spu_read_in_mbox(code)

# And send it back
code.add(spu.wrch(SPU_WrOutMbox, reg))

# Execute the synthetic program
spe_id = proc.execute(code, mode='async')
write_in_mbox(spe_id, 0x88CAFE)

while stat_out_mbox(spe_id) != 0:
    msg = read_out_mbox(spe_id))

proc.join(spe_id)
```



# Example: BLAST

## **Step 1: Find all high-scoring words in the query**

|  |          |   |            |   |   |            |   |    |            |   |    |
|--|----------|---|------------|---|---|------------|---|----|------------|---|----|
| <b>AABCD</b> FNNMPMLLASLDA<br> | $\times$ | <table style="margin-left: auto; margin-right: 0;"> <tr> <td><b>AAA</b></td> <td>=</td> <td>2</td> </tr> <tr> <td><b>AAB</b></td> <td>=</td> <td>-3</td> </tr> <tr> <td><b>AAC</b></td> <td>=</td> <td>14</td> </tr> </table> | <b>AAA</b> | = | 2 | <b>AAB</b> | = | -3 | <b>AAC</b> | = | 14 |
| <b>AAA</b>   | =        | 2   |            |   |   |            |   |    |            |   |    |
| <b>AAB</b>   | =        | -3  |            |   |   |            |   |    |            |   |    |
| <b>AAC</b>   | =        | 14  |            |   |   |            |   |    |            |   |    |

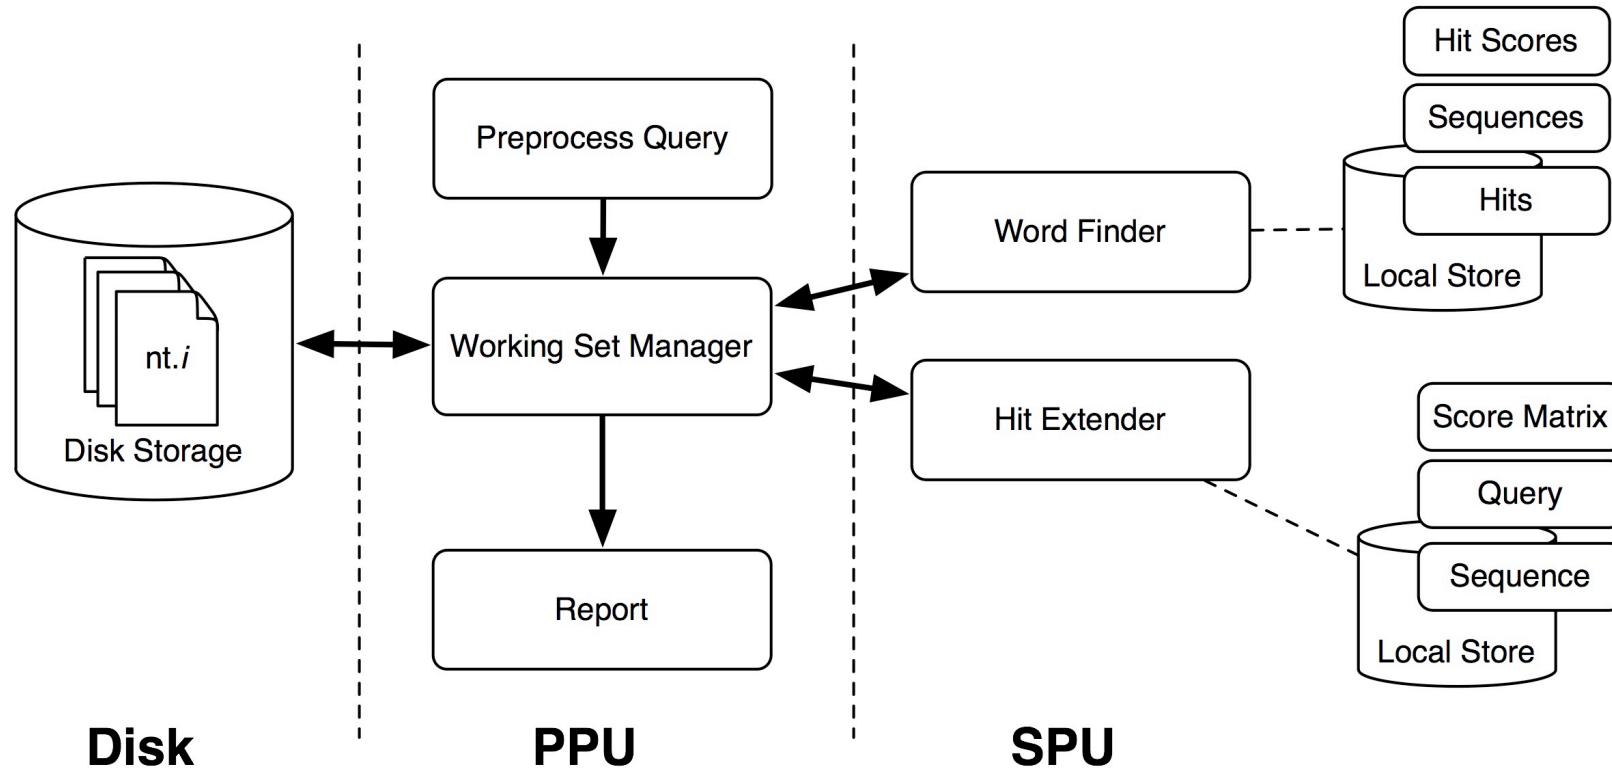
**Step 2:** For each database subject, find all instances of the high scoring words.

**Step 3:** Extend the matches in each direction to generate high scoring pairs.

**AABCD**FNNMPMLLASLDA****  
**ABMNSM**MDN**SBCDEMLAPPMLL...******



# Synthetic BLAST on the Cell BE



# Conclusion

---

- Synthetic Programming
  - Rapid development at the hardware level
  - Python libraries available for admin tasks
