

Formal Representation, Scheduling, and Extracting Parallelism

Erik Saule

esaule@uncc.edu

Parallel and Distributed Programming

08/30/17

Learning Outcomes

At the end of this session you will know how to

- Give two representations of parallel codes

- Compute metrics on dependency graphs

- Interpret metrics of dependency graphs in term of parallel execution

- Represent a schedule of a parallel code on some processors

- Apply an algorithm to build a schedule

- Find the dependencies from a sequential code and express the dependencies as a graph

Learning Outcomes

At the end of this session you will know how to

- Give two representations of parallel codes

- Compute metrics on dependency graphs

- Interpret metrics of dependency graphs in term of parallel execution

- Represent a schedule of a parallel code on some processors

- Apply an algorithm to build a schedule

- Find the dependencies from a sequential code and express the dependencies as a graph

The associated assignment will show you how to

- Formulate the parallelism of simple problems as a DAG

- Extract additional parallelism from classical problem that do not seem to exhibit parallelism

Outline

- 1 Representations
- 2 Scheduling
- 3 Extracting Parallelism
- 4 Assignment (start in class)
- 5 Further

The conflict graph representation

Conflict graph

Used to represent a set of tasks that can be executed in any order but that use a common resource.

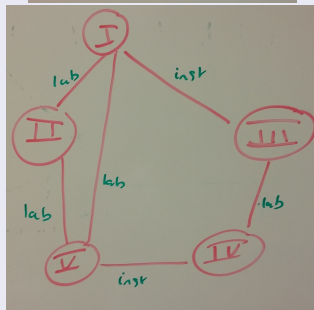
Undirected graph with an edge that connect tasks with a conflict.

Often, the problem to solve is to color the vertices with different colors so that two neighboring vertices have a different color.

NP-Complete problem but greedy heuristics are good.

Example

class	Instn	Lab.
I	A	1
II	B	1
III	A	2
IV	C	2
V	C	1



The conflict graph representation

Conflict graph

Used to represent a set of tasks that can be executed in any order but that use a common resource.

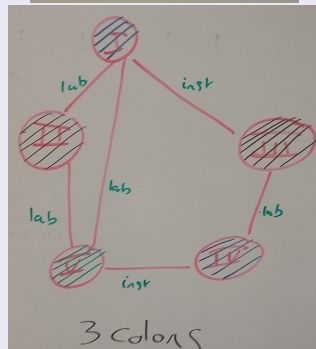
Undirected graph with an edge that connect tasks with a conflict.

Often, the problem to solve is to color the vertices with different colors so that two neighboring vertices have a different color.

NP-Complete problem but greedy heuristics are good.

Example

class	Instn	Lab.
I	A	1
II	B	1
III	A	2
IV	C	2
V	C	1



The dependency graph representation


DAG representation

Represent tasks as vertices.

Represent x before y using a $x \rightarrow y$ directed edge.

The graph is always without cycles.

Example



The dependency graph representation

DAG representation


Represent tasks as vertices.

Represent x before y using a $x \rightarrow y$ directed edge.

The graph is always without cycles.


Processing time required associated with vertices, often denoted p_i

Example



A lemon pie recipe

- (1) break 2 eggs and split the white and yoke
- (2) cut 125g of butter in cubes
- (3) mix yoke and 70g of sugar+5cl of water
- (4) mix 250g of flour with butter
- (5) mix (3) and (4) and make a ball
- (6) spread (5)
- (7) heat oven to 180C
- (8) put crust (6) in pie pan
- (9) wash 4 lemons
- (10) peel two lemons from (9) and finely cut them
- (11) press the four lemons from (9) and (10)
- (12) mix lemons(11), peel(10, 160g of sugar, 1 sp of flour
- (13) cook slowly (12)
- (14) whip 3 eggs
- (15) mix (14) and (13) and cook fast whipping
- (16) empty (15) in (17)
- (17) cook (8) for 20 minutes
- (18) wait until (16) cools



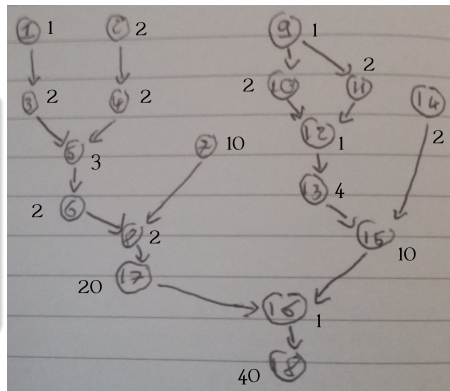
Metrics: Work

Work

Total amount of work that is to perform on the application.

