

Introduction to Concurrency

Adapted by BCP from lectures
by Maurice Herlihy at Brown

From the New York Times ...


SAN FRANCISCO, May 7, 2004 -
Intel said on Friday that it was
scrapping its development of two
microprocessors, a move that is a shift
in the company's business strategy....



© 2007 Herlihy and Shavit


2

Moore's Law



3


On Your Desktop: The Uniprocessor



© 2007 Herlihy and Shavit

4

In the Enterprise: The Shared Memory Multiprocessor (SMP)




© 2007 Herlihy and Shavit

5

Your New Desktop: The Multicore processor (CMP)

All on the
same chip



Sun
T2000
Niagara



© 2007 Herlihy and Shavit

6

Multicores Are Here

- "Intel ups ante with 4-core chip. New microprocessor, due this year, will be faster, use less electricity..." [San Fran Chronicle]
- "AMD will launch a dual-core version of its Opteron server processor at an event in New York on April 21." [PC World]
- "Sun's Niagara...will have eight cores, each core capable of running 4 threads in parallel, for 32 concurrently running threads." [The Inquirer]



© 2007 Herlihy and Shavit

7

Why do we care?

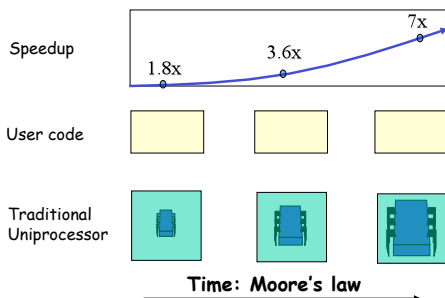
- Time no longer cures software bloat
 - The "free ride" is over
- When you double the work your program is doing...
 - ...you can't just wait 6 months for it to run the same speed again!
 - Your software must somehow exploit twice as much concurrency



© 2007 Herlihy and Shavit

8


Traditional Scaling Process



© 2007 Herlihy and Shavit

9


Multicore Scaling Process



© 2007 Herlihy and Shavit

10


Real-World Scaling Process



© 2007 Herlihy and Shavit


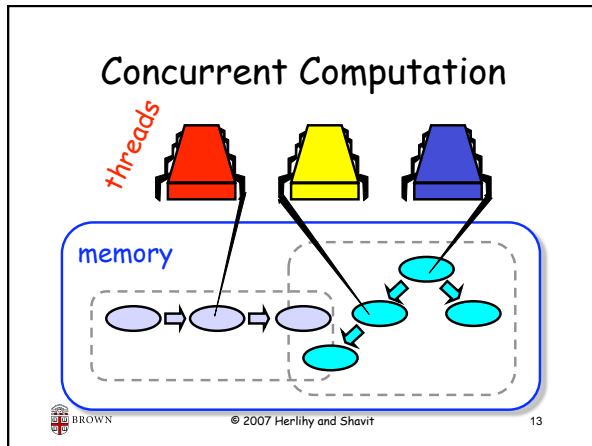
11

Sequential Computation



© 2007 Herlihy and Shavit

12




- ## Model Summary
- Multiple **threads**
 - Sometimes called **processes**
 - Single shared **memory**
 - Unpredictable asynchronous delays
- © 2007 Herlihy and Shavit 15

- ## Road Map
- Today: background on concurrency
 - Monday: semantics of Haskell's basic concurrency primitives (threads/ MVars)
 - Wednesday: thread programming
 - Following week: Software Transactional Memory (STM)
- © 2007 Herlihy and Shavit 16

- ## Concurrency Jargon
- Hardware
 - Processors
 - Software
 - Threads, processes
 - Sometimes OK to confuse them, sometimes not.
- © 2007 Herlihy and Shavit 17

- ## Parallel Primality Testing
- Challenge
 - Print primes from 1 to 10^{10}
 - Given
 - Ten-processor multiprocessor
 - One thread per processor
 - Goal
 - Get ten-fold speedup (or close)
- © 2007 Herlihy and Shavit 18

