


## Universality of Consensus

  
BROWN  
Maurice Herlihy  
CS176  
Fall 2003

## Turing Computability




- A mathematical model of computation
- Complexity irrelevant to real machines
- Computable = Computable on a T-Machine



© 2003 Herlihy and Shavit

2

## Shared-Memory Computability?




- A model of concurrent computation
- Complexity irrelevant to real machines
- Wait-free/Lock-free computable = ????



© 2003 Herlihy and Shavit

3

## Asynchronous Computability



- A model of concurrent computation
- Complexity irrelevant to real machines
- Wait-free/Lock-free computable = ????



© 2003 Herlihy and Shavit

4

## Theorem: Universality

- Consensus is universal
- From n-thread consensus build a
  - Wait-free
  - Linearizable
  - n-threaded implementation
  - Of any sequentially specified object



© 2003 Herlihy and Shavit

5

## Proof Outline

- Construct a universal object
  - From n-consensus objects
  - And atomic registers
- Object implements any sequentially specified object canonically



© 2003 Herlihy and Shavit

6

## Like a Turing Machine

- This construction
  - Not intended to be practical
  - But enlightening
- Correctness, not efficiency
  - Why does it work? (Asks the scientist)
  - How does it work? (Asks the engineer)
  - Would you like fries with that? (Asks the liberal arts major)



© 2003 Herlihy and Shavit

7

## Universal Construction

- Object implemented as a
  - List of method calls
- Method call
  - Find end of list
  - Atomically append your call



© 2003 Herlihy and Shavit

8

## Naïve Idea


- Use consensus object to store pointer to cell with current state
- Each thread creates new cell
  - computes outcome,
  - and tries to switch pointer to its outcome
- Unfortunately not...
  - consensus objects can be used once only



© 2003 Herlihy and Shavit

9

## Basic Idea: Linked-List Representation



© 2003 Herlihy and Shavit

10

## Basic Idea

- Use one-time consensus object as next pointer
- Challenges
  - How to avoid starvation?
  - What if a thread stops in the middle?



© 2003 Herlihy and Shavit

11

## Stylized Sequential Objects

- Object itself is *immutable*
  - State never changes
- Method call produces
  - New copy of object in new state
  - Return value (or exception)



© 2003 Herlihy and Shavit

12

## Stylized Sequential Objects

- **Invocation**
  - Method name
  - Arguments
- **Response**
  - Copy of modified object
  - Return value
- **Assume methods are deterministic**
  - Only one response (can relax restriction)



© 2003 Herlihy and Shavit

13

## Invocation

```
public class Invoc {  
    public String method;  
    public Object[] args;  
}
```



© 2003 Herlihy and Shavit

14

## Invocation

```
public class Invoc {  
    public String method;  
    public Object[] args;  
}
```

Method name



© 2003 Herlihy and Shavit

15

## Invocation

```
public class Invoc {  
    public String method;  
    public Object[] args;  
}
```

Arguments



© 2003 Herlihy and Shavit

16

## Response

```
public class Response {  
    public SeqObject object;  
    public Object value;  
}
```



© 2003 Herlihy and Shavit

17

## Response

```
public class Response {  
    public SeqObject object;  
    public Object value;  
}
```

New object state



© 2003 Herlihy and Shavit

18

## Response

```
public class Response {
    public SeqObject object;
    public Object value;
}
```

Return value

BROWN © 2003 Herlihy and Shavit 19

## Sequential Objects

```
public class SeqObject {
    public Response apply(Invoc invoc) {
        ...
    }
}
```

BROWN © 2003 Herlihy and Shavit 20

## Sequential Objects

```
public class SeqObject {
    public Response apply(Invoc invoc) {
        ...
    }
}
```

Applies method invocation,  
returns new state and result

BROWN © 2003 Herlihy and Shavit 21

## Basic Data Structures

```
public class Cell {
    int seq;
    Invoc invoc;
    Response response;
    Consensus next;
}
```

BROWN © 2003 Herlihy and Shavit 22

## Basic Data Structures

```
public class Cell {
    int seq;
    Invoc invoc;
    Response response;
    SeqObject next;
}
```

Sequence number: zero  
while in play

BROWN © 2003 Herlihy and Shavit 23

## Basic Data Structures

```
public class Cell {
    int seq;
    Invoc invoc;
    Response response;
    Consensus next;
}
```