Simply the sum of all processing times.

Often denoted $\sum p_i$



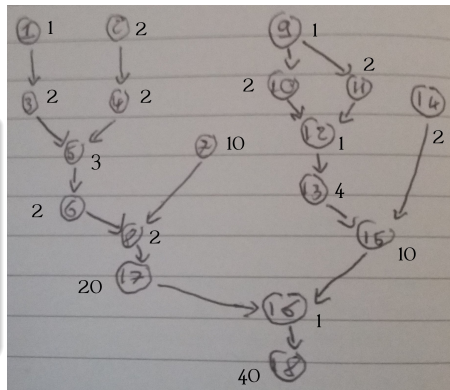
Metrics: Work

Work

Total amount of work that is to perform on the application.

Simply the sum of all processing times.

Often denoted $\sum p_i$



Here $\sum p_i = 107$

Metrics: Work

Work

Total amount of work that is to perform on the application.

Simply the sum of all processing times.


Often denoted $\sum p_i$

Usage

On P processors, the application can not be processed faster than $\frac{\sum p_i}{P}$.

$\frac{\sum p_i}{P}$ is a **lower bound** of the **makespan**.

$$C_{max} \geq \frac{\sum p_i}{P}$$



Here $\sum p_i = 107$

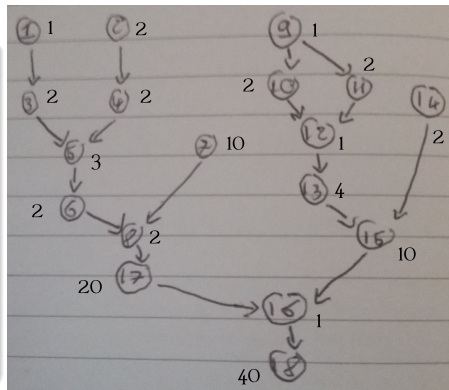
Metrics: Width

Width

Maximum number of tasks that do not have direct dependencies, or transitive dependencies.

Maximum number of independent tasks.

Sometimes called the longest antichain.




Metrics: Width

Width

Maximum number of tasks that do not have direct dependencies, or transitive dependencies.

Maximum number of independent tasks.

Sometimes called the longest antichain.



Here the width is 6.

Metrics: Width

Width

Maximum number of tasks that do not have direct dependencies, or transitive dependencies.

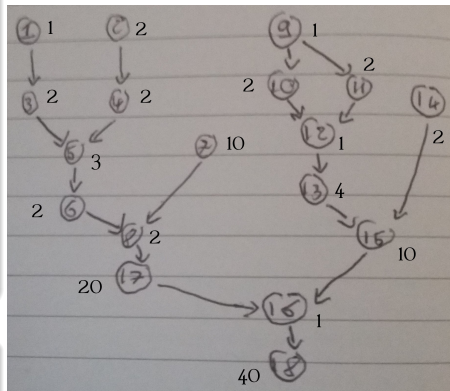
Maximum number of independent tasks.

Sometimes called the longest antichain.

Usage

Maximum number of useful processors.

$\forall P > \text{Width}, S(P) = S(\text{Width})$



Here the width is 6.


Metrics: Width

Width

Maximum number of tasks that do not have direct dependencies, or transitive dependencies.

Maximum number of independent tasks.

Sometimes called the longest antichain.



Here the width is 6.

How to find ?

Usage

Maximum number of useful processors.

$$\forall P > Width, S(P) = S(Width)$$

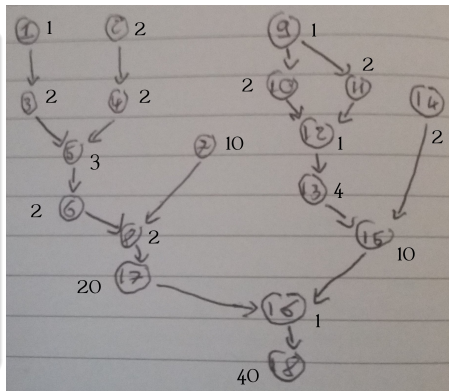
Metrics: Width

Width

Maximum number of tasks that do not have direct dependencies, or transitive dependencies.

