

Estimation of Parallel Complexity with Rewriting Techniques

Christophe Alias^{*}, Carsten Fuhs[†], Laure Gonnord[‡]

^{*} INRIA & LIP (UMR CNRS/ENS Lyon/UCB Lyon1/INRIA), Lyon, France,
christophe.alias@ens-lyon.fr

[†] Birkbeck, University of London, United Kingdom,
carsten@dcs.bbk.ac.uk

[‡] University of Lyon & LIP (UMR CNRS/ENS Lyon/UCB Lyon1/INRIA), Lyon,
France, laure.gonnord@ens-lyon.fr

15th International Workshop on Termination
September 5, 2016, Obergurgl, Austria



Why Automatic Parallelization?

- Most computers are parallel (end of Dennard scaling...)
 - Writing/debugging a parallel program is (horribly) difficult
- ~> Automation is required...

Why Automatic Parallelization?

- Most computers are parallel (end of Dennard scaling...)
- Writing/debugging a parallel program is (horribly) difficult
- ↪ Automation is required...

Challenges:

- How to represent the computation? ↪ data dependencies
- How much parallelism? ↪ data dependencies
- Which parallelism (scheduling)? ↪ data dependencies
- Which resource allocation? ↪ data dependencies

Why Automatic Parallelization?

- Most computers are parallel (end of Dennard scaling...)
- Writing/debugging a parallel program is (horribly) difficult
- ↪ Automation is required...

Challenges:

- How to represent the computation? ↪ data dependencies
 - How much parallelism? ↪ data dependencies
 - Which parallelism (scheduling)? ↪ data dependencies
 - Which resource allocation? ↪ data dependencies
- ☹ Bad news: checking data dependencies is undecidable.

Focus on regular imperative programs (polyhedral model)

- 😊 Unifying framework for program parallelization
- 😊 Exact set of dependencies, all the parallelism is found
- 😞 Scalability issues

Focus on regular imperative programs (polyhedral model)

- ☺ Unifying framework for program parallelization
- ☺ Exact set of dependencies, all the parallelism is found
- ☹ Scalability issues

Over-approximate the data dependencies

- ☺ Scalable
- ☹ Can be (very) rough and miss most of the parallelism

Focus on regular imperative programs (polyhedral model)

- ☺ Unifying framework for program parallelization
- ☺ Exact set of dependencies, all the parallelism is found
- ☹ Scalability issues

Over-approximate the data dependencies

- ☺ Scalable
- ☹ Can be (very) rough and miss most of the parallelism


Contributions:

- Assess the parallel complexity of (some) recursive programs
- ... using **monotone interpretations**
- Extends of the polyhedral model to recursive programs

Parallel Complexity

```
//Compute  $y = Ax$   
for  $i := 0$  to  $N-1$   
   $l_1: y[i] := 0;$   
  for  $j := 0$  to  $N-1$   
     $l_2: y[i] := y[i] + a[i][j]*x[j];$ 
```


$$\lambda_N = \mathcal{O}(N)$$



- Minimum number of (parallel) computation steps assuming unbounded parallel resources.
- Solved on regular programs (**polyhedral model**)
- **Goal: recursive programs on trees!**


```
for i := 0 to N-1
  ℓ1: y[i] := 0;
  for j := 0 to N-1
    ℓ2: y[i] := y[i] + a[i][j]*x[j];
```


$$\lambda_N = \mathcal{O}(N)$$



- 1 Divide an execution e into a sequence of operations \mathcal{O}_e
→ How to represent/approximate e ? which grain?

```
for i := 0 to N-1
  l1: y[i] := 0;
  for j := 0 to N-1
    l2: y[i] := y[i] + a[i][j]*x[j];
```

$$\lambda_N = \mathcal{O}(N)$$




- 1 Divide an execution e into a sequence of operations \mathcal{O}_e
→ How to represent/approximate e ? which grain?
- 2 Compute the dependencies: $\rightarrow_e \subseteq \mathcal{O}_e \times \mathcal{O}_e$
→ Impact of approximation?

