

More Definite Results From the Pluto Scheduling Algorithm

By

Athanasios Konstantinidis

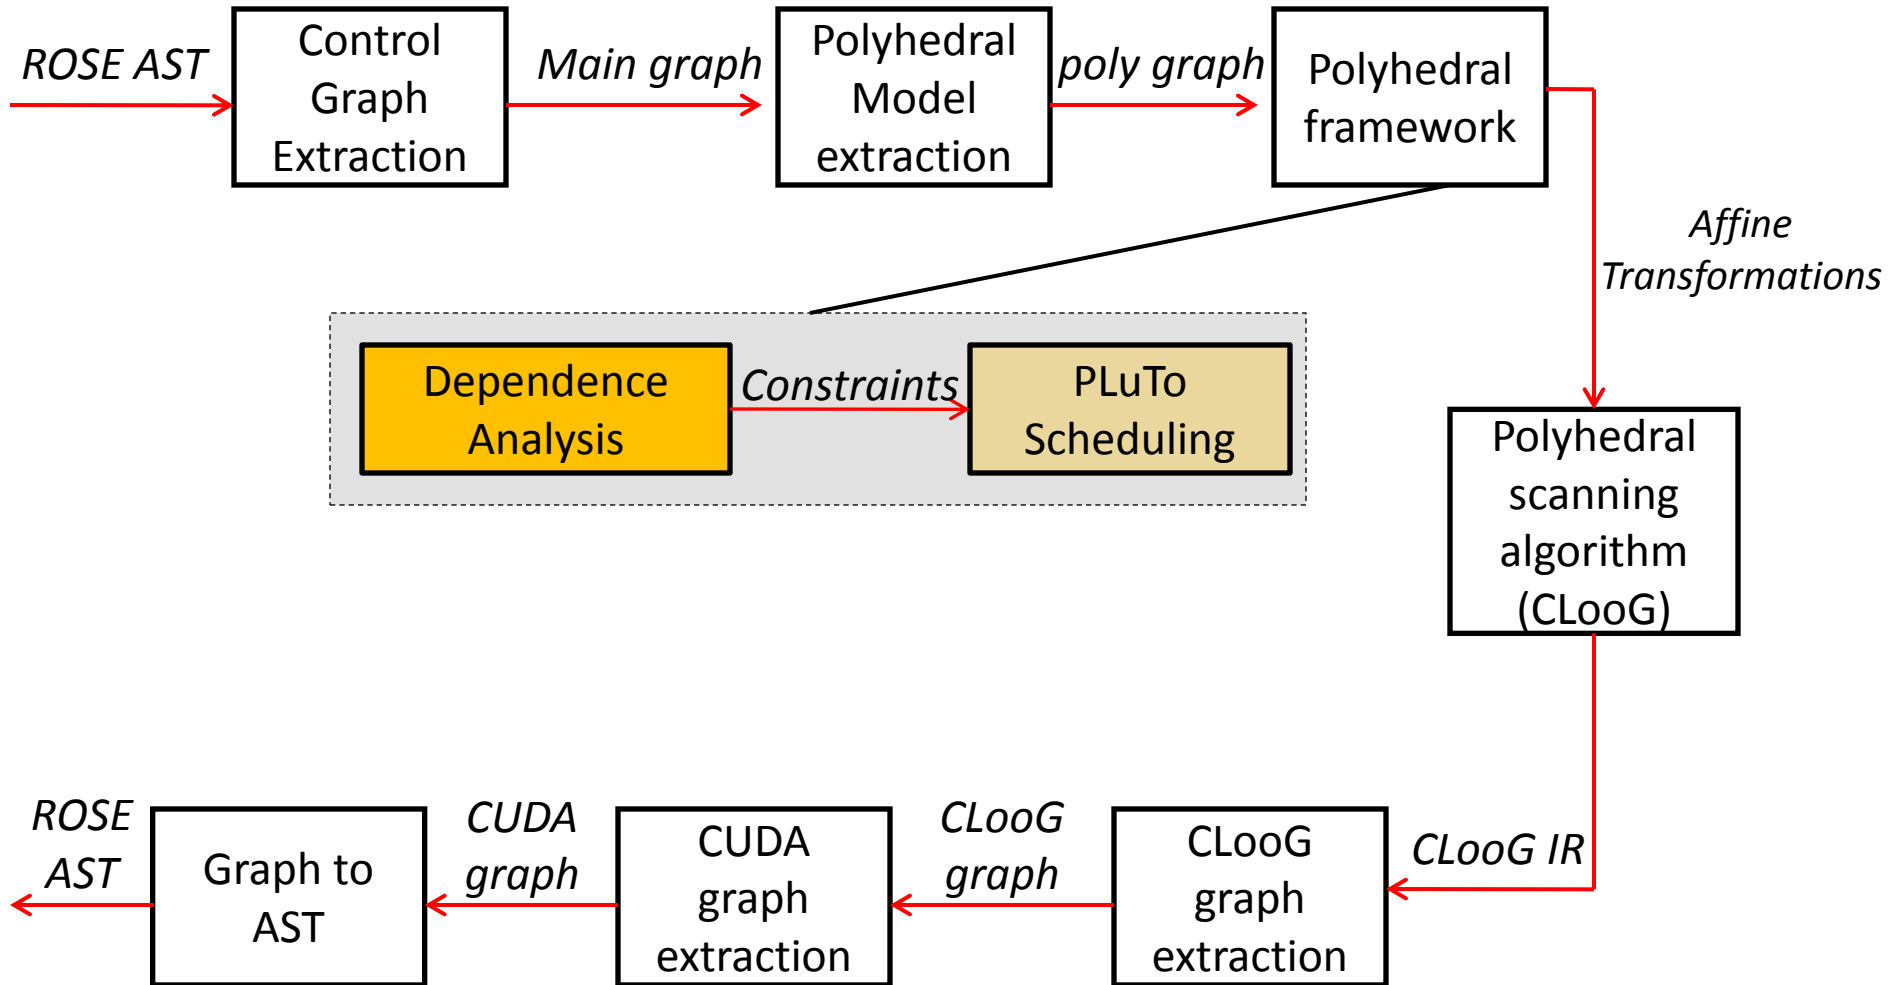
Supervisor

Paul H. J. Kelly

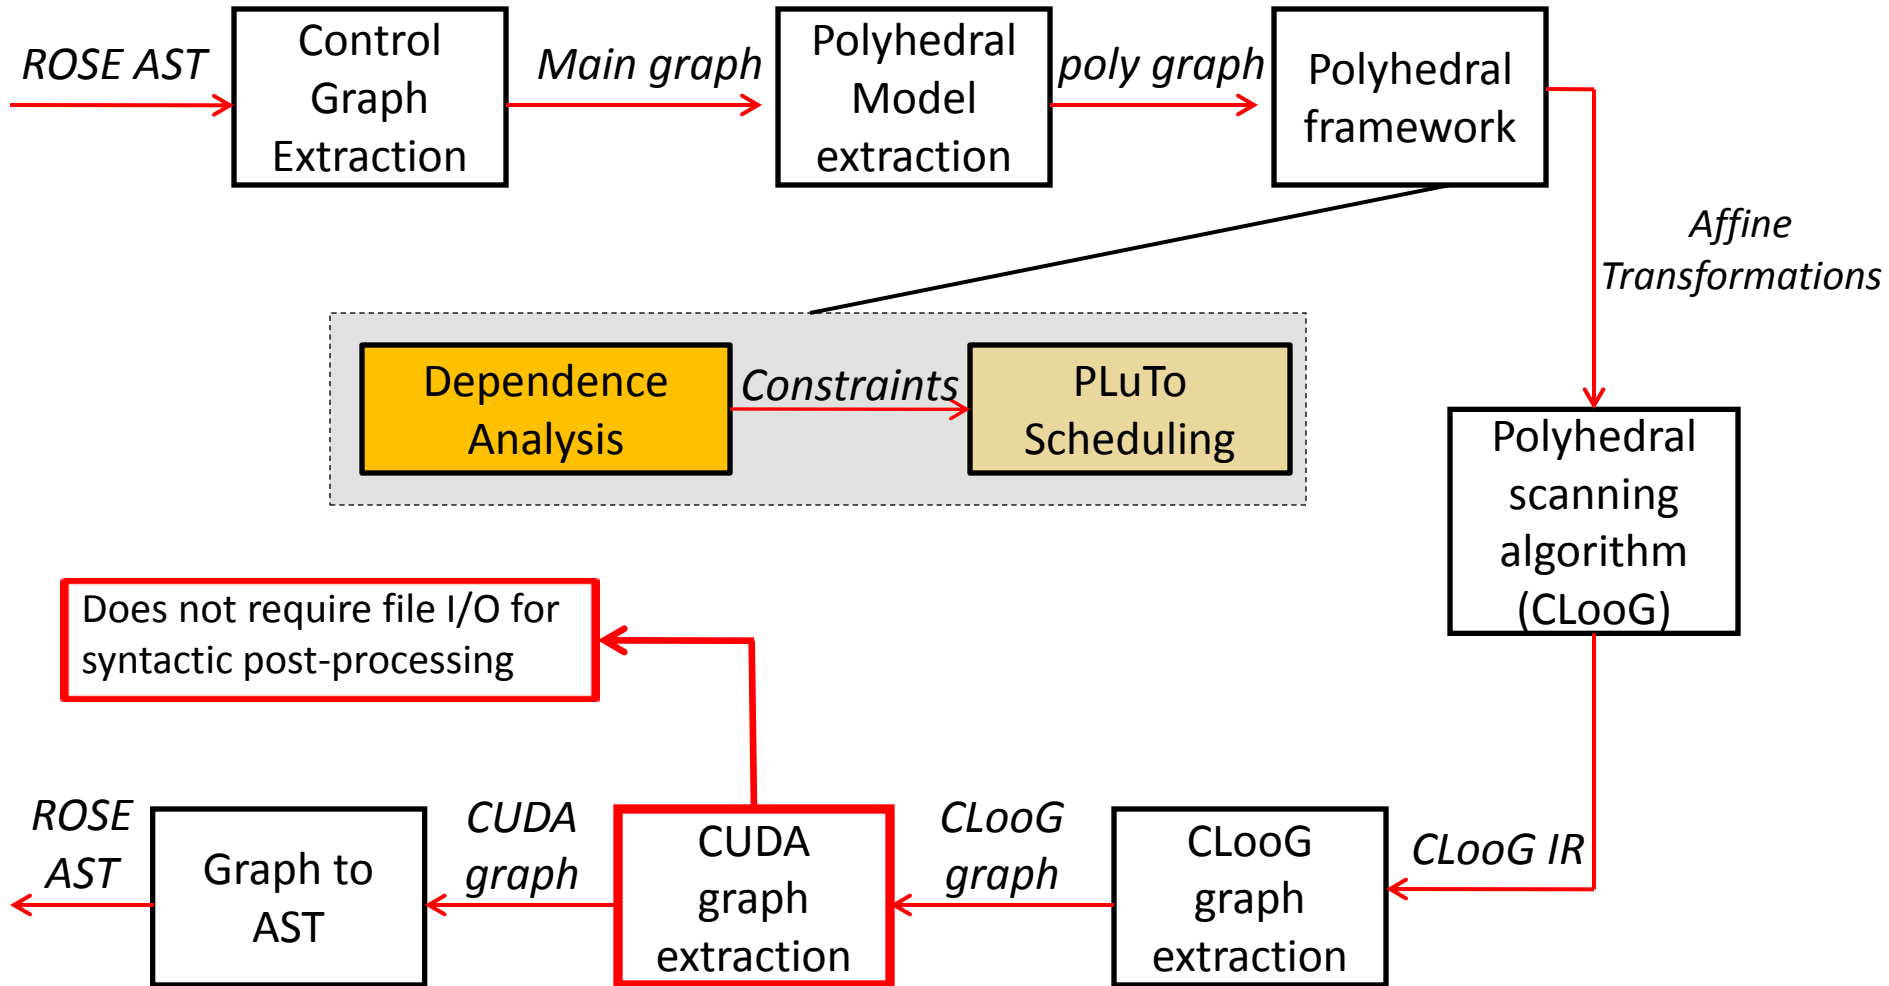
About Me

- PhD student at Imperial College London supervised by Paul H. J. Kelly.
- Compiler and Language support for Heterogeneous parallel architectures (e.g. GPGPUs, Cell BE, Multicore etc.).
- Developing our own source-to-source polyhedral compiler (CUDA back-end).
- Sponsored by EPSRC and Codeplay Software Ltd.