Maximum number of independent tasks.

Sometimes called the longest antichain.



Usage

Maximum number of useful processors.

$\forall P > \text{Width}, S(P) = S(\text{Width})$

Here the width is 6.

How to find it ?

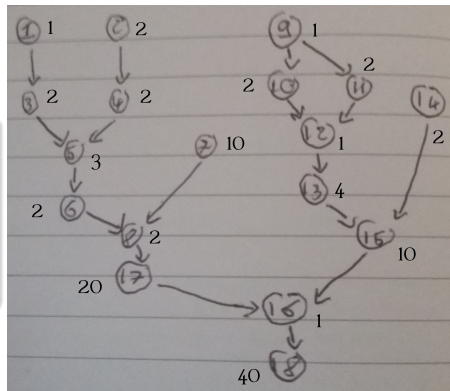
The problem is NP-Complete.
So trial and error.

Metrics: Critical Path

Critical Path

Longest chain of dependency
(in term of processing time)

The length of the chain is often
denoted CP , or T_{∞} .




Metrics: Critical Path

Critical Path

Longest chain of dependency
(in term of processing time)

The length of the chain is often
denoted CP , or T_{∞} .



Here $7 \rightarrow 8 \rightarrow 17 \rightarrow 16 \rightarrow 18$.

$CP = 73$

Metrics: Critical Path

Critical Path

Longest chain of dependency
(in term of processing time)


The length of the chain is often denoted CP , or T_{∞} .

Usage

Whichever way the algorithm unfolds, the critical path will have to be done.

The length of the critical path is a lower bound to the makespan

$$C_{max} \geq CP$$



Here $7 \rightarrow 8 \rightarrow 17 \rightarrow 16 \rightarrow 18$.
 $CP = 73$

Metrics: Critical Path

Critical Path

Longest chain of dependency
(in term of processing time)

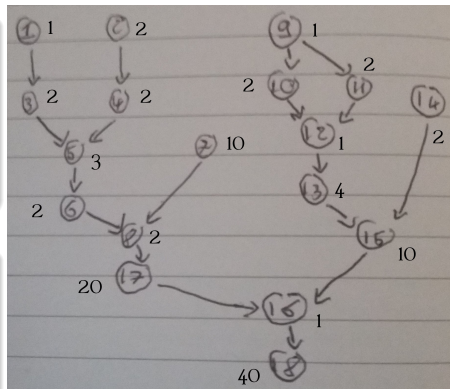
The length of the chain is often denoted CP , or T_{∞} .

Usage

Whichever way the algorithm unfolds, the critical path will have to be done.

The length of the critical path is a lower bound to the makespan

$$C_{max} \geq CP$$



Here $7 \rightarrow 8 \rightarrow 17 \rightarrow 16 \rightarrow 18$.

$$CP = 73$$

How to find it?

Metrics: Critical Path

Critical Path

Longest chain of dependency
(in term of processing time)

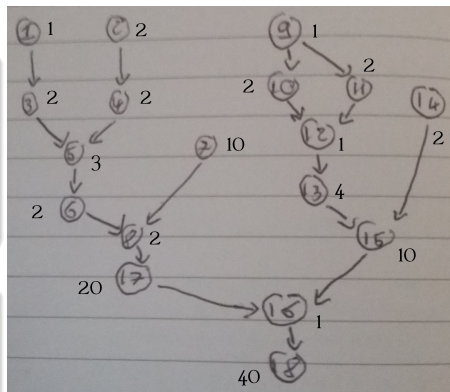
The length of the chain is often denoted CP , or T_{∞} .

Usage

Whichever way the algorithm unfolds, the critical path will have to be done.

The length of the critical path is a lower bound to the makespan

$$C_{max} \geq CP$$



Here $7 \rightarrow 8 \rightarrow 17 \rightarrow 16 \rightarrow 18$.

$$CP = 73$$

How to find it?

Recursively, from the roots down

$$L(x) = p_x + \max_{y \in \Gamma^-(x)} L(y)$$

Activities

(see handout)

Outline

- 1 Representations
- 2 Scheduling**
- 3 Extracting Parallelism
- 4 Assignment (start in class)
- 5 Further

Scheduling

Problem

A DAG of tasks

Processing times p_i

A number of processors

A schedule

Two functions mapping

task t to processor $\pi(t)$

task t to a time interval

$[\sigma(t); C(t)[$. $C(t) = \sigma(t) + p_t$

no two tasks execute on the same processor simultaneously.