Method name & args

BROWN © 2003 Herlihy and Shavit 24

## Basic Data Structures

```
public class Cell {  
    int seq;  
    Invoc invoc;  
    Response response;  
    Consensus next;  
}
```

New object state & return value



© 2003 Herlihy and Shavit

25

## Basic Data Structures

```
public class Cell {  
    int seq;  
    Invoc invoc;  
    Response response;  
    Consensus next;  
}
```

Consensus object that  
indicates next cell in list



© 2003 Herlihy and Shavit

26

## Cell Constructor

```
public Cell(Invoc invoc) {  
    this.seq = 0;  
    this.invoc = invoc;  
    this.response = null;  
    this.next = new Consensus();  
}
```



© 2003 Herlihy and Shavit

27

## Cell Constructor

```
public Cell(Invoc invoc) {  
    this.seq = 0;  
    this.invoc = invoc;  
    this.response = null;  
    this.next = new Consensus();  
}
```

Sequence number zero means  
call still incomplete



© 2003 Herlihy and Shavit

28

## Cell Constructor

```
public Cell(Invoc invoc) {  
    this.seq = 0;  
    this.invoc = invoc;  
    this.response = null;  
    this.next = new Consensus();  
}
```

Record method name & args



© 2003 Herlihy and Shavit

29

## Cell Constructor

```
public Cell(Invoc invoc) {  
    this.seq = 0;  
    this.invoc = invoc;  
    this.response = null;  
    this.next = null;  
}
```

Winning thread decides



© 2003 Herlihy and Shavit

30

## Cells are Comparable

```
class Cell implements
    java.util.Comparable {
    int compareTo(Cell other) {
        if (this.seq > aCell.seq)
            return 1;
        else if (aCell.seq > this.seq)
            return -1;
        else
            return 0;
    }
}
```



© 2003 Herlihy and Shavit

31

## Cells are Comparable

```
class Cell implements
    java.util.Comparable {
    int compareTo(Cell other) {
        if (this.seq > aCell.seq)
            return 1;
        else if (aCell.seq > this.seq)
            return -1;
        else
            return 0;
    }
}
```

**Standard interface for class whose objects are totally ordered**



© 2003 Herlihy and Shavit

32

## Cells are Comparable

```
class Cell implements
    java.util.Comparable {
    int compareTo(Cell other) {
        if (this.seq > aCell.seq)
            return 1;
        else if (aCell.seq > this.seq)
            return -1;
        else
            return 0;
    }
}
```

**Returns +1 if I'm greater, -1 if I'm lesser, and 0 otherwise**



© 2003 Herlihy and Shavit

33

## Later is Greater

```
class Cell implements
    java.util.Comparable {
    int compareTo(Cell other) {
        if (this.seq > aCell.seq)
            return 1;
        else if (aCell.seq > this.seq)
            return -1;
        else
            return 0;
    }
}
```

**I'm greater if my sequence number is higher**



© 2003 Herlihy and Shavit

34

## Cells are Comparable

```
class Cell implements
    java.util.Comparable {
    int compareTo(Cell other) {
        if (this.seq > aCell.seq)
            return 1;
        else if (aCell.seq > this.seq)
            return -1;
        else
            return 0;
    }
}
```

**And vice-versa ...**



© 2003 Herlihy and Shavit

35

## Utility method

```
class Cell implements comparable {
    ...
    static Cell max(Cell[] array) {
        Cell max = array[0];
        for (int i=0; i<array.length; i++)
            if (max.compareTo(array[i]) < 0)
                max = array[i];
        return max;
    }
}
```



© 2003 Herlihy and Shavit

36

## Utility method

```
class Cell implements Comparable {
    ...
    static Cell max(Cell[] array) {
        Cell max = array[0];
        for (int i=0; i<array.length; i++)
            if (max.compareTo(array[i]) < 0)
                max = array[i];
        return max;
    }
}
```