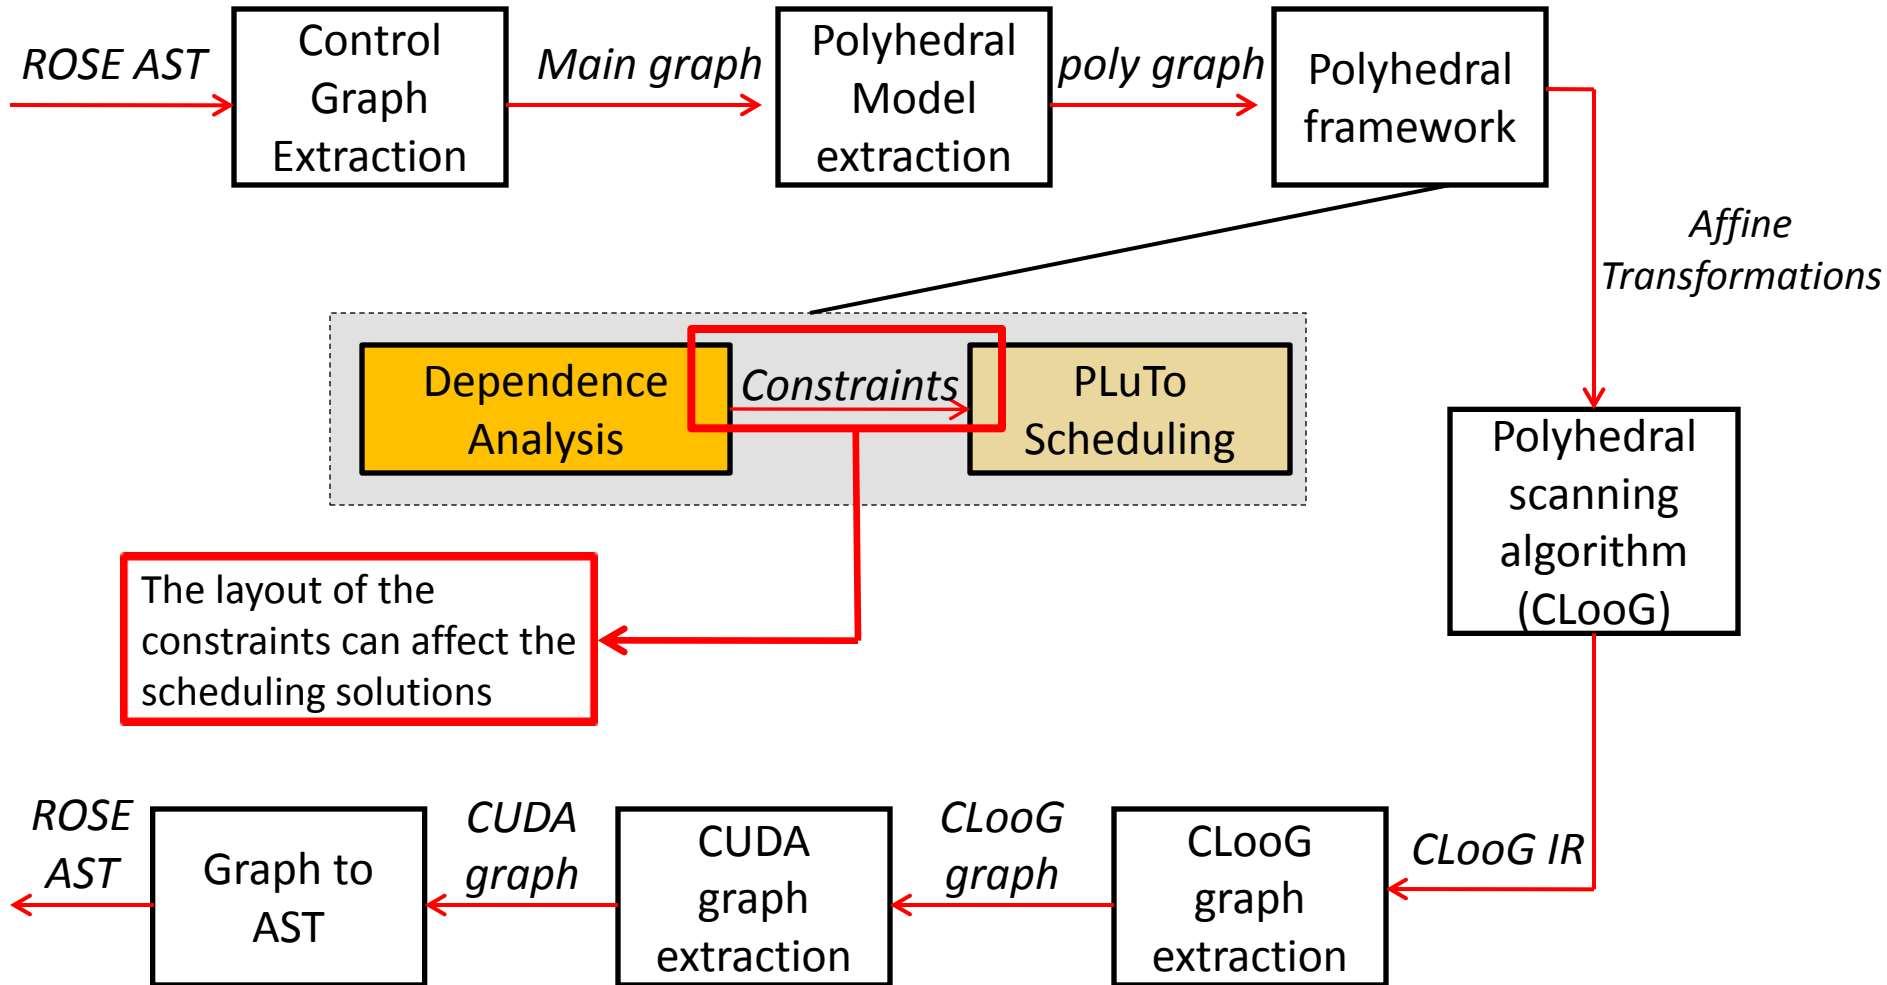
Our Polyhedral Framework



Our Polyhedral Framework

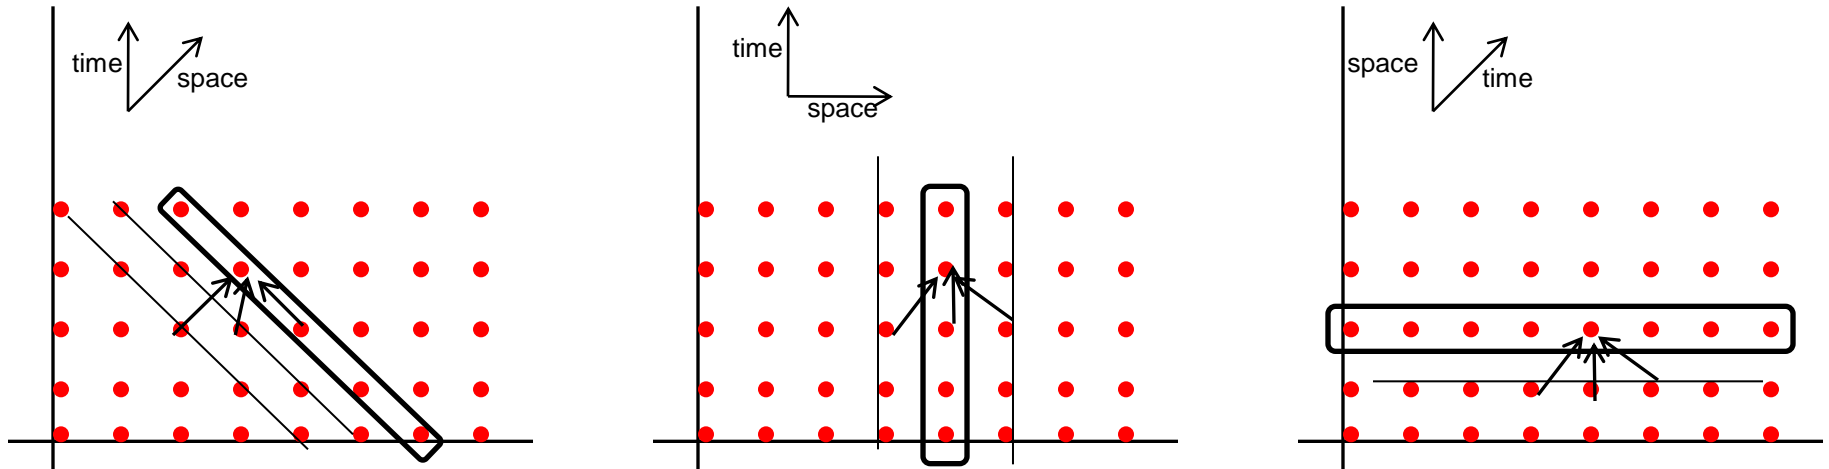


Our Polyhedral Framework



PLuTo Scheduling Algorithm 1

- Iteratively looks for a maximal set of linearly independent affine transforms of the original iteration space.
- An affine transform is a hyperplane representing a loop in the transformed iteration space.
- Each hyperplane needs to respect a set of constraints that guarantee legality and minimum communication between hyperplane instances (i.e. between different loop iterations).



PLuTo Scheduling Algorithm 2

- For each Statement S_i find a schedule $\Phi_i(\vec{x}_i) = \underbrace{[\Phi_{i0} \ \Phi_{i1} \ \cdots \ \Phi_{iN}]}_{MAX + \text{scalar dimensions}}^T \cdot \vec{x}_i$
- Each hyperplane Φ_{ij} of statement S_i is a solution to a global constraint matrix consisted of the following constraints for each dependence edge :

Global Constraint Matrix $M \rightarrow$ **Empty**

$M \leftarrow$ Legality Constraints $[\Phi_{dj} \cdot \vec{x}_d - \Phi_{sj} \cdot \vec{x}_s \geq 0]$

$M \leftarrow$ Communication Bounding [**cost** $\geq \Phi_{dj} \cdot \vec{x}_d - \Phi_{sj} \cdot \vec{x}_s$]

$M \leftarrow$ Non-Trivial Solution Constraints $[\Phi_{ij} \cdot \vec{x}_i \geq 1]$

Solution $\Phi_{ij} \leftarrow$ **Solve(M)**

- **Solve(M)** \rightarrow Uses a Parametric Integer Programming Library (PIP) to find the lexicographic minimum solution.

