

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

Languages, Compilation and Semantics, LIP Seminar
November, 3rd 2016



- 1 Introduction
- 2 Parallel complexity of **regular programs**
Parallelization \mapsto Termination
- 3 Parallel complexity of **recursive programs**
Termination \mapsto Parallelization
- 4 Conclusions

Parallel Systems






↪ Many models: KPN (and variants), SDF (and variants), etc

Correctness

- Determinism
- Termination

Efficiency

- Latency
- Bandwidth



From the source code:

- Target **latency**? \geq **termination**
- Target **bandwidth**?

Trace $(\Omega_{\mathcal{I}}, \prec_{\text{seq}})$: sequence of operations executed by the program on the input \mathcal{I} .


Undecidable in general!

- Restrict program model \rightsquigarrow Polyhedral model
- Over-approximate $\Omega_{\mathcal{I}}$ \rightsquigarrow Abstract interpretation

Regular Program $\Omega_{\mathcal{I}}$ does not depend on \mathcal{I} .
 \rightsquigarrow Polyhedral model

Polyhedral Model: Affine Loop Nests


```
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];
```



- for loop with arrays + affine constraints

Polyhedral Model: Affine Loop Nests


```
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];
```



- for loop with arrays + affine constraints
- $\Omega_{\mathcal{I}}$ can be encoded with integer polyhedra:
 - $\langle l_1, i \rangle : i \in \llbracket 0, 2N \rrbracket$
 - $\langle l_2, i, j \rangle : (i, j) \in \llbracket 0, N \rrbracket^2$
- Static analysis with ILP: dependences, scheduling, allocation

Tree Model: Tree Traversals

```
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 Ω_t
- Not exactly regular, but Ω_t is decidable!

Data Dependence relate communicating/conflicting operations,

$$\rightarrow \subseteq \Omega \times \Omega$$

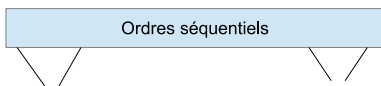
\rightsquigarrow usually split into flow-, anti-, output-dependences

Schedule Assign dates to operations, $\theta : \Omega \rightarrow (\mathcal{D}, <)$


Induced order $\prec_{\theta} = \{(s, t), \theta(s) < \theta(t)\}$

Correctness $s \rightarrow t \implies \theta(s) < \theta(t)$

Hence: $\rightarrow \subseteq \prec_{\theta}$




...



Latency

```
//Compute  $y = Ax$   
for  $i := 0$  to  $N-1$   
   $\ell_1: y[i] := 0;$   
  for  $j := 0$  to  $N-1$   
     $\ell_2: y[i] := y[i] + a[i][j]*x[j];$ 
```

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



Latency λ longest chain of \rightarrow

Degree of sequentiality smallest $s \in \mathbb{N}$ s.t. $\lambda = \mathcal{O}(N^s)$

Counterpart in the Tree Model?

Transition system

- Σ set of states
- $\Sigma_0 \subseteq \Sigma$ initial states
- $\rightarrow \subseteq \Sigma \times \Sigma$ transition relation

Termination Exhibit a **ranking function** $\rho : \Sigma \rightarrow (\mathcal{D}, <)$ s.t.

$$s \rightarrow t \Rightarrow \rho(t) < \rho(s)$$

WCET


$$\lambda = \max\{n, \exists \sigma_0 \in \Sigma_0 : \sigma_0 \rightarrow^n \sigma\}$$

Upper bound:

$$\lambda \leq \#\{\rho(\sigma), \exists \sigma_0 \in \Sigma_0 : \sigma_0 \rightarrow^* \sigma\}$$

Instance: Integer Transition Systems

```
while( $p \neq q$ )  
  if( $p < q$ )  
     $p := p - q$ ;  
  else  
     $q := q - p$ ;
```



- Each state of Σ is a pair $\langle l, \vec{i} \rangle$
- ρ is affine per control point:


$$\rho : \langle l, \vec{i} \rangle \mapsto A_i + b, \quad \Sigma \rightarrow (\mathbb{Z}^p, \ll)$$

Parallelization	Termination
Dependence Schedule Latency	Transition system Ranking function WCET

- 1 Introduction
- 2 Parallel complexity of **regular programs**
Parallelization \mapsto Termination
- 3 Parallel complexity of **recursive programs**
Termination \mapsto Parallelization
- 4 Conclusions

Dependence Analysis


```
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.