Goal

Minimize $C_{\max} = \max C(i)$

Scheduling

Problem

A DAG of tasks

Processing times p_i

A number of processors

A schedule

Two functions mapping

task t to processor $\pi(t)$

task t to a time interval
 $[\sigma(t); C(t)[$. $C(t) = \sigma(t) + p_t$

no two tasks execute on the
same processor simultaneously.

Goal

Minimize $C_{\max} = \max C(i)$

In practice

Lots of variants:

with preemption (interrupting tasks).

with migration (moving a task to an other processor).

with restricted processor allocation (some tasks can only go on some processors).

with unknown p_i

sometimes the dependencies are only known at runtime.


Representing schedules

Gantt Chart

A 2D depiction with time and processors as axes.

makes it easy to see the
makespan

and idle times




Representing schedules

Gantt Chart

A 2D depiction with time and processors as axes.

makes it easy to see the makespan


and idle times



Constrained DAG

Extract one chain per processor that respect dependency.

Highlights dependencies between processors.




Representing schedules

Gantt Chart

A 2D depiction with time and processors as axes.

makes it easy to see the makespan


and idle times



Constrained DAG

Extract one chain per processor that respect dependency.

Highlights dependencies between processors.



Scheduling optimally is hard!

NP-Completeness

Finding the best schedule is a NP-hard problem.

(That means, you are unlikely to find the optimal in a short amount of time)

Scheduling optimally is hard!

NP-Completeness

Finding the best schedule is a NP-hard problem.

(That means, you are unlikely to find the optimal in a short amount of time)

Independent tasks variants

If all tasks are independent
still NP-Hard

If all tasks are independent and the
number of processors is fixed
still NP-Hard (but weakly)

Scheduling optimally is hard!

NP-Completeness

Finding the best schedule is a NP-hard problem.

(That means, you are unlikely to find the optimal in a short amount of time)

Independent tasks variants

If all tasks are independent
still NP-Hard

If all tasks are independent and the number of processors is fixed
still NP-Hard (but weakly)

Variants

If all tasks have $p_i = 1$ (sometimes called UET)

still NP-Hard

If all tasks have $p_i = 1$, and makespan is 3

still NP-Hard

If all tasks have $p_i = 1$, and makespan is 3, and the graph is bipartite

still NP-Hard

Scheduling optimally is hard!

NP-Completeness

Finding the best schedule is a NP-hard problem.

(That means, you are unlikely to find the optimal in a short amount of time)

Independent tasks variants

If all tasks are independent

still NP-Hard

If all tasks are independent and the number of processors is fixed

still NP-Hard (but weakly)

Variants

If all tasks have $p_i = 1$ (sometimes called UET)

still NP-Hard

If all tasks have $p_i = 1$, and makespan is 3

still NP-Hard

If all tasks have $p_i = 1$, and makespan is 3, and the graph is bipartite

still NP-Hard

Scheduling is HARD!

List Scheduling for independent tasks is a 2-approximation algorithm

A greedy algorithm

Consider the tasks in any (unspecified) order.


Pick a task and schedule it as early as possible on the first available processor.

List Scheduling for independent tasks is a 2-approximation algorithm

A greedy algorithm

Consider the tasks in any (unspecified) order.

Pick a task and schedule it as early as possible on the first available processor.



Proof

Let i be the last task to complete.

i starts at time $\sigma(i)$.

At time $\sigma(i)$, all the machines were busy.

$$\sigma(i) \leq \frac{\sum_j p_j}{m}$$

$$C_{max} = C(i) = \sigma(i) + p_i$$

$$C_{max} \leq \frac{\sum_j p_j}{m} + p_i$$

$$C_{max} \leq C_{max}^* + C_{max}^* \leq 2C_{max}^*$$

And if you do the equation well:


$$C_{max} \leq \frac{\sum_j p_j - p_{max}}{m} + p_{max}$$

How bad can List Scheduling really be ?

The idea

List Scheduling says "Given any order of the task". Let us use it to make it the worse possible.

Counter example



List Scheduling is a $(2 - \frac{1}{m})$ -approximation algorithm and a counter example reaching this bound exists. The bound is said to be tight.

Is this the best algorithm ?

