Lectures

TDDA69 Data and Program Structure Declarative Concurrency Cyrille Berger



14Macro

1 Introduction

2Concepts and models of programming languages

5Declarative Computation Implementation

10Imperative Programming Implementation

8 Explicit State and Imperative Model 9 Imperative Programming Techniques

3Declarative Computation Model 4Declarative Programming Techniques

6 **Declarative Concurrency** 7 Message Passing Concurrency

11Shared-State Concurrency 12Relational Programming 13Constraint Programming

15Running natively and JIT 16Garbage Collection 17Summary

2/53

Lecture content

- Data-driven concurrent model
- Lazy execution
- MapReduce
- Stream

Concurrency

- Natural for some applications:
 - ^o GUI: event thread, computation threads...
 - ^o Games: graphics thread, ai thread...
 - ^o Operating Systems: driver threads...
 - ^o Data processing...
- Distributed applications: clientserver

Declarative Concurrency

No observable nondeterminism
 ^o The result is only dependent of the inputs
 Communication through shared dataflow variables



I.U LENSER 5 / 53

Sequential Model

Data-driven concurrent model



Concurrent Model



1.U HARVENN 9 / 53

Basic concepts

- The model allows multiple statements to execute "at the same time"
- Reading and writing different variables can be done simultaneously by different threads, as well as reading the same variable
- What about writting to the same variable?

Extension to the syntax of KL

• STATEMENT := (...|THREAD) • THREAD := 'thread' '('(EXPR(,EXPR)+)) | () ')'('->' IDENTIFIER) BLOCK • Example: define my_thread = thread(a) -> b { b = fib(a); }; define v1 = my_thread(1000); define v2 = my_thread(2000); std.print(v1 + v2);

10/53

Deterministic vs Nondeterministic

- An execution is nondeterministic if there is a computation step in which there is a choice what to do next
- Nondeterminism appears naturally when there are multiple concurrent states
- A declarative language needs to be deterministic



Example of non determinism

define x; thread() {
 x = 1;
}();
thread() {
 x = 3;
}();
The first thread th

- The first thread that bind x succeed, the other one fails
- Commonly called "data race".

13 / 53

15/53

How to solve?

```
• Use exception?
define x;
thread() {
    catch(x = 1);
}();
thread() {
    catch(x = 3);
}();
• Not deterministic!
```

- Two solutions:
- $^\circ$ Leave it as an error to solve by the programmer
- $^{\circ}$ An alternative model is to assign a value store for each thread

14/53

Multiple value stores



Thread synchronisation

Implicit synchronisation

Wait for the computation to be available
define my_thread = thread(a) -> b {
 b = fib(a);
};
define v1 = my_thread(1000);
define v2 = my_thread(2000);
std.print(v1 + v2);
The main thread is blocked until v1 and v2 are finished.

• v1 and v2 are called *dataflow variable*



Dead-lock



Sequential / Concurrent Execution





17 / 53

Since concurrency is easy...

• Declarative programs can be easily made concurrent • There is no real drawbacks (except cost of starting a thread) •define fib = function(n) -> v { $v = cond(n \le 2, 1,$ fib(x-1) + fib(x-2));•define fib = function(n) -> v { $v = cond(n \le 2, 1,$ $(thread() \rightarrow v \{ v = fib(x-1) \})()$ + fib(x-2)); }

Execution of threaded fib(6)

18/53



19/53

Lazy execution

Lazy execution

- Up to now the execution order of each thread follows textual order, when a statement comes as the first in sequence it will execute, whether or not its results are needed later
- The execution scheme is called **eager execution**, or **supplydriven** execution
- Another execution order is that a statement is executed only if its results are needed some in the program
- This scheme is called **lazy evaluation**, or **demand-driven** evaluation

22/53

24/5

Example

- define B = f1(X); define C = f2(X); define A = B + C;
- If f1 and f2 are pure, delay evaluation of B and C until it is needed to evaluate A

Lazy Evaluation and Stor

• Without

lazy
 define fib =
 function(n) -> r { ...
 };
 define a = fib(10);

With lazy
 define fib =
 function(n) -> r { ...
 };
 define a = fib(10);

store not lazy

fib: function(n) -> r { ... } a: 55

store lazy

fib: function(n) -> r { ... } a: 55

a

Concurrency

```
define my_thread =
thread(a) -> b {
    b = fib(a);
};
```

```
define v1 =
my_thread(1000);
define v2 =
my_thread(2000);
std.print(v1 + v2);
```

define v1 =
fib(1000);
define v2 =
fib(2000);
std.print(v1 + v2);

25 / 53

Map/Reduce

- Map is a higher-order function that applies a given functor to each element of a list map([1, 2, 3, 4, 5], square) -> [1, 4, 9, 16, 25]
- Reduce is a higher-order function that applies a reduction to a data set reduce([1, 2, 3, 4, 5], add) -> 15

MapReduce

Concurrent Map

• Split the dataset • function tmap(list, functor) -> value { define rest = cond(length(list) == 0, [], (thread() -> value { value = tmap(list[1:], functor); })()); value = functor(list[0]) + rest; };



Concurrent Reduce

• Split the dataset and apply the functor one more time

```
•function treduce(list, functor) -> value {
  cond(length(list) == 2,
    value = functor(list[0], list[1]),
       define threduce = thread(list) -> value
        value = treduce(list, functor);
      3;
      define half = length(list) / 2;
      define v1 = threduce(list[0:half - 1])
      define v2 = threduce(list[half:length(list)]);
      value = functor(v1, v2);
    });
};
```

Big Data Processing (1/2)

• MapReduce is a framework for batch processing of big data.