PLuTo Scheduling Algorithm 3

- Iteratively find as many linearly independent solution as possible

```
Global Constraint Matrix M  $\rightarrow$  Empty  
M  $\leftarrow$  Legality  
M  $\leftarrow$  Communication Bounding  
M  $\leftarrow$  Non-Trivial solution  
While ( Solve(M) ) {  
    M  $\leftarrow$  Linear Independence  
}
```


PLuTo Scheduling Algorithm 4

- If **NO MORE** solutions can be found \rightarrow remove any killed dependences
- If **NO** solution was found \rightarrow cut the dependence graph into Strongly Connected Components (SCC) – *loop distribution* – and remove the killed dependences

```
Global Constraint Matrix M  $\rightarrow$  Empty  
M  $\leftarrow$  Legality  
M  $\leftarrow$  Communication Bounding  
M  $\leftarrow$  Non-Trivial solution  
While ( Solve(M) ) {  
    M  $\leftarrow$  Linear Independence  
}
```

Cut in SCC If **NO** solution is found
Remove Killed dependences

PLuTo Scheduling Algorithm 5

- Iteratively find bands of fully permutable loop nests

do {

Global Constraint Matrix $M \rightarrow \text{Empty}$

$M \leftarrow \text{Legality}$

$M \leftarrow \text{Communication Bounding}$

$M \leftarrow \text{Non-Trivial solution}$

While ($\text{Solve}(M)$) {

$M \leftarrow \text{Linear Independence}$

}

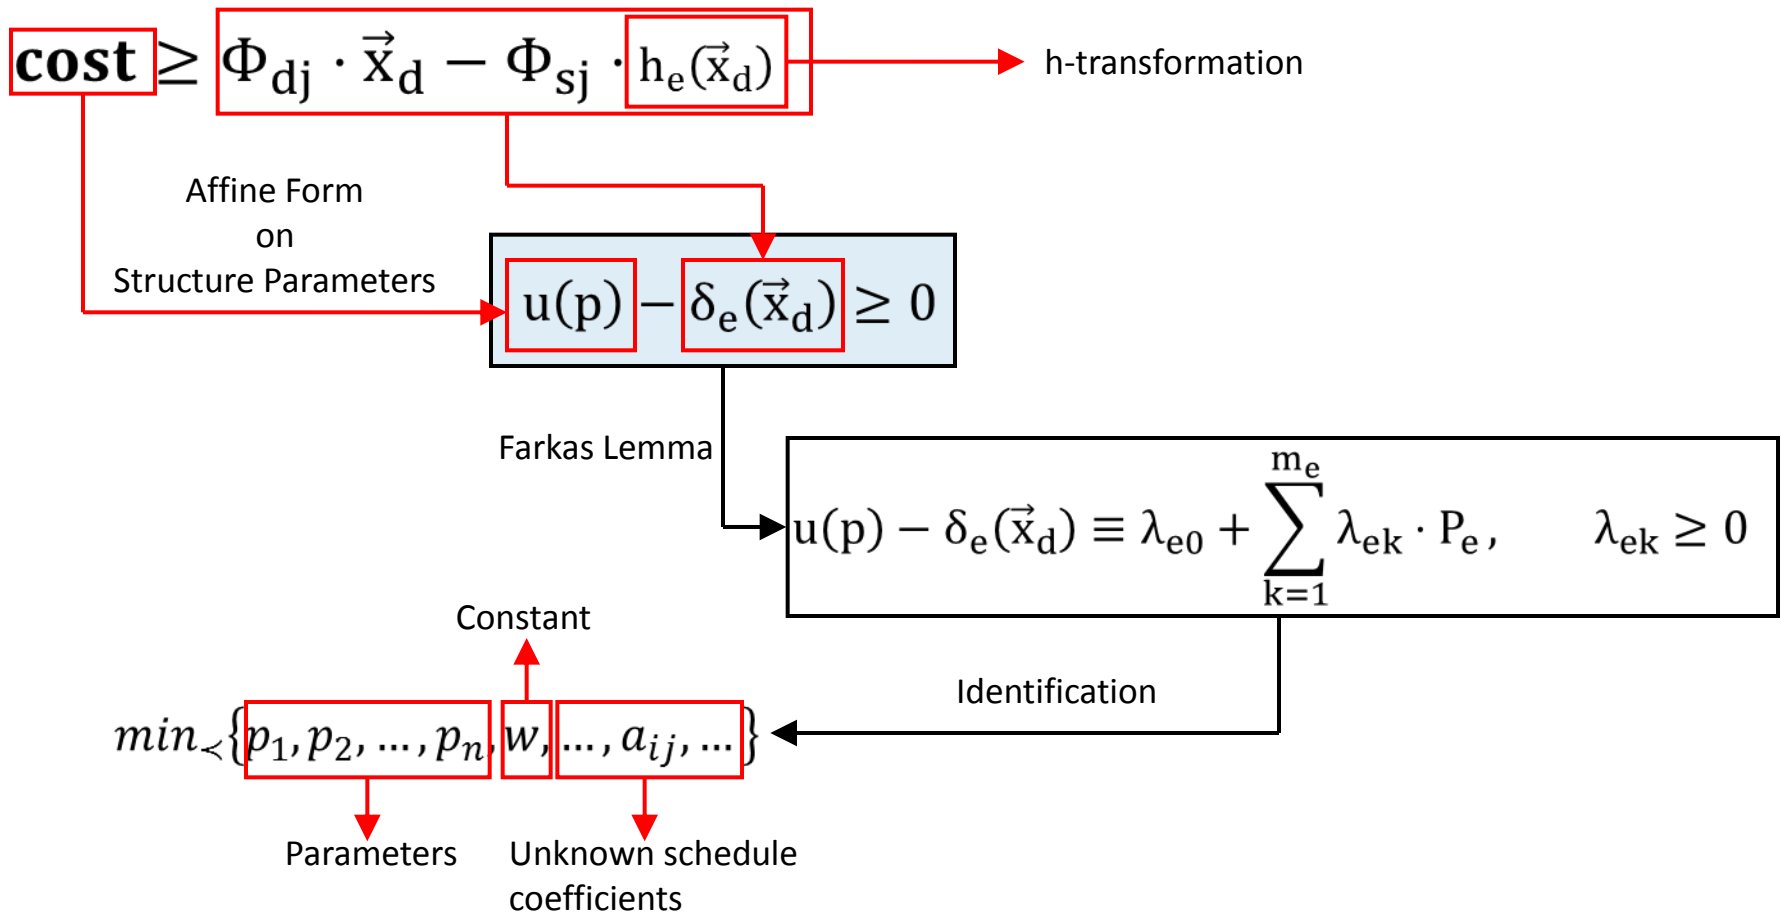
Cut in SCC If **NO** solution is found

Remove Killed dependences

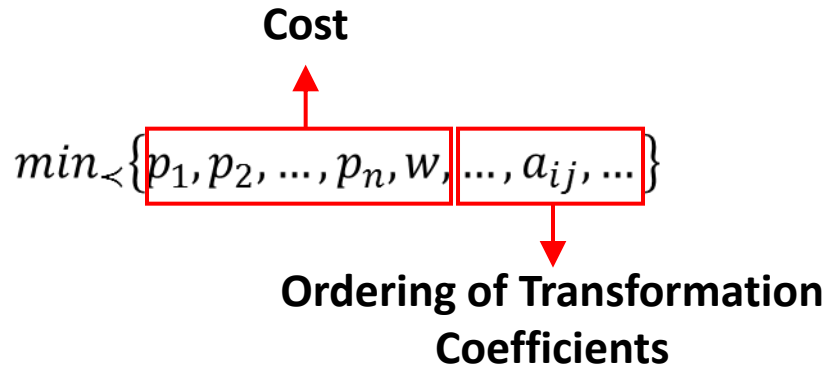
} While ($(total_sols < MAX)$ **OR** $(deps \neq 0)$)