Is there a polynomial algorithm with guaranteed performance better than 2

Largest Processing Time first

The idea

List Scheduling reaches 2 because the last task is the largest one.

Algorithm

Sort task by non-increasing order of processing times. Use List Scheduling.
Complexity $O(n \log n + nm)$

Approximation ratio of LPT

If the most loaded machine has one task : $C_{max} = p_{max}$ is optimal.

If the most loaded machine has two tasks : one can show the mapping is optimal (skipped)

If the most loaded machine has k tasks : the imbalance is less than $\frac{C_{max}}{k}$ and the ratio is $\frac{k+1}{k}$.


Theorem

LPT is a $(\frac{4}{3} - \frac{1}{3m})$ -approximation algorithm and the bound is tight.

List Scheduling for DAG

Algorithm

For DAGs, list scheduling is similar to the one for independent task. The algorithm is greedy, look at which tasks are free, and pick one of them to schedule on a free processor.



LS is a 2-approximation

The makespan is decomposed in red and green part.

Red is one chain of the graph and is smaller than the longest chain in the graph (critical path).

Green is a fully occupied machine and is smaller the total work of the problem.

Both are smaller than a lowerbound on the best makespan.

$$C_{max} \leq \frac{\sum p_i}{m} + \sum_{i \in CP} p_i$$

$$C_{max} \leq \frac{T_1}{m} + T_{\infty}$$

$$C_{max} \leq 2C_{max}^*$$

By working the equation a bit harder: $C_{max} \leq \frac{T_1 - T_{\infty}}{m} + T_{\infty}$

Scheduling: Summary

Scheduling

Finding for each task a processor for it to run on.

Finding for each task a time interval for it to run on.

Finding the best schedule is hard!
(even in sub-cases.)

Scheduling: Summary


Scheduling

Finding for each task a processor for it to run on.

Finding for each task a time interval for it to run on.

Finding the best schedule is hard!
(even in sub-cases.)

Gantt Chart



Scheduling: Summary


Scheduling

Finding for each task a processor for it to run on.

Finding for each task a time interval for it to run on.

Finding the best schedule is hard!
(even in sub-cases.)

Gantt Chart



Three heuristic algorithms

For independent tasks:

List Scheduling (pick any available task and schedule it)

LPT (sort by decreasing processing time. Then LS)

For DAGs:

List Scheduling (pick any available task and schedule it)

(see handout)

can you find the optimal schedule ?

Outline

- 1 Representations
- 2 Scheduling
- 3 Extracting Parallelism**
- 4 Assignment (start in class)
- 5 Further

More art than science!

Answering questions like :

- Does it matter if these two things are swapped?

- Would the code still be correct if... ?

- Can we do something completely different?

- Is there a different expression that can compute the same value?

Analyze dependencies in sequential code

Code analysis

Granularize code
statement
loop iteration
function call

The sequential code induces strict ordering. “Do these two tasks have to be this strictly ordered?”

Weird Fibonacci

```
int fibo_v[N];


void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```

Analyze dependencies in sequential code - example

Weird Fibonacci

```
int fibo_v[N];

void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```



Let's start with the first loop.

Make one vertex per loop iteration.


There are no dependencies.

Analyze dependencies in sequential code - example

Weird Fibonacci

```
int fibo_v[N];

void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```



Let's use the two nested loops

Each iteration of j depends on the previous iteration

Maybe some dependencies between different i iteration


Let's forget the nested loop j and coarsen the graph (higher level)

Analyze dependencies in sequential code - example

Weird Fibonacci

```
int fibo_v[N];

void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```



Let's look at some iteration of i , say until 23.

These are just the tasks

What are the dependencies ?


Let's consider just task $i = 17$ for the moment

Analyze dependencies in sequential code - example

Weird Fibonacci

```
int fibo_v[N];

void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```



$i = 17$ will read $\text{fibo_v}[12]$, $\text{fibo_v}[13]$, $\text{fibo_v}[14]$

$i = 12$ writes $\text{fibo_v}[12]$

So $i = 17$ can not happen before task $i = 12$ completes


Similarly, $i = 17$ depends on $i = 13$ and $i = 14$ completions

Analyze dependencies in sequential code - example

Weird Fibonacci