Parallel Complexity – Methodology

```
for i := 0 to N-1
  l1: y[i] := 0;
  for j := 0 to N-1
    l2: y[i] := y[i] + a[i][j]*x[j];
```

$$\lambda_N = \mathcal{O}(N)$$




- 1 Divide an execution e into a sequence of operations \mathcal{O}_e
→ How to represent/approximate e ? which grain?
- 2 Compute the dependencies: $\rightarrow_e \subseteq \mathcal{O}_e \times \mathcal{O}_e$
→ Impact of approximation?
- 3 Compute the parallel complexity: $\lambda_e := \text{height}(\rightarrow_e)$
→ How to express λ_e ?

- 1 Parallel complexity of **regular programs**
- 2 Parallel complexity of **recursive programs**

Polyhedral Model at a Glance


```
for i := 0 to 2*N
  l1: c[i] := 0;
for i := 0 to N
  for j := 0 to N
    l2: c[i+j] := c[i+j] + a[i]*b[j];
```



- Automatic parallelization of regular loop nests with arrays

Polyhedral Model at a Glance


```
for i := 0 to 2*N
  l1: c[i] := 0;
for i := 0 to N
  for j := 0 to N
    l2: c[i+j] := c[i+j] + a[i]*b[j];
```



- Automatic parallelization of **regular loop nests with arrays**
- e is decidable and can be analyzed (e.g. with ILP)
 - $\langle l_1, i \rangle : i \in \llbracket 0, 2N \rrbracket$
 - $\langle l_2, i, j \rangle : (i, j) \in \llbracket 0, N \rrbracket^2$
- Key analysis: **array dependencies**, **affine scheduling**

(Affine) Array Dependencies


```
for i := 0 to 2*N
  l1: c[i] := 0;
for i := 0 to N
  for j := 0 to N
    l2: c[i+j] := c[i+j] + a[i]*b[j];
```



- Idea: Given a consumer, find the last producer \rightsquigarrow ILP.



(Affine) Array Dependencies

```
for i := 0 to 2*N
  l1: c[i] := 0;
for i := 0 to N
  for j := 0 to N
    l2: c[i+j] := c[i+j] + a[i]*b[j];
```



- Idea: Given a consumer, find the last producer \rightsquigarrow **ILP**.
- \rightarrow_N is an **affine relation**:
 - $\langle l_2, i-1, j+1 \rangle \rightarrow_N \langle l_2, i, j \rangle \quad : i > 0 \wedge j < N$
 - $\langle l_1, i \rangle \rightarrow_N \langle l_2, 0, i \rangle \quad : 0 \leq i \leq N$
 - $\langle l_1, i \rangle \rightarrow_N \langle l_2, i-N, N \rangle \quad : N < i \leq 2N$

(Affine) Array Dependencies



- Idea: Given a consumer, find the last producer \rightsquigarrow ILP.



- \rightarrow_N is an affine relation:

$$\langle l_2, i - 1, j + 1 \rangle \rightarrow_N \langle l_2, i, j \rangle \quad : i > 0 \wedge j < N$$

$$\langle l_1, i \rangle \rightarrow_N \langle l_2, 0, i \rangle \quad : 0 \leq i \leq N$$


$$\langle l_1, i \rangle \rightarrow_N \langle l_2, i - N, N \rangle \quad : N < i \leq 2N$$

(Affine) Scheduling



- Assign each operation $\langle \ell, \vec{x} \rangle$ with a timestamp $\theta_\ell(\vec{x}) \in \mathbb{N}^{d_\ell}$.
- Correctness: $\langle \ell, \vec{x} \rangle \rightarrow_N \langle \ell', \vec{y} \rangle \Rightarrow \theta_\ell(\vec{x}) \ll \theta_{\ell'}(\vec{y})$
- **Affine schedule**: $\theta_\ell(\vec{x}) = A\vec{x} + \vec{b} \rightsquigarrow \text{ILP}$.

