

Multithreaded Programming in Cilk

LECTURE 3

Charles E. Leiserson

Supercomputing Technologies Research Group
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Minicourse Outline

- **LECTURE 1**
Basic Cilk programming: Cilk keywords, performance measures, scheduling.
- **LECTURE 2**
Analysis of Cilk algorithms: matrix multiplication, sorting, tableau construction.
- **LABORATORY**
Programming matrix multiplication in Cilk — Dr. Bradley C. Kuszmaul
- **LECTURE 3**
Advanced Cilk programming: inlets, abort, speculation, data synchronization, & more.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 2

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 3

Operating on Returned Values

Programmers may sometimes wish to incorporate a value returned from a spawned child into the parent frame by means other than a simple variable assignment.

Example: `x += spawn foo(a,b,c);`

Cilk achieves this functionality using an internal function, called an *inlet*, which is executed as a secondary thread on the parent frame when the child returns.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 4

Semantics of Inlets

```
int max, ix = -1;
inlet void update ( int val, int index ) {
  if (idx == -1 || val > max) {
    ix = index; max = val;
  }
}
...
for (i=0; i<1000000; i++) {
  update ( spawn foo(i), i );
}
sync; /* ix now indexes the largest foo(i) */
```

- The **inlet** keyword defines a **void** internal function to be an inlet.
- In the current implementation of Cilk, the inlet definition may not contain a **spawn**, and only the first argument of the inlet may be spawned at the call site.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 5

Semantics of Inlets

```
int max, ix = -1;
inlet void update ( int val, int index ) {
  if (idx == -1 || val > max) {
    ix = index; max = val;
  }
}
...
for (i=0; i<1000000; i++) {
  update ( spawn foo(i), i );
}
sync; /* ix now indexes the largest foo(i) */
```

1. The non-**spawn** args to **update ()** are evaluated.
2. The Cilk procedure **foo (i)** is spawned.
3. Control passes to the next statement.
4. When **foo (i)** returns, **update ()** is invoked.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 6

Semantics of Inlets

```
int max, ix = -1;
inlet void update ( int val, int index ) {
    if (idx == -1 || val > max) {
        ix = index; max = val;
    }
}
...
for (i=0; i<1000000; i++) {
    update ( spawn foo(i), i );
}
sync; /* ix now indexes the largest foo(i) */
```

Cilk provides *implicit atomicity* among the threads belonging to the same frame, and thus no locking is necessary to avoid data races.

Implicit Inlets

```
cilk int wfib(int n) {
    if (n == 0) {
        return 0;
    } else {
        int i, x = 1;
        for (i=0; i<=n-2; i++) {
            x += spawn wfib(i);
        }
        sync;
        return x;
    }
}
```

For assignment operators, the Cilk compiler automatically generates an *implicit inlet* to perform the update.

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

Computing a Product

$$p = \prod_{i=0}^n A_i$$

```
int product(int *A, int n) {
    int i, p=1;
    for (i=0; i<n; i++) {
        p *= A[i];
    }
    return p;
}
```

Optimization: Quit early if the partial product ever becomes 0.

Computing a Product

$$p = \prod_{i=0}^n A_i$$

```
int product(int *A, int n) {
    int i, p=1;
    for (i=0; i<n; i++) {
        p *= A[i];
        if (p == 0) break;
    }
    return p;
}
```

Optimization: Quit early if the partial product ever becomes 0.

Computing a Product in Parallel

$$p = \prod_{i=0}^n A_i$$

```
cilk int prod(int *A, int n) {
    int p = 1;
    if (n == 1) {
        return A[0];
    } else {
        p *= spawn product(A, n/2);
        p *= spawn product(A+n/2, n-n/2);
        sync;
        return p;
    }
}
```

How do we quit early if we discover a zero?

Cilk's Abort Feature

```
cilk int product(int *A, int n) {
  int p = 1;
  inlet void mult(int x) {
    p *= x;
    return;
  }

  if (n == 1) {
    return A[0];
  } else {
    mult( spawn product(A, n/2) );
    mult( spawn product(A+n/2, n-n/2) );
    sync;
    return p;
  }
}
```

1. Recode the implicit inlet to make it explicit.

Cilk's Abort Feature

```
cilk int product(int *A, int n) {
  int p = 1;
  inlet void mult(int x) {
    p *= x;

    return;
  }

  if (n == 1) {
    return A[0];
  } else {
    mult( spawn product(A, n/2) );
    mult( spawn product(A+n/2, n-n/2) );
    sync;
    return p;
  }
}
```