Communication Bounding Constraints

- For every dependence edge e :

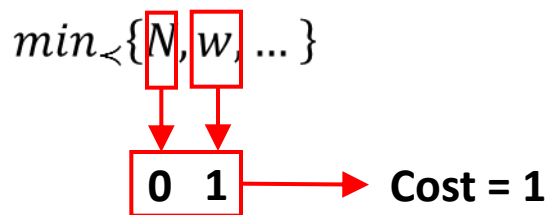


Ordering Sensitivity 1

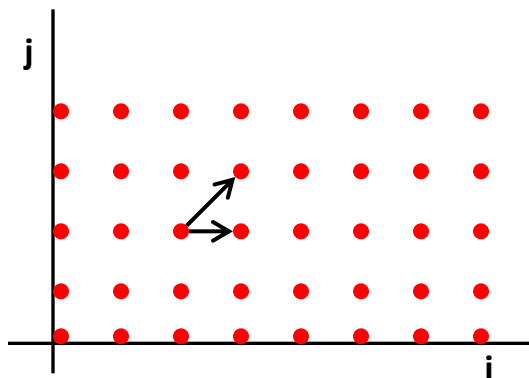


- For the same **Cost** the solution we will get from the PIP solver will eventually depend on the **ordering of the transformation coefficients**.

Ordering Sensitivity (example)



- Minimum **Cost** is **1**.
- No outer parallel loop.

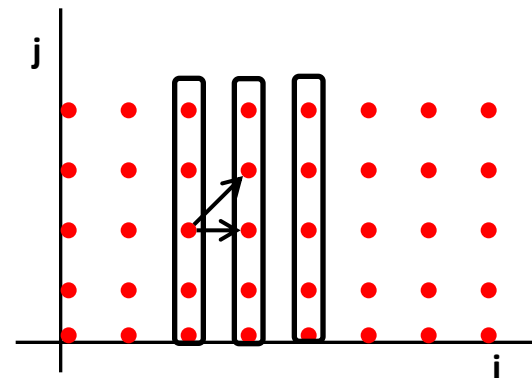
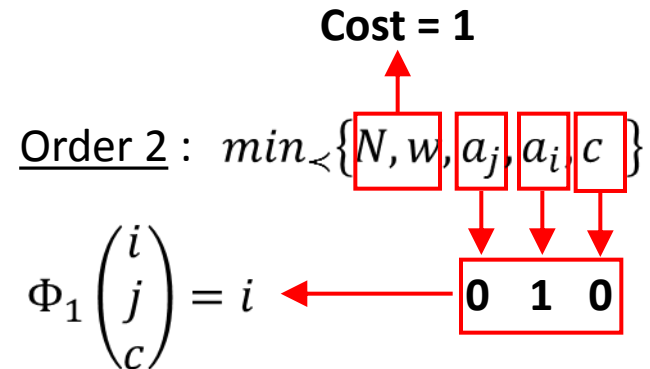
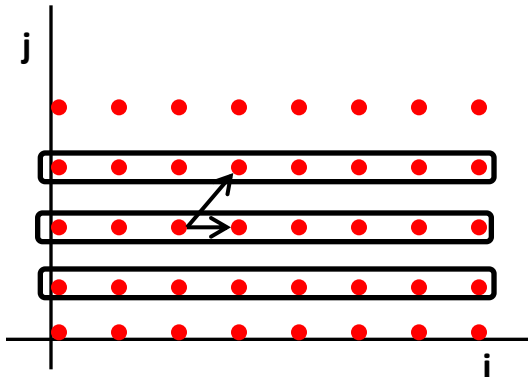
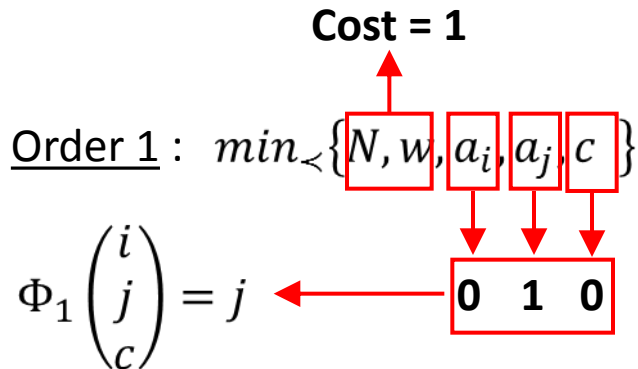


```

for i = 0, N
  for j = 0, N
    A[i][j] = A[i-1][j]*A[i-1][j-1]
  
```

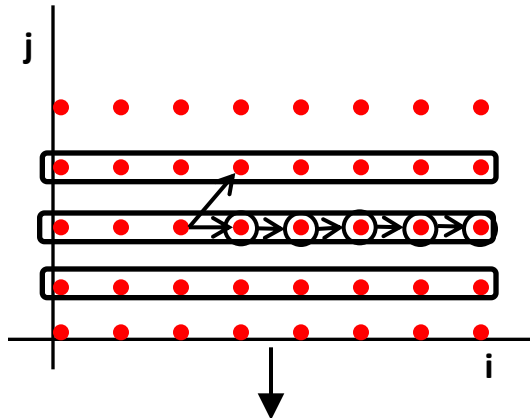
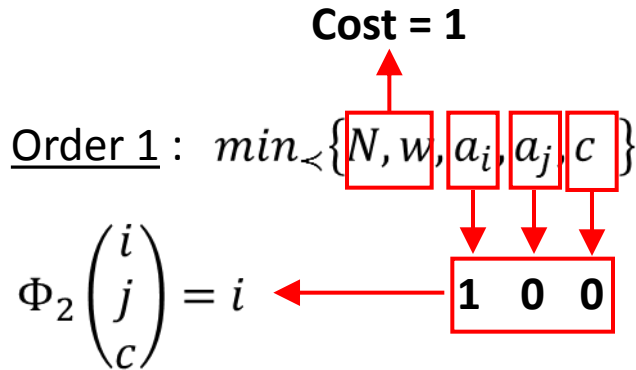
Ordering Sensitivity (example)

- By changing the order of the transformation coefficients we get **two different solutions** both having **Cost = 1**.



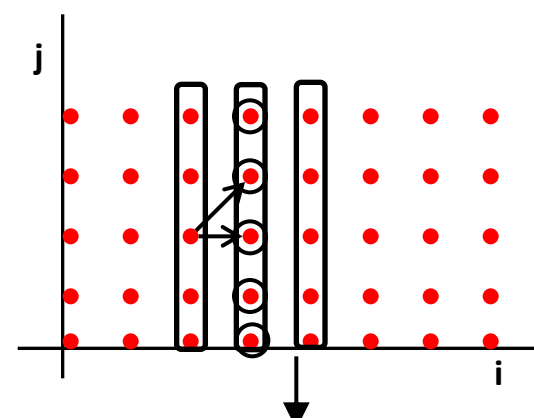
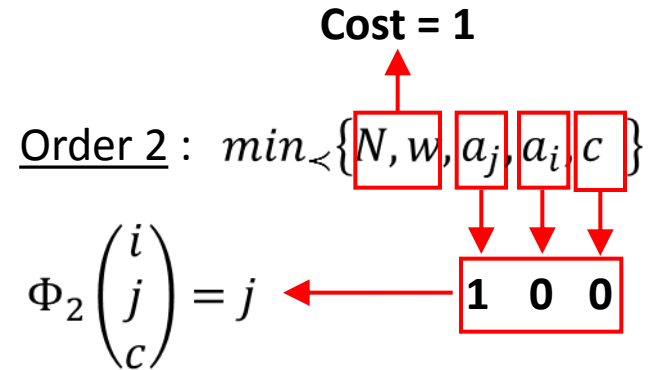
Ordering Sensitivity (example)

- By adding the **linear independence constraints** we get a **second solution**.
- Order 2 yields an **inner loop** that is **fully parallel**.
- Which solution/order is better ?