Load Balancing



Idea:

- Split the work evenly
- Each thread tests range of 10^9



© 2007 Herlihy and Shavit

19

Procedure for Thread i

```
void primePrint {  
  int i = ThreadID.get(); // IDs in {0..9}  
  for (j = i*109+1, j<(i+1)*109; j++) {  
    if (isPrime(j))  
      print(j);  
  }  
}
```

Note Herlihy's slightly awkward pseudocode notation for Haskell



© 2007 Herlihy and Shavit

20

Issues

- Higher ranges have fewer primes
- Yet larger numbers harder to test
- Thread workloads
 - Uneven
 - Hard to predict
- Need **dynamic** load balancing



© 2007 Herlihy and Shavit

21

Issues

- Higher ranges have fewer primes
- Yet larger numbers harder to test
- Thread workloads
 - Uneven
 - Hard to predict
- Need **dynamic** load balancing


rejected



© 2007 Herlihy and Shavit

22

Shared Counter



© 2007 Herlihy and Shavit

23

Procedure for Thread i

```
int counter = new Counter(1);  
  
void primePrint {  
  long j = 0;  
  while (j < 1010) {  
    j = counter.getAndIncrement();  
    if (isPrime(j))  
      print(j);  
  }  
}
```



© 2007 Herlihy and Shavit

24

Procedure for Thread i

```
Counter counter = new Counter(1);
```

```
void primePrint {
    long j = 0;
    while (j < 1010) {
        j = counter.getAndIncrement();
        if (isPrime(j))
            print(j);
    }
}
```


Shared counter object



© 2007 Herlihy and Shavit

25

Where Things Reside



© 2007 Herlihy and Shavit

26

Procedure for Thread i

```
Counter counter = new Counter(1);
```

```
void primePrint {
    long i = 0;
    while (j < 1010) {
        j = counter.getAndIncrement();
        if (isPrime(j))
            print(j);
    }
}
```

Stop when every value taken



© 2007 Herlihy and Shavit

27

Procedure for Thread i

```
Counter counter = new Counter(1);
```

```
void primePrint {
    long j = 0;
    while (j < 1010) {
        j = counter.getAndIncrement();
        if (isPrime(j))
            print(j);
    }
}
```

Increment & return each new value



© 2007 Herlihy and Shavit

28

Counter Implementation

```
public class Counter {
    private long value;

    public long getAndIncrement() {
        return value++;
    }
}
```



© 2007 Herlihy and Shavit

29

Counter Implementation

```
public class Counter {
    private long value;

    public long getAndIncrement() {
        return value++;
    }
}
```

OK for single thread,
not for concurrent threads
(i.e., not "thread safe")



© 2007 Herlihy and Shavit

30

What It Means

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        return value++;  
    }  
}
```



© 2007 Herlihy and Shavit

31

What It Means

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        return value++;  
    }  
}
```


temp = value;
value = value + 1;
return temp;



© 2007 Herlihy and Shavit

32


Not so good...



© 2007 Herlihy and Shavit

33

Is this problem inherent?



If we could only glue reads and writes...



© 2007 Herlihy and Shavit

34

Challenge

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        temp = value;  
        value = temp + 1;  
        return temp;  
    }  
}
```



© 2007 Herlihy and Shavit

35

Challenge

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        temp = value;  
        value = temp + 1;  
        return temp;  
    }  
}
```

Make these steps atomic (indivisible)




© 2007 Herlihy and Shavit

36

Hardware Solution

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        temp = value;  
        value = temp + 1;  
        return temp;  
    }  
}
```



ReadModifyWrite()
instruction



© 2007 Herlihy and Shavit

37

An Aside: Java™

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        synchronized {  
            temp = value;  
            value = temp + 1;  
        }  
        return temp;  
    }  
}
```

Mutual Exclusion



© 2007 Herlihy and Shavit

38

An Aside: Java™

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        synchronized {  
            temp = value;  
            value = temp + 1;  
        }  
        return temp;  
    }  
}
```



BROWN

© 2007 Herlihy and Shavit

An Aside: Java™