2. Check for 0 within the inlet.

Cilk's Abort Feature

```
cilk int product(int *A, int n) {
  int p = 1;
  inlet void mult(int x) {
    p *= x;
    if (p == 0) {
      abort; /* Aborts existing children, */
            /* but not future ones. */
    }
    return;
  }

  if (n == 1) {
    return A[0];
  } else {
    mult( spawn product(A, n/2) );
    mult( spawn product(A+n/2, n-n/2) );
    sync;
    return p;
  }
}
```

2. Check for 0 within the inlet.

Cilk's Abort Feature

```
cilk int product(int *A, int n) {
  int p = 1;
  inlet void mult(int x) {
    p *= x;
    if (p == 0) {
      abort; /* Aborts existing children, */
            /* but not future ones. */
    }
    return;
  }

  if (n == 1) {
    return A[0];
  } else {
    mult( spawn product(A, n/2) );

    mult( spawn product(A+n/2, n-n/2) );
    sync;
    return p;
  }
}
```

Cilk's Abort Feature

```
cilk int product(int *A, int n) {
  int p = 1;
  inlet void mult(int x) {
    p *= x;
    if (p == 0) {
      abort; /* Aborts existing children, */
            /* but not future ones. */
    }
    return;
  }

  if (n == 1) {
    return A[0];
  } else {
    mult( spawn product(A, n/2) );
    if (p == 0) { /* Don't spawn if we've */
      return 0; /* already aborted! */
    }
    mult( spawn product(A+n/2, n-n/2) );
    sync;
    return p;
  }
}
```

Implicit atomicity eases reasoning about races.

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

Min-Max Search

- Two players: MAX ■ and MIN ●.
- The game tree represents all moves from the current position within a given search depth.
- At leaves, apply a static evaluation function.
- MAX chooses the maximum score among its children.
- MIN chooses the minimum score among its children.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 19

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 20

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 21

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 22

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 23

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 24

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 25

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 26

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 27

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 28

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 29

Alpha-Beta Pruning

IDEA: If MAX ■ discovers a move so good that MIN ● would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 30

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 31

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 32

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 33

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 34

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 35

Alpha-Beta Pruning

IDEA: If MAX \blacksquare discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in C++ — LECTURE 3 July 17, 2006 36

Alpha-Beta Pruning

IDEA: If MAX \square discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 37

Alpha-Beta Pruning

IDEA: If MAX \square discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 38

Alpha-Beta Pruning

IDEA: If MAX \square discovers a move so good that MIN \bullet would never allow that position, MAX's other children need not be searched — *beta cutoff*.

Unfortunately, this heuristic is inherently serial.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 39

Parallel Min-Max Search

OBSERVATION: In a best-ordered tree, the degree of every internal node is either **1** or *maximal*.

IDEA: [Feldman-Mysliwicz-Monien 91] If the first child fails to generate a cutoff, *speculate* that the remaining children can be searched in parallel without wasting any work: “*young brothers wait.*”

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 40

Parallel Alpha-Beta (I)

```

cilk int search(position *prev, int move, int depth) {
    position cur;           /* Current position */
    int bestscore = -INF;   /* Best score so far */
    int num_moves;         /* Number of children */
    int mv;                /* Index of child */
    int sc;                 /* Child's score */
    int cutoff = FALSE;    /* Have we seen a cutoff? */

```

- View from MAX's perspective; MIN's viewpoint can be obtained by negating scores — *negamax*.
- The node generates its children from its parent's position `prev` *#Cilk keywords used so far*
- The `alpha` and `beta` limits and the move list are fields of the `position` data structure. 1

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 41

Parallel Alpha-Beta (II)

```

inlet void get_score(int child_sc) {
    child_sc = -child_sc; /* Negamax */

    if (child_sc > bestscore) {
        bestscore = child_sc;
        if (child_sc > cur.alpha) {
            cur.alpha = child_sc;
            if (child_sc >= cur.beta) { /* Beta cutoff */
                cutoff = TRUE; /* No need to search more */
                abort; /* Terminate other children */
            }
        }
    }
}

```

3

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 42

Parallel Alpha-Beta (III)

```
/* Create current position and set up for search */
make_move(prev, move, &cur);

sc = eval(&cur); /* Static evaluation */
if ( abs(sc)>=MATE || depth<=0 ) { /* Leaf node */
    return (sc);
}

cur.alpha = -prev->beta; /* Negamax */
cur.beta = -prev->alpha;

/* Generate moves, hopefully in best-first order*/
num_moves = gen_moves(&cur);
```

3

© 2006 by Charles E. Leiserson

Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 43

Parallel Alpha-Beta (IV)

```
/* Search the moves */
for (mv=0; !cutoff && mv<num_moves; mv++) {
    get_score( spawn search(&cur, mv, depth-1) );
    if (mv==0) sync; /* Young brothers wait */
}
sync;
return (bestscore);
}
```

- Only 6 Cilk keywords need be embedded in the C program to parallelize it.
- In fact, the program can be parallelized using only 5 keywords at the expense of minimal obfuscation.