Dependence Analysis

```
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**:
 - $\langle l_2, i-1, j+1 \rangle \rightarrow_N \langle l_2, i, j \rangle \quad : i > 0 \wedge j < N$
 - $\langle l_1, i \rangle \rightarrow_N \langle l_2, 0, i \rangle \quad : 0 \leq i \leq N$
 - $\langle l_1, i \rangle \rightarrow_N \langle l_2, i-N, N \rangle \quad : N < i \leq 2N$



Dependence Analysis – Representation



- \rightarrow_N is an **affine relation**:
 $\langle l_2, i - 1, j + 1 \rangle \rightarrow_N \langle l_2, i, j \rangle \quad : i > 0 \wedge j < N$
 $\langle l_1, i \rangle \rightarrow_N \langle l_2, 0, i \rangle \quad : 0 \leq i \leq N$
 $\langle l_1, i \rangle \rightarrow_N \langle l_2, i - N, N \rangle \quad : N < i \leq 2N$
- ... finitely represented as a graph (PRDG)



- Timestamps are vectors of $(\mathbb{N}^{d_\ell}, \ll)$.
- **Affine schedule:** $\theta_\ell : \vec{x} \mapsto A_\ell \vec{x} + B_\ell \vec{N} + \vec{c}_\ell \quad \mathcal{D}_\ell \rightarrow (\mathbb{N}^{d_\ell}, \ll)$
- Can be computed with **ILP**




- Timestamps are vectors of $(\mathbb{N}^{d_\ell}, \ll)$.
- **Affine schedule:** $\theta_\ell : \vec{x} \mapsto A_\ell \vec{x} + B_\ell \vec{N} + \vec{c}_\ell \quad \mathcal{D}_\ell \rightarrow (\mathbb{N}^{d_\ell}, \ll)$
- Can be computed with **ILP**
- **Bonus:** reverse the order: termination algorithm! [RanK, 2010]

- 1 Introduction
- 2 Parallel complexity of **regular programs**
Parallelization \mapsto Termination
- 3 Parallel complexity of **recursive programs**
Termination \mapsto Parallelization
- 4 Conclusions

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 Ω_t .
- \rightarrow 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?
 \rightsquigarrow With monotone interpretations! [AProVE, KoAT]

Given a TRS \rightarrow over a term algebra $\mathcal{T}(\Sigma)$:


- Assign each symbol $a/n \in \Sigma$ with $[a] : \mathbf{R}^n \rightarrow \mathbf{R}$.
- The **monotone interpretation** of a term is:

$$\llbracket f(t_1, \dots, t_n) \rrbracket = [f](\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket)$$

- Correctness: $t \rightarrow u \Rightarrow \llbracket t \rrbracket > \llbracket u \rrbracket$.
- Polynomial interpretations can be computed by 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
$[\text{dep}](x_1) = x_1$ $[\text{Tree}](x_1, x_2, x_3) = x_2 + x_3 + 1$	$\lambda = \mathcal{O}(\text{height}(t))$
$[\text{dep}](x_1) = x_1$ $[\text{Tree}](x_1, x_2, x_3) = \max(x_2, x_3) + 1$	$\lambda = \mathcal{O}(\text{height}(t))$

What happens on polyhedral 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: $[\text{dep}](x_1, x_2) = x_1 + x_2 \quad \lambda \leq 2n$

Same as in the polyhedral model!

Position:

- Two successful experiences of cross fertilization
Parallelization \leftrightarrow *Termination*.
- Monotonic interpretations can benefit to automatic parallelization.
- Natural 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?

Thanks!