• Big data:



29/53

Big Data Processing (2/2)

The MapReduce idea:

- ^o Data sets are too big to be analyzed by one machine
- ^o Using multiple machines has the same complications, regardless of the application/ analysis
- ^o Pure functions enable an abstraction barrier between data processing logic and coordinating a distributed application

MapReduce Evaluation Model (1/2)

- Map phase: Apply a mapper function to all inputs, emitting intermediate key-value pairs
 - ^o The mapper takes an iterable value containing inputs, such as lines of text
 - ^o The mapper yields zero or more key-value pairs for each input





31 / 53

I.U



30 / 5

MapReduce Evaluation Model (2/2)

- **Reduce phase**: For each intermediate key, apply a reducer function to accumulate all values associated with that key
- ^o The reducer takes an iterable value containing intermediate key-value pairs
- ^o All pairs with the same key appear consecutively
- ^o The reducer yields zero or more values, each associated with that intermediate key



33 / 53

35 /

MapReduce Execution Model (1/2)



34/53

36/5

MapReduce Execution Model (2/2)



ManReduce example

- From a 1.1 billion people database (facebook?), we want to know the average number of friends per age
- In SQL: SELECT age, AVG(friends) FROM users GROUP BY age
 In MapReduce:

o the total set of users in splitted different
users_set = { user... }
function age_to_friends(user)
{
 send(user.age,
 user.friends.size);
}

```
o The keys are shuffled and assigned to reducers
function average_friends(age,
friends_counts):
{
    var r = reduce(friends_counts, 0,
        function(friends_count, s) -> r
        {
            r = friends_count.size + s
        });
        send(age, r / friends.size);
}
```

MapReduce Assumptions

- Constraints on the mapper and reducer:
- ^o The mapper must be equivalent to applying a deterministic pure function to each input independently
- ^o The reducer must be equivalent to applying a deterministic pure function to the sequence of values for each key
- Benefits of functional programming:

^o When a program contains only pure functions, call expressions can be evaluated in any order, lazily, and in parallel

 In MapReduce, these functional programming ideas allow:

° Consistent results, however computation is partitioned

^o Re-computation and caching of results, as needed

37 / 53

MapReduce Benefits

 Fault tolerance: A machine or hard drive might crash

^o The MapReduce framework automatically re-runs failed tasks

 Speed: Some machine might be slow because it's overloaded

 $^{\circ}$ The framework can run multiple copies of a task and keep the result of the one that finishes first

- Network locality: Data transfer is
 - The framework tries to schedule map tasks on the machines that hold the data to be processed

38/53

Stream

Stream (1/2)

 So far, wait for the end of the execution of a thread to do the next computation define a = generator(X); define b = map1(a, Y); define c = map2(b, Z);

 No reason to wait for a to be finished to start on b (and b for c...)



. . .

Stream (2/2)

A stream is a sequence of message
A stream is first-in-first-out (FIFO)
The producer augments the stream with new messages, the consumer reads the messages, one by one

Stream communication

- Producer, that produces incremently the elements
- **Transducer(s)**, that transforms the elements of the stream
- **Consumer**, that accumulate the results



Consumer Pattern

```
thread consumer(State, Input, Final, Consume) -> ret
{
    define r = Input.reader();
    function c(State)
    {
        cond(r.wait(),
            c(Consume(r.next(), State)),
            ret = Final(State));
    };
    c(State);
}
```

Consumer Pattern: Example

```
define c = consume(0, I,
  function(State) {
  return(State); },
  function(Next, State) {
  return(State + Next); });
```


45 / 53

Transducer Pattern

thread transducer(State, Input, More, Produce, Transform) -> ret

```
{
    define s = stream();
    ret = s.ouput();
    define r = Input.reader();
    function t(State)
    {
        cond(r.wait(), {
            var nState = Transform(r.next(), State);
            cond(not More or More(State), {
               s.send(Produce(nState));
               t(nState)} ) );
    };
    t(state);
}
```

Transducer Pattern: Examples

```
define t1 = transducer(null, I, null,
  function(State) { return 2 * State; },
  function(Next, State) { return Next;
});
define t2 = transducer(null, I,
  function(State) { return State < 5; },
  function(State) { return State; },
  function(Next, State) { return Next;
});
```





46 / 5

Limitation of eager stream processing streams

- The producer might be much faster than the consumer
- This will produce a large intermediate stream that requires potentially unbounded memory storage

49/53

Applications

 Image, Video, Sound processing
 Large numerical computation, for instance, prime number computation

Solutions

• Three

- 1Play with the speed of the different threads, i.e. play with the scheduler to make the producer slower
- 2Create a bounded buffer, say of size N, so that the producer waits automatically when the buffer is full
- 3Use demand-driven approach, where the consumer activates the producer when it need a new element (Lazy evaluation)
- The last two approaches introduce the notion of flow-control between concurrent activities (very common)

50/53

Stream and purity

• The full stream chain is definitely pure function my_stream(...) -> ret { define p1 = producer(...); define t1 = transducer(..., p1, ...); define t2 = transducer(..., t1, ...); ret = consumer(..., t2, ...); } • No-side effects, fully deterministic • But: thread f(...) -> ret { var s = stream(); ret = s.ouput(); s.send(...); // side-effect? }



Conclusion

- Data-driven concurrency
 - $^{\circ}$ Dataflow variables
 - $^{\circ}$ Implicit synchronisation
- Lazy execution
- Stream

53 / 53