6

© 2006 by Charles E. Leiserson

Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 44

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

© 2006 by Charles E. Leiserson

Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 45

Mutual Exclusion

Cilk's solution to mutual exclusion is no better than anybody else's.

Cilk provides a library of spin locks declared with `Cilk_lockvar`.

- To avoid deadlock with the Cilk scheduler, a lock should only be held within a Cilk thread.
- *I.e.*, `spawn` and `sync` should not be executed while a lock is held.

Fortunately, Cilk's control parallelism often mitigates the need for extensive locking.

© 2006 by Charles E. Leiserson

Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 46

Cilk's Memory Model

Programmers may also synchronize through memory using lock-free protocols, although Cilk is agnostic on consistency model.

- If a program contains no data races, Cilk effectively supports sequential consistency.
- If a program contains data races, Cilk's behavior depends on the consistency model of the underlying hardware.

To aid portability, the `Cilk_fence()` function implements a memory barrier on machines with weak memory models.


© 2006 by Charles E. Leiserson

Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 47

Debugging Data Races

Cilk's *Nondeterminator* debugging tool provably guarantees to detect and localize data-race bugs.



A *data race* occurs whenever two logically parallel threads, holding no locks in common, access the same location and one of the threads modifies the location.

© 2006 by Charles E. Leiserson

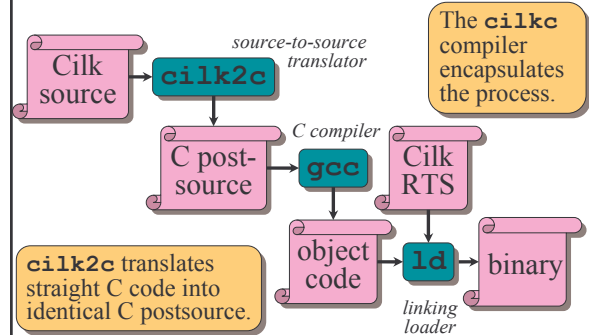
Multithreaded Programming in Cilk — LECTURE 3

July 17, 2006 48

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

Compiling Cilk




Cilk's Compiler Strategy

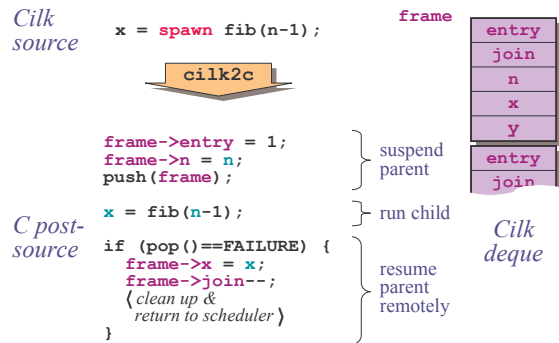
The `cilk2c` translator generates two “clones” of each Cilk procedure:

- **fast clone**—serial, common-case code.
- **slow clone**—code with parallel bookkeeping.


- The **fast clone** is always spawned, saving live variables on Cilk's work deque (shadow stack).
- The **slow clone** is resumed if a thread is stolen, restoring variables from the shadow stack.
- A check is made whenever a procedure returns to see if the resuming parent has been stolen.



Compiling `spawn` — Fast Clone

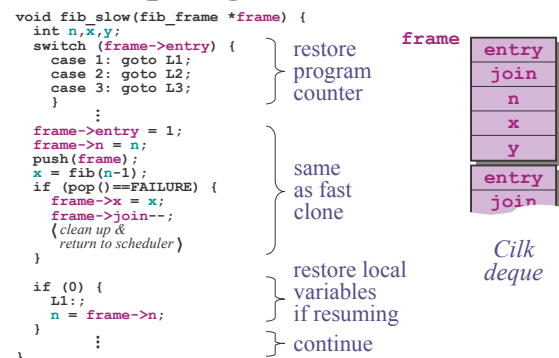


Compiling `sync` — Fast Clone




No synchronization overhead in the fast clone!

