

Lecture 04: Productivity in Parallel Programming

Vivek Kumar

Computer Science and Engineering

IIIT Delhi

vivekk@iiitd.ac.in



Last Lecture

- Concurrency v/s parallelism
 - “Dealing” v/s “Doing”
- Mapping tasks to threads
 - Static and dynamic mapping
- Parallel performance
 - Critical path
 - Amdahl’s law
 - Speedup
- Task decomposition techniques
 - Data (SIMD)
 - Recursive
 - Exploratory
 - Speculative

Today's Lecture

- ➔ ● Issues with explicit multithreading
- Tasks based parallel programming model

Fibonacci using Pthread

```
#include <inttypes.h>
#include <pthread.h>
#include <stdio.h>
#include <stdlib.h>

uint64_t fib(uint64_t n) {
    if (n < 2) {
        return n;
    } else {
        uint64_t x = fib(n-1);
        uint64_t y = fib(n-2);
        return (x + y);
    }
}

typedef struct {
    uint64_t input;
    uint64_t output;
} thread_args;

void *thread_func(void *ptr) {
    uint64_t i =
        ((thread_args *) ptr)->input;
    ((thread_args *) ptr)->output = fib(i);
    return NULL;
}
```

**What are the issues
in this program?**

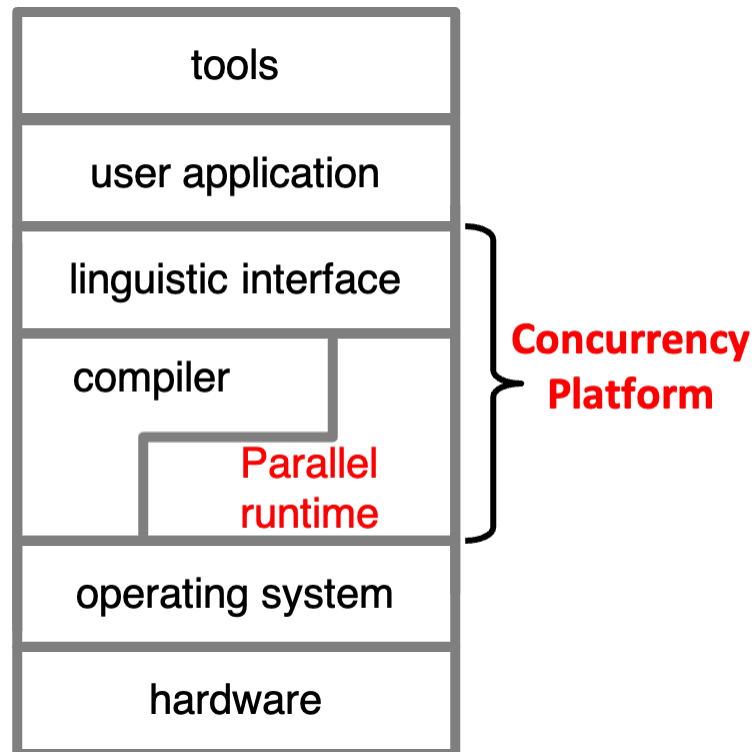
```
int main(int argc, char *argv[]) {
    pthread_t thread;
    thread_args args;
    int status;
    uint64_t result;

    if (argc < 2) { return 1; }
    uint64_t n = strtoul(argv[1], NULL, 0);
    if (n < 30) {
        result = fib(n);
    } else {
        args.input = n-1;
        status = pthread_create(&thread,
                                NULL,
                                thread_func,
                                (void*) &args);
        // main can continue executing
        if (status == NULL) { return 1; }
        result = fib(n-2);
        // wait for the thread to terminate.
        status = pthread_join(thread, NULL);
        if (status != NULL) { return 1; }
        result += args.output;
    }
    printf("Fibonacci of %" PRIu64 " is %" PRIu64 ".\n",
          n, result);
    return 0;
}
```

Issues with Pthreads

Overhead	The cost of creating a thread $>10^4$ cycles \Rightarrow coarse-grained concurrency. (Thread pools can help.)
Scalability	Fibonacci code gets at most about 1.5 speedup for 2 cores. Need a rewrite for more cores. Although, array sum is scalable with increasing number of codes (although limited by DRAM access latency)
Modularity	The program logic is no longer neatly encapsulated in the original sequential function.
Code Simplicity	Programmers must marshal arguments (shades of 1958!) and engage in error-prone protocols in order to load-balance.