(Affine) Scheduling




- Assign each operation $\langle \ell, \vec{x} \rangle$ with a timestamp $\theta_\ell(\vec{x}) \in \mathbb{N}^{d_\ell}$.
- Correctness: $\langle \ell, \vec{x} \rangle \rightarrow_N \langle \ell', \vec{y} \rangle \Rightarrow \theta_\ell(\vec{x}) \ll \theta_{\ell'}(\vec{y})$
- Affine schedule: $\theta_\ell(\vec{x}) = A\vec{x} + \vec{b} \rightsquigarrow \text{ILP}$.
- Bonus: reverse the order: termination algorithm! [RanK, 2010]

HPC community	TCS community
Data Dependence Graph	Integer transition system
Schedule	Ranking function
Latency	Computational complexity
Recursive schedule	Monotonic interpretations

Target: recursive programs on trees

```
public int treeMax() {  
    int leftMax = Integer.MIN_VALUE;  
    int rightMax = Integer.MIN_VALUE;  
    if (this.left != null)  
        leftMax = this.left.treeMax();  
    if (this.right != null)  
        rightMax = this.right.treeMax();  
    return Math.max(this.val,  
                    Math.max(leftMax, rightMax));  
}
```



- Each node (subtree) of t is an operation of e .
- \rightarrow_e can be encoded as a term rewrite system (TRS):


$$\text{dep}(\text{Tree}(\text{val}, \text{left}, \text{right})) \rightarrow \text{dep}(\text{left})$$
$$\text{dep}(\text{Tree}(\text{val}, \text{left}, \text{right})) \rightarrow \text{dep}(\text{right})$$

- How to schedule (check the termination of) a TRS?

~> With monotone interpretations! **[AProVE, KoAT]**

Putting it all together

```
public int treeMax() {  
    int leftMax = Integer.MIN_VALUE;  
    int rightMax = Integer.MIN_VALUE;  
    if (this.left != null)  
        leftMax = this.left.treeMax();  
    if (this.right != null)  
        rightMax = this.right.treeMax();  
    return Math.max(this.val,  
        Math.max(leftMax, rightMax));  
}
```




Monotone interpretation	Parallel complexity
$[\text{dep}](x_1) = x_1$	$\lambda_t = \mathcal{O}(\text{height}(t))$
$[\text{Tree}](x_1, x_2, x_3) = x_2 + x_3 + 1$	
$[\text{dep}](x_1) = x_1$	$\lambda_t = \mathcal{O}(\text{height}(t))$
$[\text{Tree}](x_1, x_2, x_3) = \max(x_2, x_3) + 1$	

What happens on regular programs?

```
for(i=0; i<=N; i++)
  for (j=0; j<=N; j++)
    //Block S
    {
      m1[i][j] = Integer.MIN_VALUE;
      for(k=1; k<=i; k++)
        m1[i][j] = max(m1[i][j], H[i-k][j] + W[k]);

      m2[i][j] = Integer.MIN_VALUE;
      for(k=1; k<=j; k++)
        m2[i][j] = max(m2[i][j], H[i][j-k] + W[k]);

      H[i][j] = max(0, H(i-1, j-1)+s(a[i], b[i]),
                   m1[i][j], m2[i][j]);
    }
}
```



$\text{dep}(i, j) \rightarrow \text{dep}(i-1, j-1) : 0 \leq i \leq n, 0 \leq j \leq n$

$\text{dep}(i, j) \rightarrow \text{dep}(i-k, j) : 0 \leq i \leq n, 0 \leq j \leq n, 1 \leq k \leq i$

$\text{dep}(i, j) \rightarrow \text{dep}(i, j-\ell) : 0 \leq i \leq n, 0 \leq j \leq n, 1 \leq \ell \leq j$

Result: $[\text{dep}](x_1, x_2) = x_1 + x_2 \quad \lambda_n \leq 2n$

Same as in the polyhedral model!

Position:

- Automatic parallelization can take profit of monotonic interpretations.
- Extension of affine scheduling to recursive programs

Locks:

- How to define/find the best schedule?
- How to count the steps?
- Steps towards a parallelizing compiler:
 - Computation partitioning?
 - Generation of the parallel code given a schedule?