Pipeline/Wavefront

$$\Phi \begin{pmatrix} i \\ j \\ c \end{pmatrix} = \{\Phi_1, \Phi_2\}$$

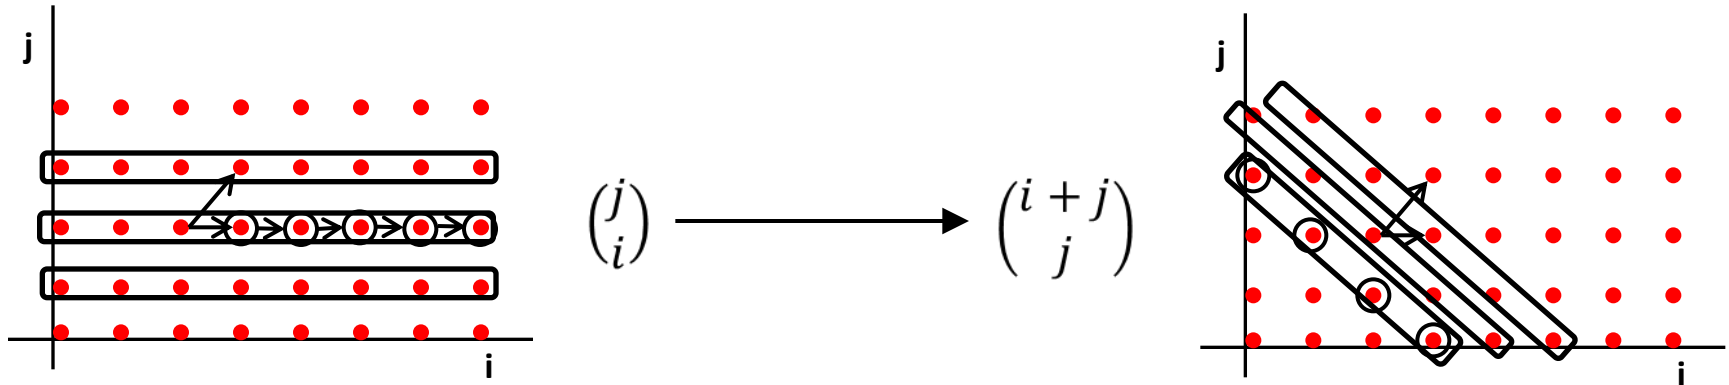


Fully Parallel Inner Loop

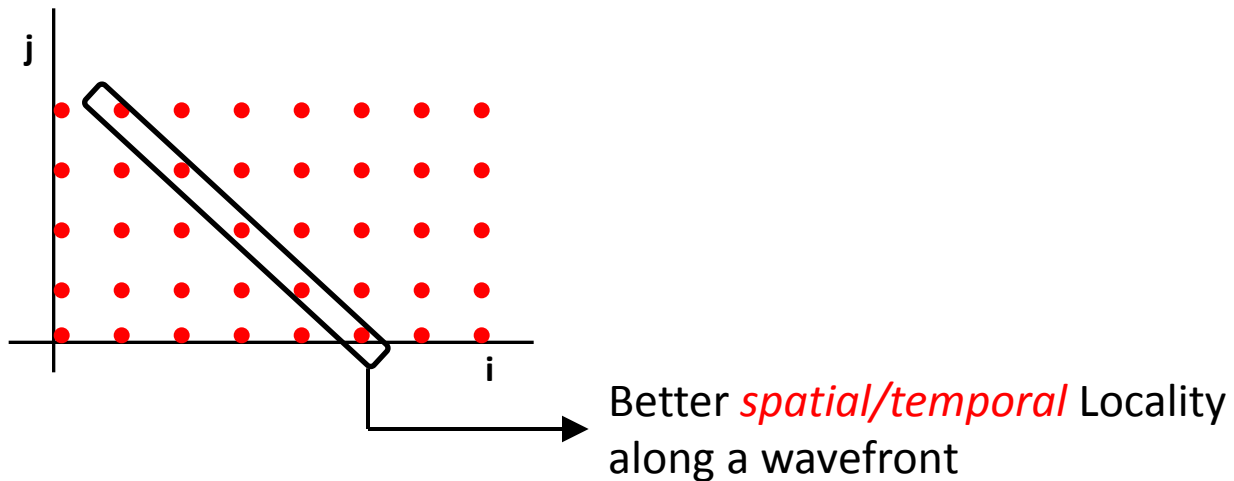
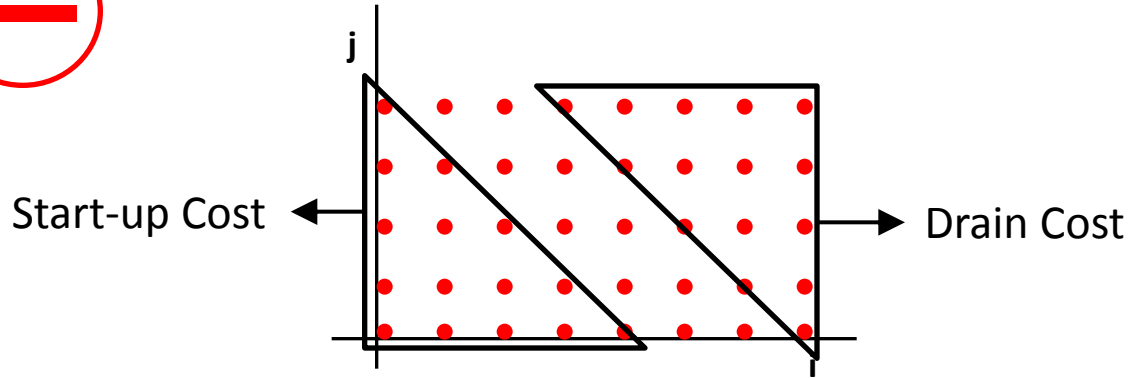
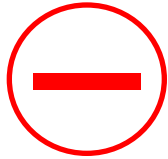
Pipeline Degrees of Parallelism

- N Non-parallel loops can be **transformed into a wavefront/pipeline** consisted of one sequential and $N-1$ parallel loops i.e. degrees of parallelism.

