

Verified Validation for Polyhedral Scheduling

Xuyang Li, Hongjin Liang, Xinyu Feng

Nanjing University

Background

- Human needs highly optimized compilation techniques for modern software.
- Nested loops are optimization targets due to its heavy numerical computation.
- Common loop optimizations:
 - Apply combination of loop fusion, distribution, interchange, skewing, reverse, tiling, ..., to improve memory locality and/or parallelizability.
 - May further apply vectorization (SIMD) and parallelization (openmp primitives).
 - Map it automatically to domain-specific hardware like GPU.

Background

- To do such loop transformations, we need to analyze dependences between instructions of different iterations.

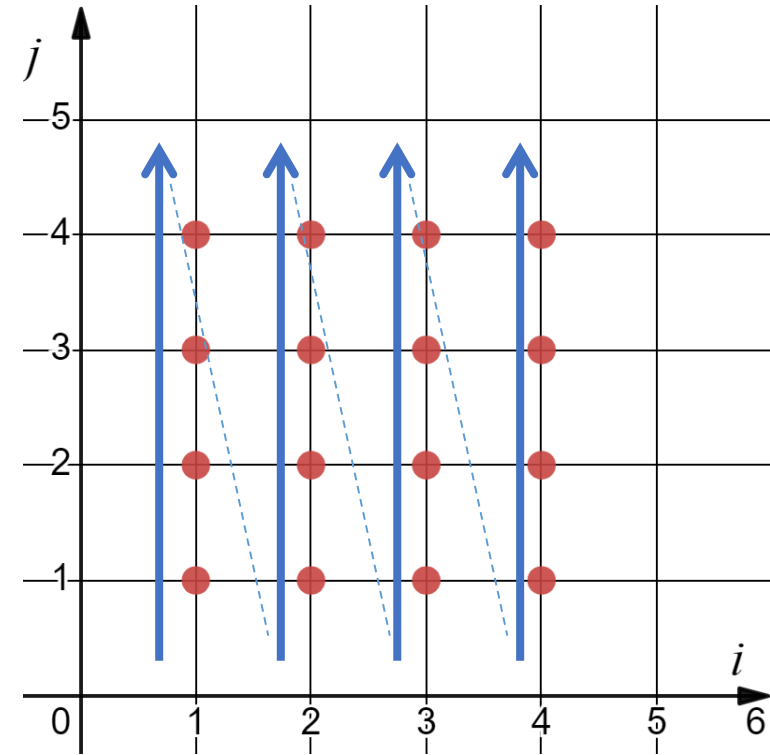
```
for i in [1, N]:  
    for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Can we parallelize it?

Example

```
for i in [1, N]:  
    for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```