**Find latest cell in array**



© 2003 Herlihy and Shavit

37

## Universal Object

```
public class Universal {
    private Cell[] announce;
    private Cell[] head;
    ...
}
```



© 2003 Herlihy and Shavit

38

## Universal Object

```
public class Universal {
    private Cell[] announce;
    private Cell[] head;
    ...
}
```

**If this thread does not succeed,  
another thread will help out**



© 2003 Herlihy and Shavit

39

## Universal Object

```
public class Universal {
    private Cell[] announce;
    private Cell[] head;
    ...
}
```

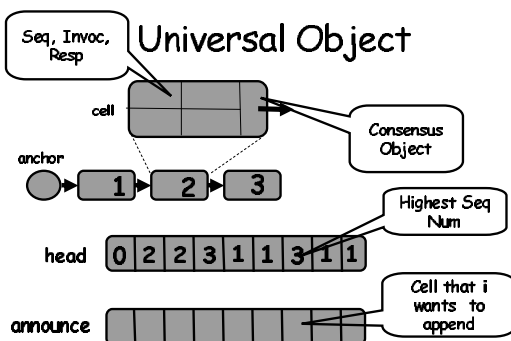
**Find the end of the list**



© 2003 Herlihy and Shavit

40

## Universal Object



© 2003 Herlihy and Shavit

41

## Thread A

- Allocates cell to represent call
- Stores (pointer to) cell in announce
  - If A doesn't execute it
  - Another thread will
- Looks for thread near end of list
  - By scanning head array
  - Choosing cell with largest sequence num



© 2003 Herlihy and Shavit

42

## Helping

- "Announcing" my intention
  - Guarantees progress
  - Even if the scheduler hates me
  - My method call will complete
- Makes protocol wait-free
- Otherwise starvation possible



© 2003 Herlihy and Shavit

43

## A Cry For Help

```
public class Universal {
    public Object apply(Invoc invoc) {
        int i = Thread.myIndex();
        this.announce[i] = new Cell(invoc);
        for (int j = 0; j < N; j++)
            this.head[i] =
                max(this.head[i], this.head[j]);
        ...
    }
}
```



© 2003 Herlihy and Shavit

44

## A Cry For Help

```
public class Universal {
    public Object apply(Invoc invoc) {
        int i = Thread.myIndex();
        this.announce[i] = new Cell(invoc);
        for (int j = 0; j < N; j++)
            this.head[i] =
                max(this.head[i], this.head[j]);
        ...
    }
}
```

Announce my intention to  
append cell to list



© 2003 Herlihy and Shavit

45

## Where does it End?

- Need to find end of list
  - Can't start from anchor
  - Starvation for slow threads
- Each thread records last cell seen
- If all collect from that array
  - Some thread has last cell



© 2003 Herlihy and Shavit

46

## Assume the Position

```
public class Universal {
    public Object apply(Invoc invoc) {
        int i = Thread.myIndex();
        this.announce[i] = new Cell(invoc);
        this.head[i] = max(this.head);
        ...
    }
}
```



© 2003 Herlihy and Shavit

47

## Assume the Position

```
public class Universal {
    public Object apply(Invoc invoc) {
        int i = Thread.myIndex();
        this.announce[i] = new Cell(invoc);
        this.head[i] = max(this.head);
        ...
    }
}
```

Look for end of list



© 2003 Herlihy and Shavit

48



## Are We Done Yet?

- Non-zero sequence number indicates success
- Thread keeps appending cells
- Until its own cell is done



© 2003 Herlihy and Shavit

49

## Keep on Keeping on

```
while (this.announce[i].seq == 0) {
  Cell before, help, prefer;
  before = this.head[i];
  help = announce[(before.seq+1) % n];
  if (help.seq == 0)
    prefer = help;
  else
    prefer = this.announce[i];
}
```



© 2003 Herlihy and Shavit

50

## Keep on Keeping on

```
while (this.announce[i].seq == 0) {
  Cell before, help, prefer;
  before = this.head[i];
  help = announce[(before.seq+1) % n];
  if (help.seq == 0)
    prefer = help;
  Keep trying until my cell gets a
  sequence number
}
```



© 2003 Herlihy and Shavit

51

## Keep on Keeping on

```
while (this.announce[i].seq == 0) {
  Cell before, help, prefer;
  before = this.head[i];
  help = announce[(before.seq+1) % n];
  if (help.seq == 0)
    prefer = help;
  else
    prefer = ... Possible end of list
}
```



© 2003 Herlihy and Shavit

52

## Altruism

- Choose a thread to "help"
- If that thread needs help
  - Try to append its cell
  - Otherwise append your own
- Worst case
  - Everyone tries to help same pitiful loser
  - Someone succeeds