Parallel Runtimes



- A parallel runtime should provide:
 - an interface for specifying the **logical parallelism** of the computation;
 - a runtime layer to automate scheduling and synchronization; and
 - guarantees of performance and resource utilization competitive with hand-tuned code

Fibonacci using C++11 `std::thread`

```

1. uint64_t fib(uint64_t n) {
2.     if (n < 2) {
3.         return n;
4.     } else {
5.         uint64_t x = fib(n-1);
6.         uint64_t y = fib(n-2);
7.         return (x + y);
8.     }
9. }
10. int main(int argc, char *argv[]) {
11.     uint64_t result = fib(40);
12.     printf("Result is %" PRIu64 ".\n", result);
13. }

```

std::thread



```

uint64_t fib(uint64_t n) {
    if (n < 2) {
        return n;
    } else {
        uint64_t x, y;
        std::thread t1([&]() {x = fib(n-1);});
        std::thread t2([&]() {y = fib(n-2);});
        t1.join(); t2.join();
        return (x + y);
    }
}

```

- The modularity and code simplicity issues are resolved using C++11 threads
 - Will the above program run for fib(40)?

Today's Lecture

- Issues with explicit multithreading
- ➔ ● Tasks based parallel programming model

Types of Tasks

- Synchronous (finish)
 - Blocks until the task execution is complete
- Asynchronous (async)
 - Doesn't blocks for the task to complete its execution

Your Sunday Tasks

Post on Facebook that you are done with all your tasks!

Wash your clothes in washing machine

Watch movies on laptop

Talk to father

Buy fruits online using your smartphone

Talk to mother

Make your bed

Complete your MPPRS project deadline

Your Sunday Tasks (Stmt. Reordering)

Complete your MPPRS project deadline

Wash your clothes in washing machine

Watch movies on laptop

Talk to father

Talk to mother

Buy fruits online using your smartphone

Make your bed

Post on Facebook that you are done with all your tasks!

Async-Finish Sunday Tasks

```

finish {
  async { Complete your MPPRS project deadline }
}
finish {
  async { Watch movies on laptop }
  async { Talk to father
          Talk to mother }
  async { Buy fruits online using your smartphone }
}
  
```

Post on Facebook that you are done with all your tasks!

Async-Finish Sunday Tasks

```

finish {
  async { Complete your MPPRS project deadline }
  async { Wash your clothes in washing machine }
}
finish {
  async { Watch movies on laptop }
  async { Talk to father
          Talk to mother }
  async { Buy fruits online using your smartphone }
  async { Make your bed }
}

```

- Statement “async S1; S2;” implies S1 and S2 could run asynchronously
- Hence, there is no need to specify “async” on S2

Post on Facebook that you are done with all your tasks!

Async-Finish Sunday Tasks

```

finish {
  async { Wash your clothes in washing machine }
}
}
finish {
  async { Complete your MPPRS project deadline }
  async { Watch movies on laptop }
  async { Talk to father
          Talk to mother }
  async { Buy fruits online using your smartphone }
  async { Make your bed }
}

```

- We applied statement reordering here
 - Alas.. MPPRS project deadline will take more time than washing machine...
- Hence, the second finish could be removed as we cannot launch async tasks until we are done with MPPRS project implementation

Post on Facebook that you are done with all your tasks!