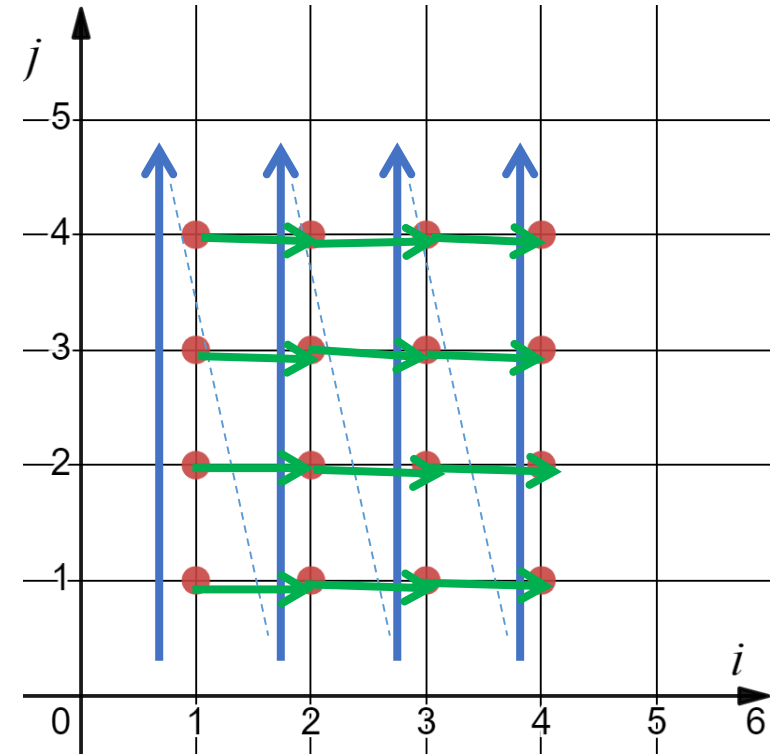


Iterations

Execution order

Example

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



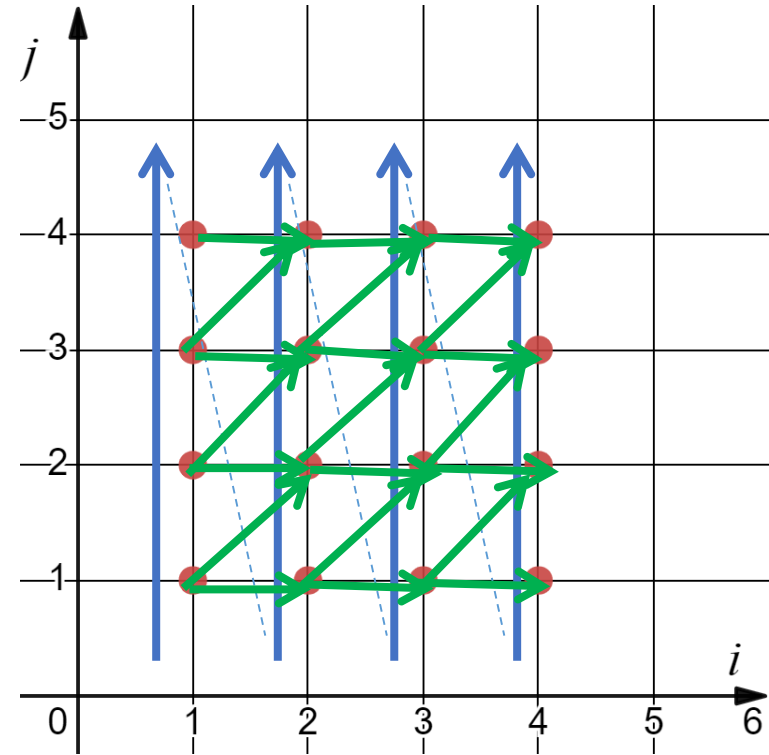
Iterations

Execution order

Dependence

Example

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



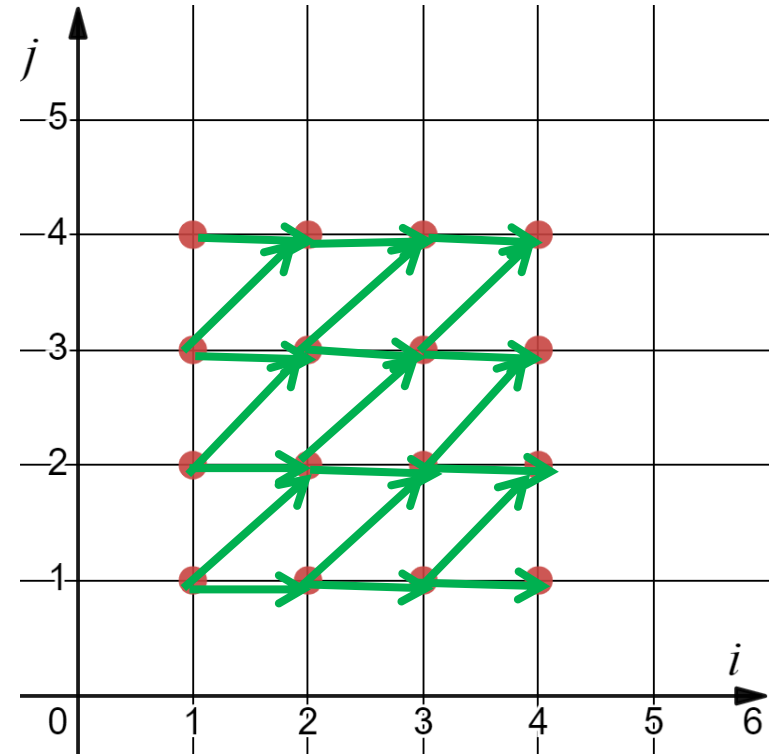
Iterations

Execution order

Dependence

Example

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Iterations

Execution order