Compiling the Slow Clone



Breakdown of Work Overhead (circa 1997)




Benchmark: fib on one processor.

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

The JCilk System



- Jgo = Java + goto.
- The Jgo compiler was built by modifying gcj to accept goto statements so that a continuation mechanism for JCilk could be implemented.

JCilk Keywords

cilk
spawn
sync
SYNCHED

Same as Cilk, except that **cilk** can also modify **try**.

~~**inlet**~~
~~**abort**~~

Eliminated!

JCilk leverages Java's exception mechanism to render two Cilk keywords unnecessary.

Exception Handling in Java

“During the process of throwing an exception, the Java virtual machine **abruptly completes**, one by one, any expressions, statements, method and constructor invocations, initializers, and field initialization expressions that have begun but not completed execution in the current thread. This process continues until a handler is found that indicates that it handles that particular exception by naming the class of the exception or a superclass of the class of the exception.”

— J. Gosling, B Joy, G. Steele, and G. Bracha, *Java Language Specification*, 2000, pp. 219–220.

Exception Handling in JCilk

```

private cilk void foo() throws IOException {
    spawn A();
    cilk try {
        spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
    spawn D();
    doYetSomethingElse();
    sync;
}
    
```

Exception Handling in JCilk

```
private cilk void foo() throws IOException {
    spawn A();
    cilk try {
        spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
    spawn D();
    doYetSomethingElse();
    sync;
}
```

Exception!

An exception causes all subcomputations dynamically enclosed by the catching clause to abort!

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 61

Exception Handling in JCilk

```
private cilk void foo() throws IOException {
    spawn A();
    cilk try {
        spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
    spawn D();
    doYetSomethingElse();
    sync;
}
```

ArithmeticEx'n

Nothing aborts.

An exception causes all subcomputations dynamically enclosed by the catching clause to abort!

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 62

Exception Handling in JCilk

```
private cilk void foo() throws IOException {
    spawn A();
    cilk try {
spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
    spawn D();
    doYetSomethingElse();
    sync;
}
```

RuntimeEx'n

An exception causes all subcomputations dynamically enclosed by the catching clause to abort!

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 63

Exception Handling in JCilk

```
private cilk void foo() throws IOException {
spawn A();
    cilk try {
spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
spawn D();
    doYetSomethingElse();
    sync;
}
```

IOException

An exception causes all subcomputations dynamically enclosed by the catching clause to abort!

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 64

Exception Handling in JCilk

```
private cilk void foo() throws IOException {
    spawn A();
    cilk try {
spawn B();
        cilk try {
            spawn C();
        } catch(ArithmeticEx'n e) {
            doSomething();
        }
    } catch(RuntimeException e) {
        doSomethingElse();
    }
    spawn D();
    doYetSomethingElse();
    sync;
}
```

RuntimeEx'n

The appropriate catch clause is executed only after all spawned methods within the corresponding try block terminate.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 65

JCilk's Exception Mechanism

- JCilk's exception semantics allow programs such as alpha-beta to be coded without Cilk's **inlet** and **abort** keywords.
- Unfortunately, Java exceptions are slow, reducing the utility of JCilk's faithful extension.

© 2006 by Charles E. Leiserson Multithreaded Programming in Cilk — LECTURE 3 July 17, 2006 66

LECTURE 3

- Inlets
- Abort
- Speculative Computing
- Data Synchronization
- Under the Covers
- JCilk
- Conclusion

Future Work

Adaptive computing

- Get rid of `--nproc`.
- Build a job scheduler that uses *parallelism feedback* to balance processor resources among Cilk jobs.

Integrating Cilk with static threads

- Currently, interfacing a Cilk program to other system processes requires arcane knowledge.
- Build linguistic support into Cilk for Cilk processes that communicate.
- Develop a job scheduler that uses *pipeloid* to allocate resources among Cilk processes.

Key Ideas

- Cilk is simple: **cilk**, **spawn**, **sync**, **SYNCHED**, **inlet**, **abort**

- JCilk is simpler

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span

- Work & span


- Work & span

- Work & span

- Work & span

- Work & span

- Work & span



Open-Cilk Consortium

- We are in the process of forming a consortium to manage, organize, and promote Cilk open-source technology.
- If you are interested in participating, please let us know.

ACM Symposium on Parallelism in Algorithms and Architectures

SPAA 2006



Cambridge, MA, USA
July 30 – August 2, 2006