```
int fibo_v[N];

void fibo() {
    for (int i=0; i<6; ++i) {
        fibo_v[i] = 1;
    }
    for(int i=6; i<N; ++i) {
        fibo_v[i] = 0;
        for (int j=0; j<3; ++j) {
            fibo_v[i] += fibo_v[i-j-3];
        }
    }
}
```



Similarly, all tasks for $i \geq 6$, have 3 inputs.

Can you find the width, work, CP of this graph?

As a function of N ?

Usual dependencies

$x \rightarrow y$

When x is before y in sequential and there is
flow dependence (read-after-write (RAW))

y reads a variable written by x

$out[x] = 12; in[y] += out[x];$

anti-dependence (write-after-read (WAR))

y writes a variable read by x

$out[x] += in[y]; in[y] += 12;$

output-dependence (write-after-write (WAW))

y writes a variable written by x

$glob = x; glob = y$

Note that input-dependence (read-after-read (RAR)) usually does not matter.

Mutual exclusion and resolution

Mutual exclusion

Appropriate when two blocks of code can not run simultaneously.

But either order is valid.

Essentially pairs a conflict graph to the DAG.


I usually just leaves a note on the graph to express mutual exclusion.

Example

```
int val[N];
int re[2];

int f (int i);

void red () {
    for (int i=0; i<N; ++i) {
        val[i] = f(i);
        re[i%2] += val[i];
    }
}
```



Mutual exclusion and resolution

Mutual exclusion

Appropriate when two blocks of code can not run simultaneously.

But either order is valid.

Essentially pairs a conflict graph to the DAG.


I usually just leaves a note on the graph to express mutual exclusion.

Example

```
int val[N];
int re[2];

int f (int i);

void red () {
    for (int i=0; i<N; ++i) {
        val[i] = f(i);
        re[i%2] += val[i];
    }
}
```



Mutual exclusion and resolution

Mutual exclusion

Appropriate when two blocks of code can not run simultaneously.

But either order is valid.

Essentially pairs a conflict graph to the DAG.

I usually just leaves a note on the graph to express mutual exclusion.


Not too much help!

Example


```
int val[N];
int re[2];

int f (int i);

void red () {
    for (int i=0; i<N; ++i) {
        val[i] = f(i);
        re[i%2] += val[i];
    }
}
```




A story of whip and oven



And now for something completely different

Example


```
void red () {  
    for (int i=0; i<N; ++i) {  
        val[i] = f(i);  
        re[i%2] += val[i];  
    }  
}
```



And now for something completely different


Example

```
void red () {  
    for (int i=0; i<N; ++i) {  
        val[i] = f(i);  
        re[i%2] += val[i];  
    }  
}
```



Do it differently

```
void red () {  
    for (int i=0; i<N; ++i)  
        val[i] = f(i);  
    for (int i=0; i<N; ++i)  
        re[i%2] += val[i];  
}
```



Changes dependency structures.

Useful if $f(i)$ is expensive.

Assumes $f(i)$ s are independent.

Outline

- 1 Representations
- 2 Scheduling
- 3 Extracting Parallelism
- 4 Assignment (start in class)
- 5 Further

(see text.)

Extract parallelism

Transform

Reduction (int +, string +, float +)

Find First (array and list)

Prefix Sum

Merge Sort

Outline

- 1 Representations
- 2 Scheduling
- 3 Extracting Parallelism
- 4 Assignment (start in class)
- 5 Further**

cilk on graphs metrics

Conflict graph and coloring:

Conflict graphs: <http://math.cmu.edu/~bkell/21110-2010s/conflict-graphs.html>

A. H Gebremedhin, F. Manne, Alex Pothén. What Color Is Your Jacobian? Graph Coloring for Computing Derivatives. Siam Review 2005.

M. Deveci, E. Boman, K. Devine, and S. Rajamanickam. Parallel Graph Coloring for Manycore Architectures. IPDPS 2016.

Scheduling:

Scheduling is NP-Hard: M. Garey and D. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. Freeman. 1979.

LS for independent tasks: R. Graham. Bounds for certain multiprocessing anomalies. Bell System Technical Journal. 1966

LPT and LS with precedence: R. Graham. Bounds on Multiprocessing Timing Anomalies. SIAM Journal on Applied Mathematics. 1969.

Typical compiler optimization:

Loop fission: https://en.wikipedia.org/wiki/Loop_fission

Loop tiling: https://en.wikipedia.org/wiki/Loop_tiling

Various: https://en.wikipedia.org/wiki/Compiler_optimization