```
public class Counter {  
    private long value;  
  
    public long getAndIncrement() {  
        synchronized {  
            temp = value;  
            value = temp + 1;  
        }  
        return temp;  
    }  
}
```

Haskell uses slightly
different primitives to
achieve the same effect


Synchronized block



© 2007 Herlihy and Shavit

40

Mutual Exclusion, or "Alice & Bob share a pond"



© 2007 Herlihy and Shavit



41

Alice has a pet



© 2007 Herlihy and Shavit

42



Formalizing the Problem

- Two types of formal properties in asynchronous computation:
 - **Safety** Properties
 - Nothing bad happens ever
 - **Liveness** Properties
 - Something good happens eventually

© 2007 Herlihy and Shavit

Formalizing our Problem

- **Mutual Exclusion**
 - Both pets never in pond simultaneously
 - This is a **safety** property
- **No Deadlock**
 - if only one wants in, it gets in
 - if both want in, one gets in
 - This is a **liveness** property

© 2007 Herlihy and Shavit

Simple Protocol

- **Idea**
 - Just look at the pond, see if it is empty, and release pet if so
- **Gotcha**
 - Both look at the same instant
 - Both release pet
 - Bad thing happens in pond

© 2007 Herlihy and Shavit

Telephone Protocol

- **Idea**
 - Bob calls Alice (or vice-versa)
- **Gotcha**
 - Alice in shower when Bob calls
 - Bob recharging phone battery when Alice calls
 - Alice out shopping for pet food when Bob calls...

© 2007 Herlihy and Shavit

Patient Telephone Protocol


- Idea
 - Bob calls Alice (or vice-versa) and lets phone ring until Alice answers
- Gotcha
 - Alice goes on vacation for a month...
- Lesson
 - Need to be able to leave **persistent** messages (like writing, not speaking)



© 2007 Herlihy and Shavit

49


Can Protocol



© 2007 Herlihy and Shavit

50


Bob conveys a bit



© 2007 Herlihy and Shavit

51

Bob conveys a bit



© 2007 Herlihy and Shavit

52

Can Protocol


- Idea
 - Cans on Alice's windowsill
 - Strings lead to Bob's house
 - Bob pulls strings, knocks over cans
- Gotcha
 - Cans cannot be reused
 - Bob runs out of cans



© 2007 Herlihy and Shavit


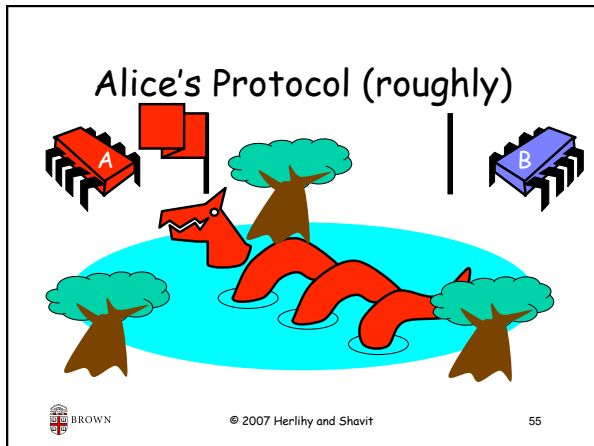
53

Flag Protocol



© 2007 Herlihy and Shavit

54



Alice's Protocol

- Raise flag
- Wait until Bob's flag is down
- Unleash pet
- Lower flag when pet returns

BROWN © 2007 Herlihy and Shavit 57

Bob's Protocol

- Raise flag
- Wait until Alice's flag is down
- Unleash pet
- Lower flag when pet returns

danger!

BROWN © 2007 Herlihy and Shavit 58

Bob's Protocol (2nd try)

- Raise flag
- While Alice's flag is up
 - Lower flag
 - Wait for Alice's flag to go down
 - Raise flag
- Unleash pet
- Lower flag when pet returns

Bob defers to Alice

BROWN © 2007 Herlihy and Shavit 59