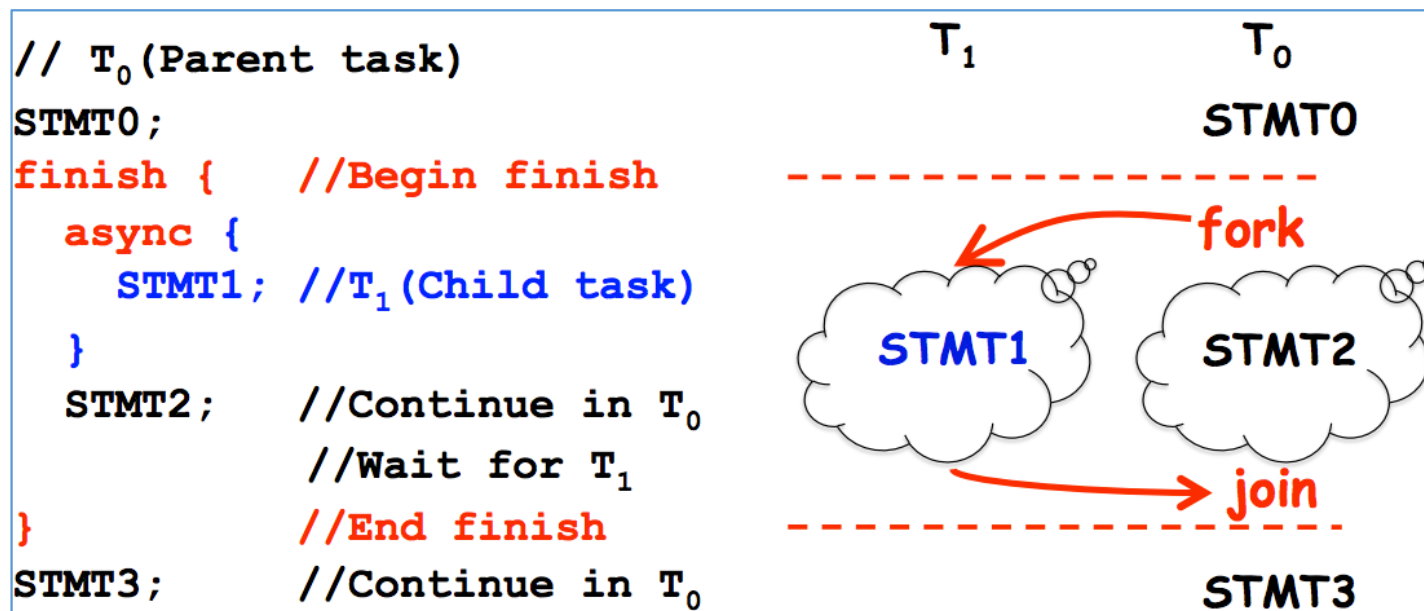
Async and Finish Statements for Task Creation and Termination (Pseudocode)

async S

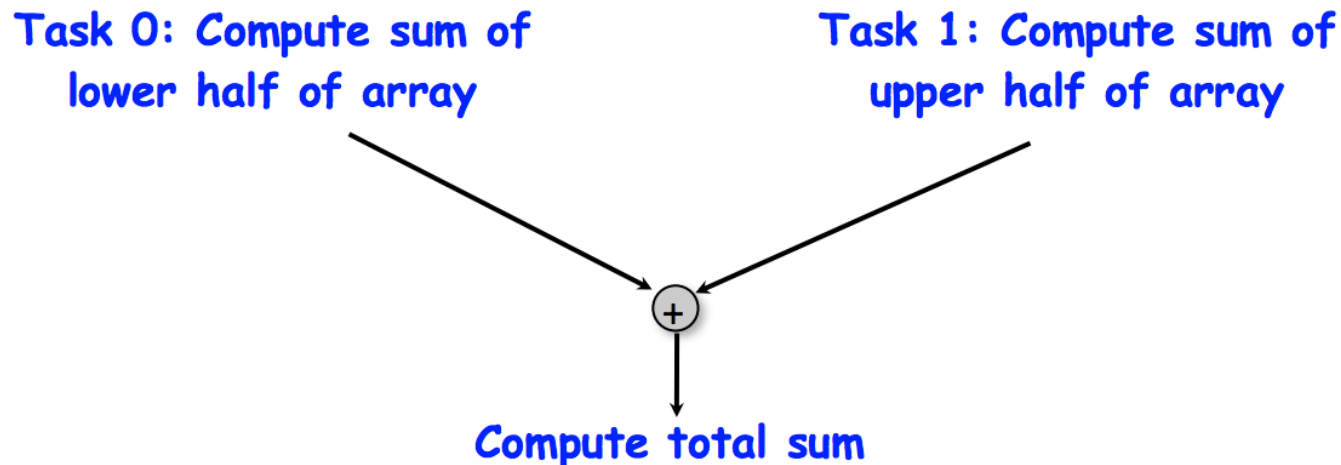
- Creates a new child task that executes statement S

finish S

- Execute S but wait until all async in S's scope have terminated



2-Way Parallel Array Sum using async-finish



- Basic idea

- Decompose the problem into tasks for partial sums
- Combine results to obtain final answer
- Parallel divide-n-conquer pattern

2-Way Parallel Array Sum using `async-finish`

Algorithm 2: Two-way Parallel ArraySum

Input: Array of numbers, X .