  - Meta-programming enables new software architectures (see also C++ Boost libraries)
- ...on the Cell BE
  - Effective model for managing SPU execution
  - More fun than cross-compiling
- Wish list
  - SPU vector lookup
  - `mfc_gets` safely exposed on PPU



Thank you!

---

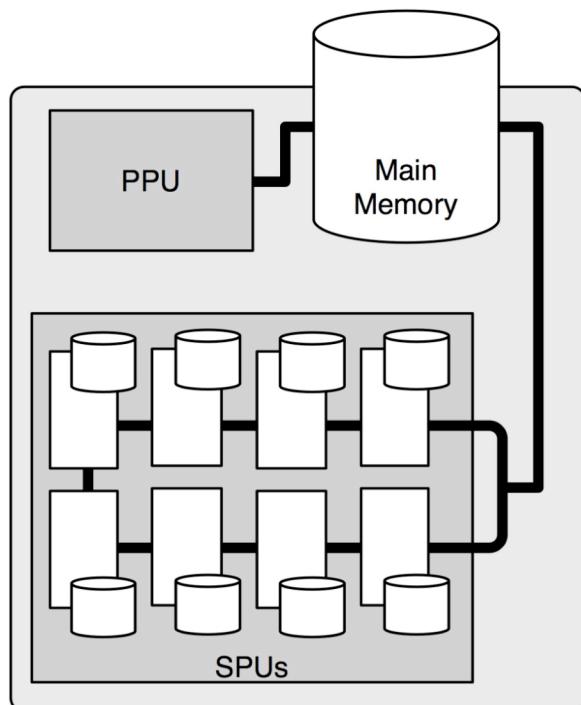
**www.synthetic-programming.org**

**chemuell@cs.indiana.edu**



# Multi-Core SPE: Cell BE (1)

The the **Cell BE** contains two general purpose PPC cores and eight special purpose **Synergistic Processing Element\*** cores.



Programmer's view of the Cell BE

## Challenges and Opportunities:

- Three ISAs - PowerPC, SPU, AltiVec
- Explicit control of data transfers between main memory and SPUs
- In-order execution on SPUs
- Limited access to Cell BE blades
- Linux port designed for C-based tool chain



\*SPE has a specific meaning in Cell BE jargon. In this discussion, SPE will refer to the Synthetic Programming Environment and SPU the Cell BE's Synergistic Processing Unit.

# Multi-Core SPE: Cell BE (2)

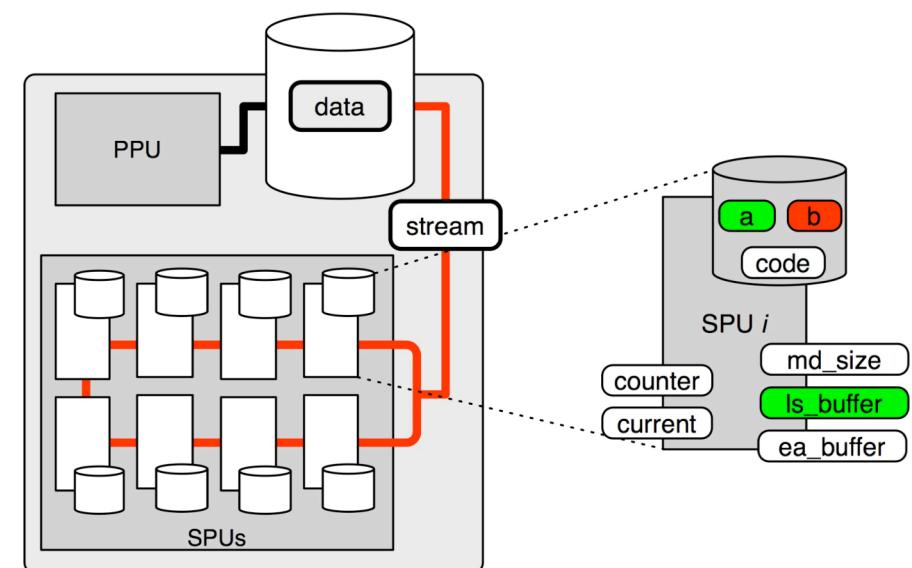
*The Synthetic Programming Environment for the Cell BE includes support for **SPU code synthesis** and **execution**, using both the **SPU ISA** and the higher-level **expression** and **iterator** libraries.*

## Double Buffered Stream Operation:

```
1. c = SPU.InstructionStream()
2. s = stream_buffer( ... )
3. stream = parallel(s)
4. md      = memory_desc( ... )

5. for buffer in stream:
6.   for x in spu_vec_iter(c, md):
7.     x.v = x + x

8. r = proc.execute(c, n_spus = 8)
```



Asynchronous execution and Cell BE communication mechanisms enable much more complex and powerful system architectures.