The Flag Principle

- Raise the flag
- Look at other's flag
- Flag Principle:
 - If each raises and looks, then
 - Last to look must see both flags up

BROWN © 2007 Herlihy and Shavit 60

Proof of Mutual Exclusion


- Assume both pets in pond
 - Derive a contradiction
 - By reasoning backwards
- Consider the last time Alice and Bob each looked before letting the pets in
- Without loss of generality assume Alice was the last to look...



© 2007 Herlihy and Shavit

61

Proof



Alice must have seen Bob's Flag. A Contradiction

Proof of No Deadlock

- If only one pet wants in, it gets in.



© 2007 Herlihy and Shavit

63

Proof of No Deadlock

- If only one pet wants in, it gets in.
- Deadlock requires both continually trying to get in.



© 2007 Herlihy and Shavit

64

Proof of No Deadlock

- If only one pet wants in, it gets in.
- Deadlock requires both continually trying to get in.
- If Bob sees Alice's flag, he gives her priority (a gentleman...)

QED



© 2007 Herlihy and Shavit

65

Remarks

- Protocol is **unfair**
 - Bob's pet might never get in
- Protocol uses **waiting**
 - If Bob is eaten by his pet, Alice's pet might never get in



© 2007 Herlihy and Shavit

66

Moral of Story

- Mutual Exclusion **cannot be solved** by
 - transient communication (cell phones)
 - interrupts (cans)
- It **can be solved** by
 - one-bit shared variables
 - that can be read or written



© 2007 Herlihy and Shavit

67

The Fable Continues

- Alice and Bob fall in love & marry



© 2007 Herlihy and Shavit

68

The Fable Continues

- Alice and Bob fall in love & marry
- Then they fall out of love & divorce
 - She gets the pets
 - He has to feed them



© 2007 Herlihy and Shavit

69

The Fable Continues


- Alice and Bob fall in love & marry
- Then they fall out of love & divorce
 - She gets the pets
 - He has to feed them
- Leading to a new coordination problem: Producer-Consumer



© 2007 Herlihy and Shavit

70


Bob Puts Food in the Pond



© 2007 Herlihy and Shavit

71

Alice releases her pets to Feed



© 2007 Herlihy and Shavit

72

Producer/Consumer

- Alice and Bob can't meet
 - Each has restraining order on other
 - So he puts food in the pond
 - And later, she releases the pets
- Avoid
 - Releasing pets when there's no food
 - Putting out food if uneaten food remains



© 2007 Herlihy and Shavit

73

Producer/Consumer


- Need a mechanism so that
 - Bob lets Alice know when food has been put out
 - Alice lets Bob know when to put out more food



© 2007 Herlihy and Shavit

74


Surprise Solution



© 2007 Herlihy and Shavit

75


Bob puts food in Pond



© 2007 Herlihy and Shavit

76

Bob knocks over Can



© 2007 Herlihy and Shavit


77

Alice Releases Pets



© 2007 Herlihy and Shavit

78



Pseudocode