Output: $sum = \text{sum of elements in array } X$.

// Start of Task T1 (main program)

$sum1 \leftarrow 0; sum2 \leftarrow 0;$

// Compute $sum1$ (lower half) and $sum2$ (upper half) in parallel.

`finish{`

`async{`

 // Task T2

 for $i \leftarrow 0$ to $X.length/2 - 1$ do

$sum1 \leftarrow sum1 + X[i];$

 };

`async{`

 // Task T3

 for $i \leftarrow X.length/2$ to $X.length - 1$ do

$sum2 \leftarrow sum2 + X[i];$

 };

`};`

// Task T1 waits for Tasks T2 and T3 to complete

// Continuation of Task T1

$sum \leftarrow sum1 + sum2;$

return $sum;$

Race condition if you
miss the **finish**

N-Way Parallel Array Sum

```
int A[SIZE]; // Initialization code elided
int array_sum(int low, int high) {
    int sum = 0;
    for (int i=low; i<high; i++) {
        sum += A[i];
    }
    return sum;
}
```

N-Way Parallel Array Sum using async-finish

```
int A[SIZE]; // Initialization code elided
int array_sum(int low, int high) {
    int sum = 0;
    for (int i=low; i<high; i++) {
        sum += A[i];
    }
    return sum;
}
```

```
int main(int argc, char *argv[]) {
    int result;
    if (SIZE < 1024) {
        result = array_sum(0, SIZE);
    } else {
        int chunk = SIZE/NTHREADS;
        finish([&result,=] () {
            for (int i=0; i<NTHREADS; i++) {
                async([&result,=] () {
                    int low = i*chunk, high = (i+1)*chunk;
                    result += array_sum(low, high);
                });
            }
        });
    }
    printf("Total Sum is %d\n", result);
    return 0;
}
```

N-Way Parallel Array Sum using async-finish

```
int A[SIZE]; // Initialization code elided
int array_sum(int low, int high) {
    int sum = 0;
    for (int i=low; i<high; i++) {
        sum += A[i];
    }
    return sum;
}
```

Do you see any issues in this version of the parallel array sum?

```
int main(int argc, char *argv[]) {
    int result;
    if (SIZE < 1024) {
        result = array_sum(0, SIZE);
    } else {
        int chunk = SIZE/NTHREADS;
        finish([&result,=]() {
            for (int i=0; i<NTHREADS; i++) {
                async([&result,=]() {
                    int low = i*chunk, high = (i+1)*chunk;
                    result += array_sum(low, high);
                });
            }
        });
        printf("Total Sum is %d\n", result);
        return 0;
    }
}
```

C++11 lambda functions

Race condition !!!

N-Way Parallel Array Sum using async-finish

```
int A[SIZE]; // Initialization code elided
int array_sum(int low, int high) {
    int sum = 0;
    for (int i=low; i<high; i++) {
        sum += A[i];
    }
    return sum;
}
```

Race condition is fixed using an array to store result from each thread


```
int main(int argc, char *argv[]) {
    int result[NTHREADS], sum;
    if (SIZE < 1024) {
        sum = array_sum(0, SIZE);
    } else {
        int chunk = SIZE/NTHREADS;
        finish([&result,=]() {
            for (int i=0; i<NTHREADS; i++) {
                async([&result,=]() {
                    int low = i*chunk, high = (i+1)*chunk;
                    result[i] = array_sum(low, high);
                });
            }
        });
        for (int i=0; i<NTHREADS; i++) {
            sum += result[i];
        }
    }
    printf("Total Sum is %d\n", sum);
    return 0;
}
```

Recursive Fibonacci using async-finish

```

uint64_t fib(uint64_t n) {
  if (n < 2) {
    return n;
  } else {
    uint64_t x, y;
    finish([&]() {
      async([&]() { x = fib(n-1); });
      y = fib(n-2);
    });
    return (x + y);
  }
}

```



Next Lecture

- Greedy scheduler
- Quiz-1
 - Syllabus: Lectures 1-4