© 2003 Herlihy and Shavit

53

## Help!

- Last cell in list has sequence number  $k$
- All threads check ...
  - Whether thread  $k+1 \bmod n$  wants help
  - If so, try to append her cell first



© 2003 Herlihy and Shavit

54

## Help!

- All threads try to help  $k+1 \pmod n$
- First time after  $k+1$  announces
  - Some may see announcement
  - Some may not
- "Many are cold but few are frozen"



© 2003 Herlihy and Shavit

55

## Help!

- First time after  $k+1$  announces
  - No guarantees
- After  $n$  more cells appended
  - Everyone sees that  $k+1$  wants help
  - Everyone tries to append that cell
  - Someone succeeds



© 2003 Herlihy and Shavit

56

## Lemma

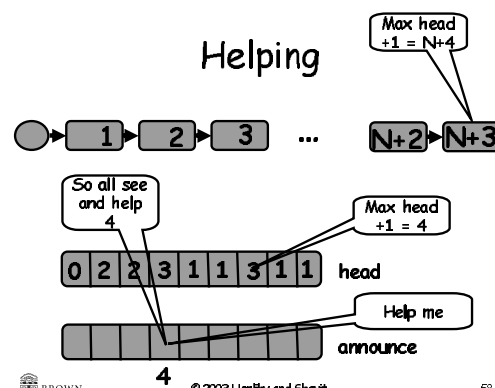
- After  $A$  announces cell
- No more than  $n$  other calls
  - Can start and finish
  - Without appending  $A$ 's cell



© 2003 Herlihy and Shavit

57

## Helping



© 2003 Herlihy and Shavit

58

## Pull Together

```
while (this.announce[i].seq == 0) {
    Cell before, help, prefer;
    before = this.head[i];
    help = announce[(before.seq+1) % n];
    if (help.seq == 0)
        prefer = help;
    else
        prefer = this.announce[i];
    ...
}
```



© 2003 Herlihy and Shavit

59

## Pull Together

```
while (this.announce[i].seq == 0) {
    Cell before, help, prefer;
    before = this.head[i];
    help = announce[(before.seq+1) % n];
    if (help.seq == 0)
        prefer = help;
    else
        prefer = this.announce[i];
    ...
    Pick another thread to help
}
```



© 2003 Herlihy and Shavit

60

## Pull Together

```
while (true) {
    Help if help required, but
    otherwise it's all about me!
    Cell before = this.head[i];
    help = announce((before.seq % n) + 1);
    if (help.seq == 0)
        prefer = help;
    else
        prefer = this.announce[i];
}
}
```



© 2003 Herlihy and Shavit

61

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    ...
    before.next.propose(prefer);
    Cell after = before.next.decide();
    SeqObject oldobject =
        before.response.object;
    Response response =
        oldobject.apply(after.invoc);
    after.nextState.propose(response);
    after.response =
        (Response)after.nextState.decide();
    ...
}
```



© 2003 Herlihy and Shavit

62

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    ...
    before.next.propose(prefer);
    Cell after = before.next.decide();
    SeqObject oldobject =
        before.response.object;
    Response response =
        oldobject.apply(after.invoc);
    Propose my favorite, then find out
    who actually won
    ...
}
```



© 2003 Herlihy and Shavit

63

## Finishing the Job

- Once we have linked in a cell
- Make sure remaining fields are filled in before moving on
- New object state in response field



© 2003 Herlihy and Shavit

64

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    ...
    before.next.propose(prefer);
    Cell after = before.next.decide();
    SeqObject oldobject =
        before.response.object;
    after.response =
        oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
}
```



© 2003 Herlihy and Shavit

65

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    ...
    before.next.propose(prefer);
    Cell after = before.next.decide();
    SeqObject oldobject =
        before.response.object;
    after.response =
        oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
}
```



© 2003 Herlihy and Shavit

66

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    Apply method and
    fill in response
    Cell after = before.next.decide();
    Seqobject oldobject =
    before.response.object;
    after.response =
    oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
```



© 2003 Herlihy and Shavit

67

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    Fill in sequence
    number
    Cell after = before.next.decide();
    Seqobject oldobject =
    before.response.object;
    after.response =
    oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
```