```
while (true) {
  while (can.isUp()){};
  pet.release();
  pet.recapture();
  can.reset();
}
```

Alice's code

BROWN © 2007 Herlihy and Shavit 80

Pseudocode

```
while (true) {
  while (can.isUp()){};
  pet.release();
  pet.recapture();
  can.reset();
}
```

Alice's code

```
while (true) {
  while (can.isDown()){};
  pond.stockWithFood();
  can.knockover();
}
```


Bob's code

BROWN © 2007 Herlihy and Shavit 81

Correctness

- **Mutual Exclusion** — safety
Pets and Bob never together in pond
- **No Starvation** — liveness
If Bob always willing to feed and pets always famished, then pets eat infinitely often.
- **Producer/Consumer** — safety
Pets never enter pond unless there is food, and Bob never provides food if there is unconsumed food.

BROWN © 2007 Herlihy and Shavit 82



Waiting

- Both solutions use waiting
 - while (mumble) {}
- Waiting is **problematic**
 - If one participant is delayed, so is everyone else!
 - But delays are common & unpredictable

BROWN © 2007 Herlihy and Shavit 84

The Fable drags on ...

- Bob and Alice still have issues



© 2007 Herlihy and Shavit

85

The Fable drags on ...

- Bob and Alice still have issues
- So they need to communicate



© 2007 Herlihy and Shavit

86

The Fable drags on ...


- Bob and Alice still have issues
- So they need to communicate
- So they agree to use billboards ...



© 2007 Herlihy and Shavit

87

Billboards are Large



© 2007 Herlihy and Shavit

88


Write One Letter at a Time ...



© 2007 Herlihy and Shavit

89

To post a message



© 2007 Herlihy and Shavit

90

Let's send another message

© 2007 Herlihy and Shavit 91

Uh-Oh

© 2007 Herlihy and Shavit 92

Readers/Writers

- Devise a protocol so that
 - Writer writes one letter at a time
 - Reader reads one letter at a time
 - Reader sees
 - Old message or new message
 - No mixed messages

© 2007 Herlihy and Shavit 93

Readers/Writers (continued)

- Easy with mutual exclusion
- But mutual exclusion requires **waiting**
 - One **waits** for the other
 - Everyone executes **sequentially**
- Remarkably
 - We can solve R/W without mutual exclusion

© 2007 Herlihy and Shavit 94

Why do we care?

- We want as much of the code as possible to execute concurrently (in parallel)
- A larger sequential part implies reduced performance
- **Amdahl's law**: this relation is not linear...

© 2007 Herlihy and Shavit 95

Amdahl's Law

$$\text{Speedup} = \frac{\text{OldExecutionTime}}{\text{NewExecutionTime}}$$

...of computation given n CPUs instead of 1

© 2007 Herlihy and Shavit 96

Amdahl's Law

Sequential fraction \nearrow $1 - p + \frac{p}{n}$ \nwarrow Parallel fraction
Number of processors \nearrow n

Speedup = $\frac{1}{1 - p + \frac{p}{n}}$

BROWN © 2007 Herlihy and Shavit 97

Example

- Ten processors
- 60% concurrent, 40% sequential
- How close to 10-fold speedup?

BROWN © 2007 Herlihy and Shavit 98

Example

- Ten processors
- 60% concurrent, 40% sequential
- How close to 10-fold speedup?

Speedup = 2.17 = $\frac{1}{1 - 0.6 + \frac{0.6}{10}}$

BROWN © 2007 Herlihy and Shavit 99

Example

- Ten processors
- 80% concurrent, 20% sequential
- How close to 10-fold speedup?

BROWN © 2007 Herlihy and Shavit 100

Example

- Ten processors
- 80% concurrent, 20% sequential
- How close to 10-fold speedup?

Speedup = 3.57 = $\frac{1}{1 - 0.8 + \frac{0.8}{10}}$

BROWN © 2007 Herlihy and Shavit 101

Example

- Ten processors
- 90% concurrent, 10% sequential
- How close to 10-fold speedup?

BROWN © 2007 Herlihy and Shavit 102

Example

- Ten processors
- 90% concurrent, 10% sequential
- How close to 10-fold speedup?

$$\text{Speedup}=5.26 = \frac{1}{1 - 0.9 + \frac{0.9}{10}}$$



© 2007 Herlihy and Shavit

103

Example

- Ten processors
- 99% concurrent, 01% sequential
- How close to 10-fold speedup?



© 2007 Herlihy and Shavit

104

Example

- Ten processors
- 99% concurrent, 01% sequential
- How close to 10-fold speedup?

$$\text{Speedup}=9.17 = \frac{1}{1 - 0.99 + \frac{0.99}{10}}$$



© 2007 Herlihy and Shavit

105

The Moral

- The small % of a program that is hard to parallelize may have a large impact on overall speedup.



© 2007 Herlihy and Shavit

106