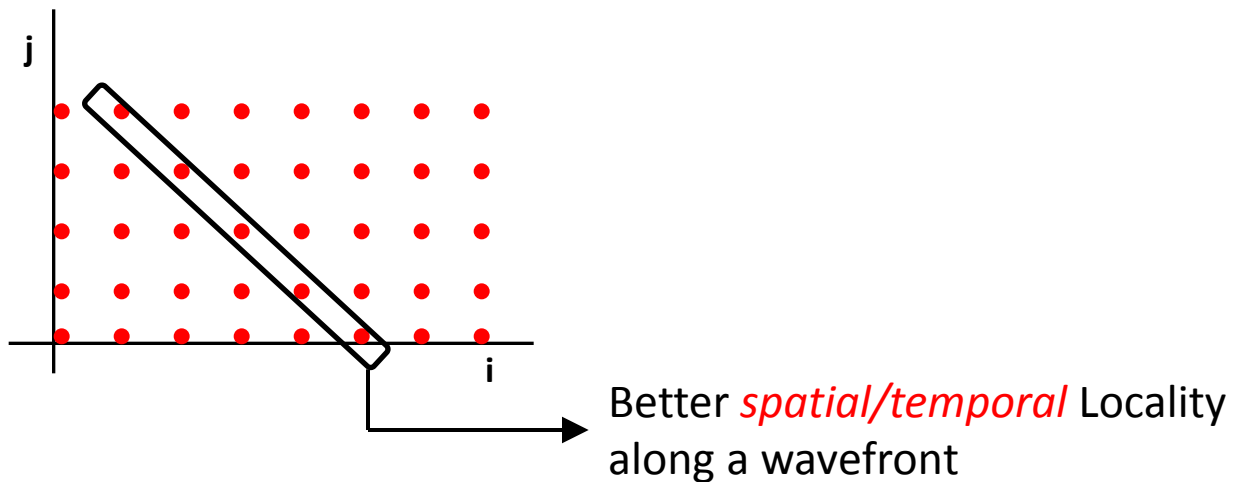
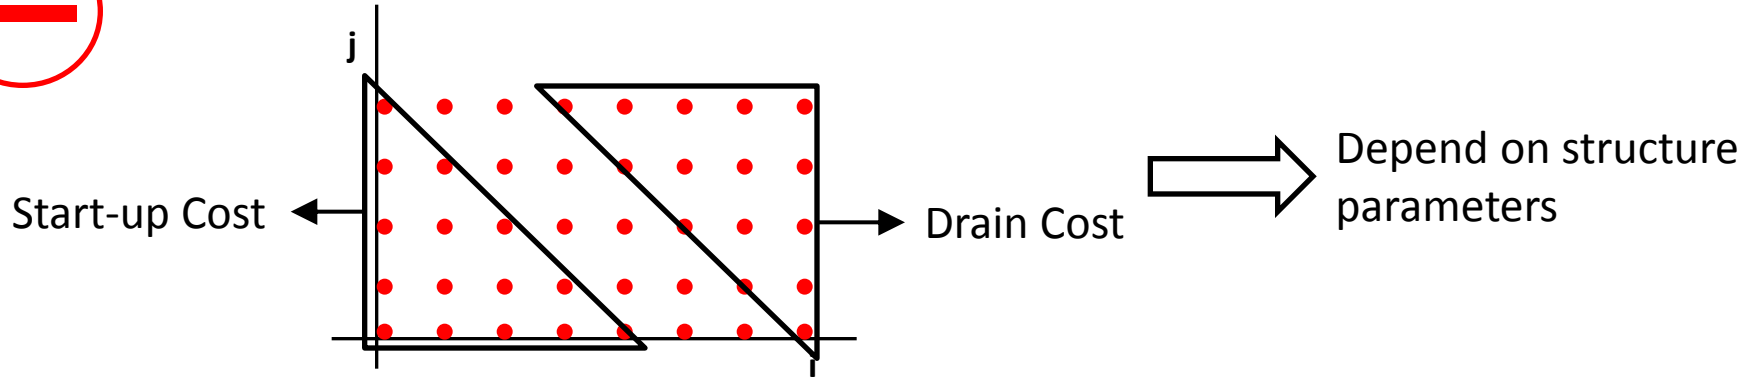
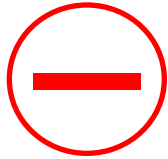
$$\text{Wavefront/pipeline} \leftarrow \begin{pmatrix} \Phi'_1 \\ \Phi'_2 \\ \vdots \\ \Phi'_n \end{pmatrix} = \begin{pmatrix} 1 & 1 & \dots & 1 & 1 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} \Phi_1 \\ \Phi_2 \\ \vdots \\ \Phi_n \end{pmatrix} \rightarrow \text{Non-parallel loops}$$



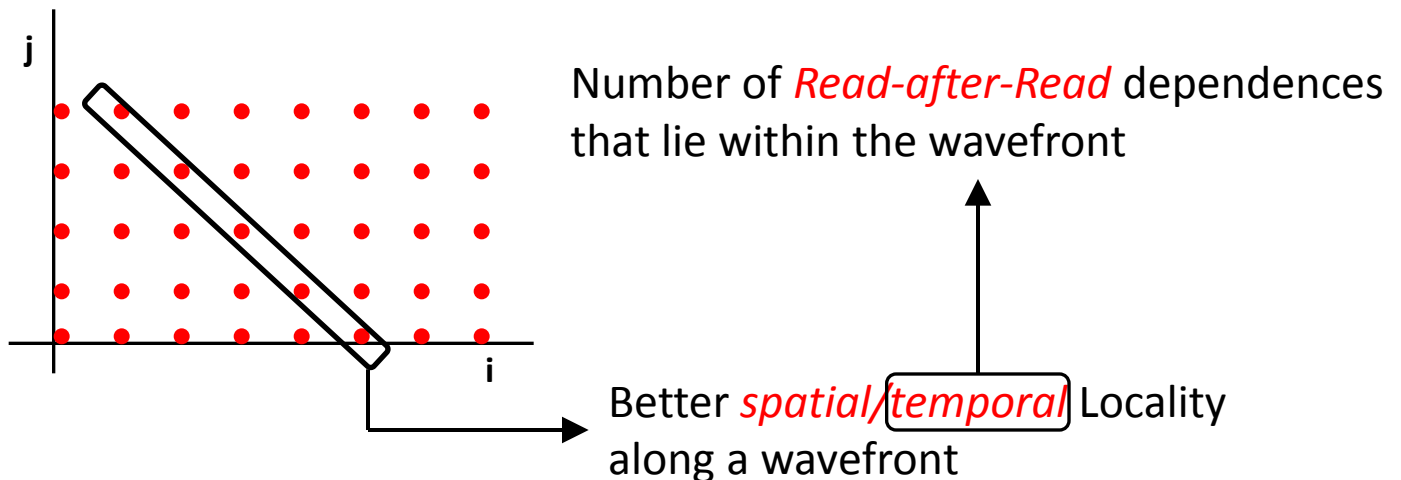
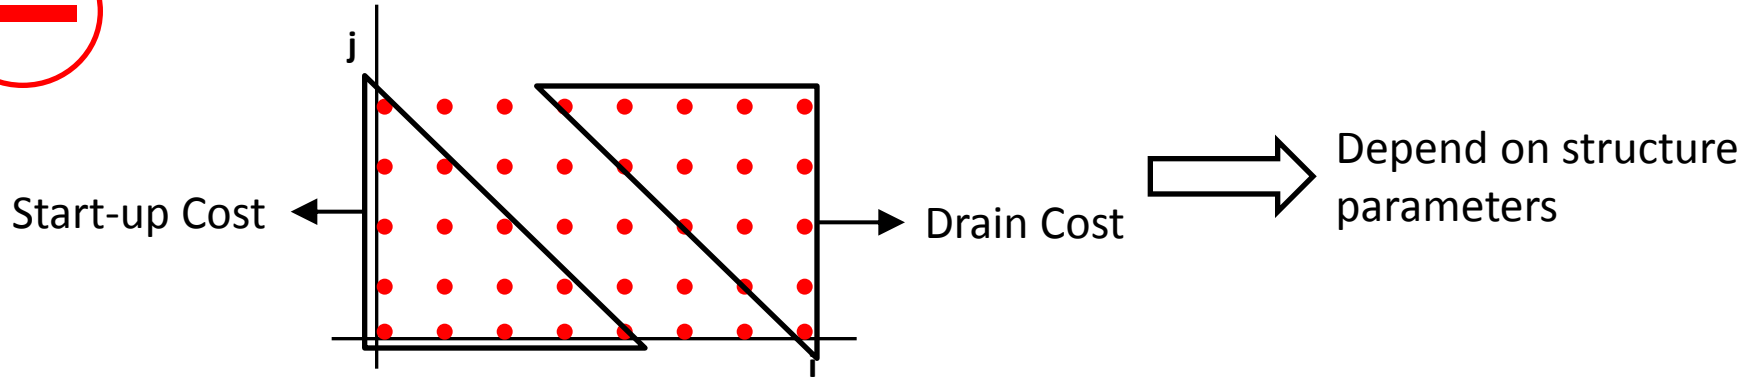
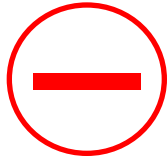
Pipeline Degrees of Parallelism