© 2003 Herlihy and Shavit

68

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    Note: Multiple threads
    could fill in these fields
    before.next.propose(prefer);
    Cell after = before.next.decide();
    Seqobject oldobject =
    before.response.object;
    after.response =
    oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
```



© 2003 Herlihy and Shavit

69

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    Note: Multiple threads
    could fill in these fields
    before.next.propose(prefer);
    Cell after = before.next.decide();
    Seqobject oldobject =
    before.response.object;
    after.response =
    oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
```



© 2003 Herlihy and Shavit

70

## Quality is Job 1.1

```
while (this.announce[i].seq == 0) {
    before.next.propose(prefer);
    Cell after = before.next.decide();
    Seqobject oldobject =
    before.response.object;
    after.response =
    oldobject.apply(after.invoc);
    after.seq = before.seq + 1;
    this.head[i] = after;
}
```



© 2003 Herlihy and Shavit

71

## Asynchronous Computability



Wait-free/Lock-free computable  
=  
Threads with methods that solve n-consensus



© 2003 Herlihy and Shavit

72

## GetAndSet is not Universal

```
public class RMWRegister {
    private int value;
    public boolean getAndSet(int update)
    {
        int prior = this.value;
        this.value = update;
        return prior;
    }
}
```



BROWN()

© 2003 Herlihy and Shavit

73

## GetAndSet is not Universal

```
public class RMWRegister {
    private int value;
    public boolean getAndSet(int update)
    {
        int prior = this.value;
        this.value = update;
        return prior;
    }
}
```

**Consensus number 2**



BROWN()

© 2003 Herlihy and Shavit

74

## GetAndSet is not Universal

```
public class RMWRegister {
    private int value;
    public boolean getAndSet(int update)
    {
        int prior = this.value;
        this.value = update;
        return prior;
    }
}
```

**Not universal for  $\geq 3$  threads**



BROWN()

© 2003 Herlihy and Shavit

75

## CompareAndSet is Universal

```
public class RMWRegister {
    private int value;
    public boolean
        compareAndSet(int expected,
                      int update) {
        int prior = this.value;
        if (this.value == expected) {
            this.value = update;
            return true;
        }
        return false;
    }
}
```



BROWN()

© 2003 Herlihy and Shavit

76

## CompareAndSet is Universal

```
public class RMWRegister {
    private int value;
    public boolean
        compareAndSet(int expected,
                      int update) {
        int prior = this.value;
        if (this.value == expected) {
            this.value = update;
            return true;
        }
        return false;
    }
}
```

**Consensus number  $\infty$**



BROWN()

© 2003 Herlihy and Shavit

77

## CompareAndSet is Universal

```
public class RMWRegister {
    private int value;
    public boolean
        compareAndSet(int expected,
                      int update) {
        int prior = this.value;
        if (this.value == expected) {
            this.value = update;
            return true;
        }
        return false;
    }
}
```

**Universal for any number of threads**



BROWN()

© 2003 Herlihy and Shavit

78

## Practical Implications

- Any architecture that does not provide a universal primitive has inherent limitations
- You cannot avoid locking for concurrent data structures ...



© 2003 Herlihy and Shavit

79

## Older Architectures

- IBM 360
  - testAndSet (getAndSet)
- NYU UltraComputer
  - getAndAdd
- Neither universal
  - Except for 2 threads



© 2003 Herlihy and Shavit

80

## Newer Architectures

- Intel x86, Itanium, SPARC
  - compareAndSet
- Alpha AXP, PowerPC
  - Load-locked/store-conditional
- All universal
  - For any number of threads
- Trend is clear ...



© 2003 Herlihy and Shavit

81

## Till Now: Correctness

- Models
  - Accurate (we never lied to you)
  - But idealized (so we forgot to mention a few things)
- Protocols
  - Elegant
  - Important
  - But naïve



© 2003 Herlihy and Shavit

82

## Next Time: Performance

- Models
  - More complicated (not the same as complex)
  - Still focus on principles (not soon obsolete)
- Protocols
  - Elegant (in their fashion)
  - Important (why else would we pay attention?)
  - And realistic (your mileage may vary)



© 2003 Herlihy and Shavit

83

## Comments on Universality

Maurice: I removed non-determinism from the Code but there are still redundant fields in each cell. Please fix code. Also, from comments and questions we need to think of a better way to define The max function of the seq numbers in the nodes since It confused several students who came up with counter examples simply because they missed the fact That the max is on the seq numbers...