Pipeline Degrees of Parallelism



Pipeline Degrees of Parallelism



Fully Parallel vs Pipeline Degrees of Parallelism 1

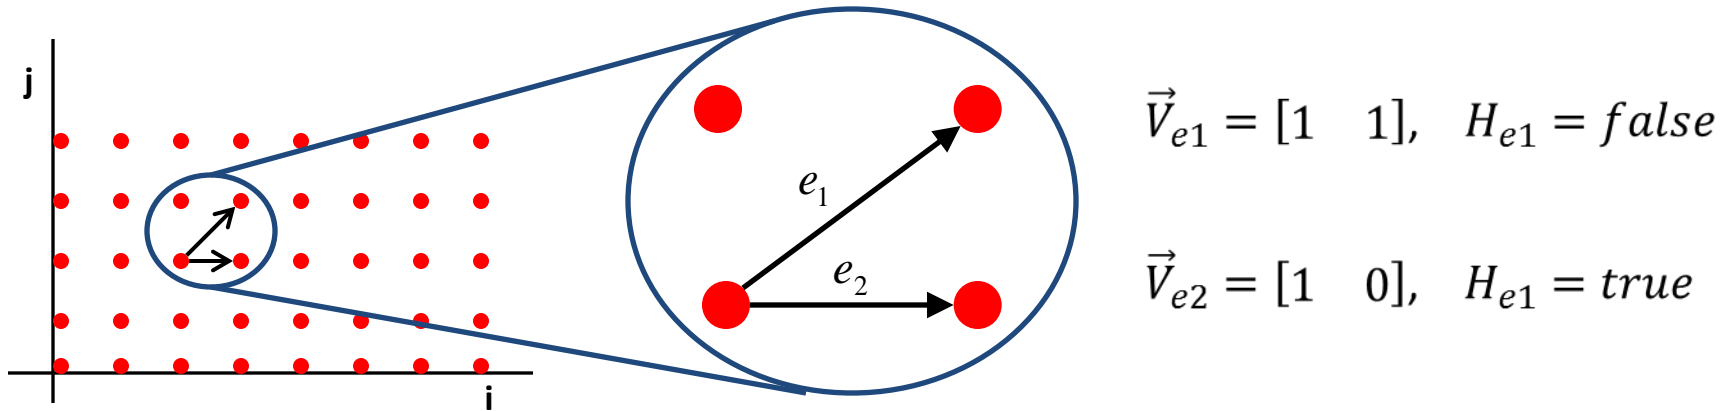
- We propose a way of distinguishing between **fully parallel** and **pipeline** degrees of parallelism.
- We use **dependence direction vectors** in order to expose **inner fully parallel** degrees of parallelism.

Direction Information :

$$\forall e \in E \xrightarrow[\substack{\text{bit vector} \\ 0 \leq i \leq \min(\dim_{dest}, \dim_{src})}]{\hspace{1.5cm}} \vec{V}_e[i] = \begin{cases} 1, & \text{If } e \text{ extends along } i \\ 0, & \text{If } e \text{ does not extend along } i \end{cases}$$

$$\forall e \in E \xrightarrow{\text{Boolean}} H_e = \begin{cases} true, & \text{If } e \text{ extends in } \mathbf{only\ 1} \text{ dimension} \\ false, & \text{If } e \text{ extends in } \mathbf{more\ than\ 1} \text{ dimensions} \end{cases}$$

Fully Parallel vs Pipeline Degrees of Parallelism 2

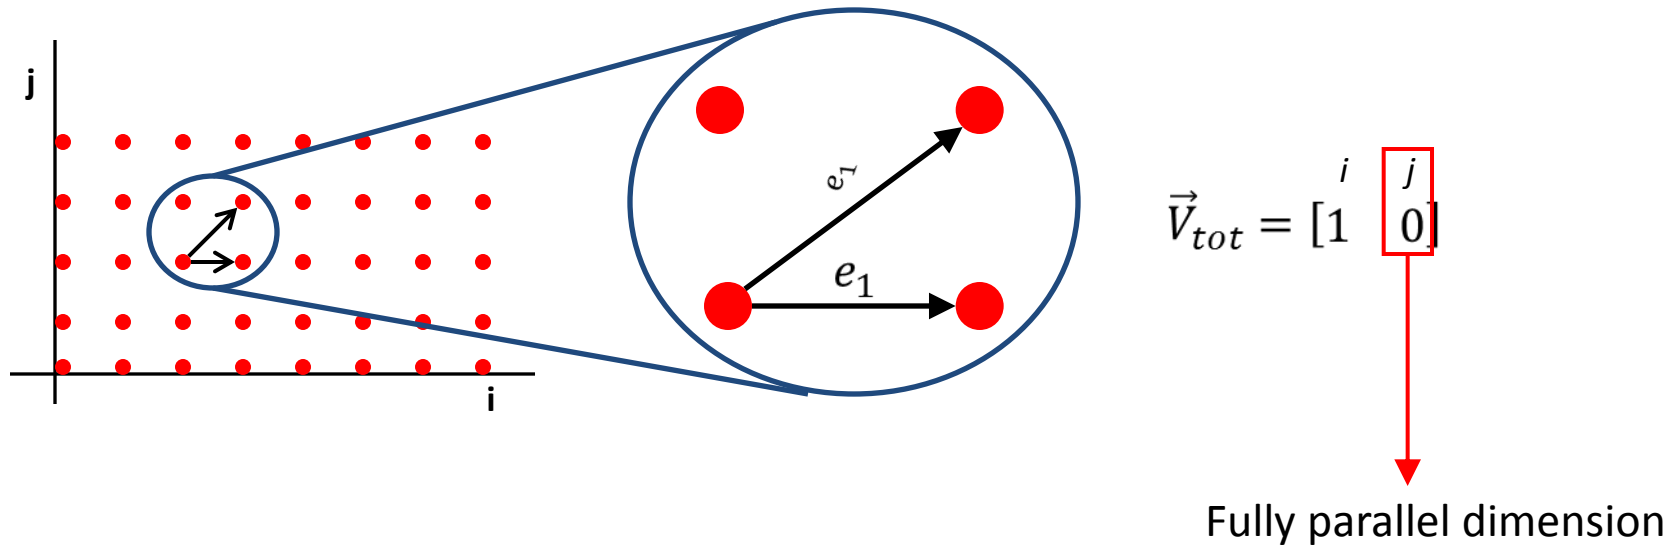


- The **disjunction** of all bit vectors that extend in **only 1** dimension ($H_e = \text{true}$) will **expose inner fully parallel loops**.

$$\vec{V}_{tot} = \bigvee_{\forall e \in E, H_e = \text{true}} \vec{V}_e$$

- If $\vec{V}_{tot}[\mathbf{i}] = \mathbf{0}$ then put a_i in **leading minimization position**.

Fully Parallel vs Pipeline Degrees of Parallelism 3



- By placing the coefficients of **fully parallel dimensions** in **leading minimization positions** we are effectively **pushing** them towards inner nest levels.
- As a result **fully parallel degrees of parallelism** can be **recovered**.

Conclusions

- The PLuTo scheduling algorithm iteratively finds **affine transformations** that **minimize communication**.
- For the **same minimum communication** the solution might be **sensitive to the ordering of the affine transformation coefficients** in the global constraint matrix.
- We might have to **choose** between **fully parallel** and **pipeline** degrees of parallelism.
- We **propose** a method for **distinguishing** between **fully parallel** and **pipeline** degrees of parallelism.
- We use **dependence direction information** in order to **expose inner fully parallel loops**.

Thank You !

Any Questions ?