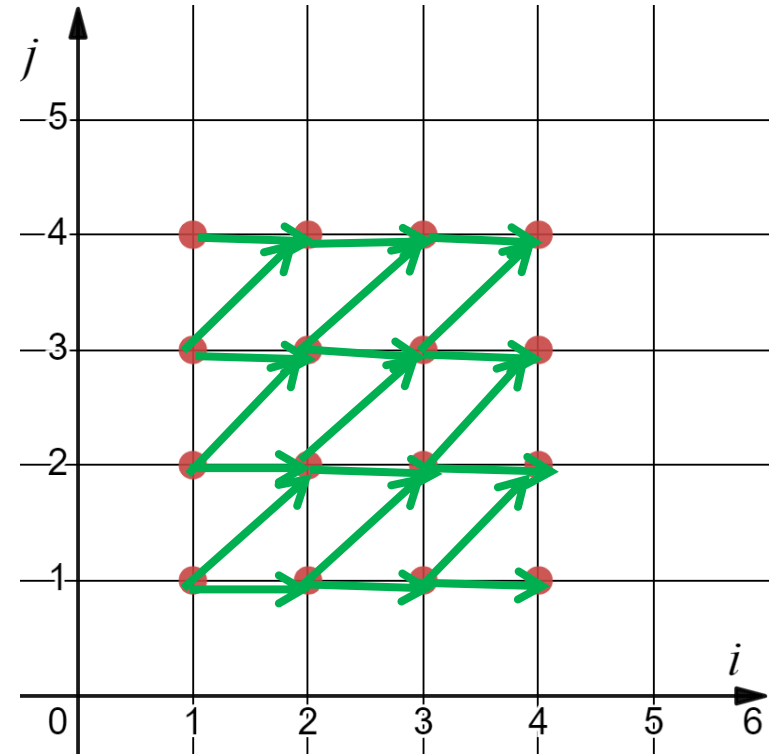
Dependence

Example

Can we parallelize it?

Respecting dependences?

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Iterations

Execution order

Dependence

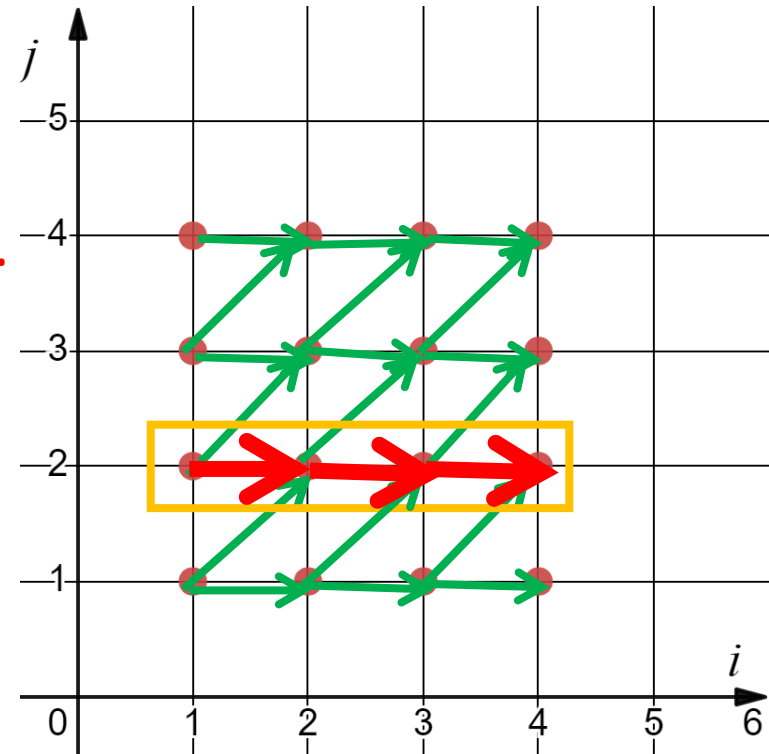
Example

Can we parallelize it?

Respecting dependences?

- Not along j . There is inner dependence.
- What about other axis?

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Iterations

Execution order

Dependence

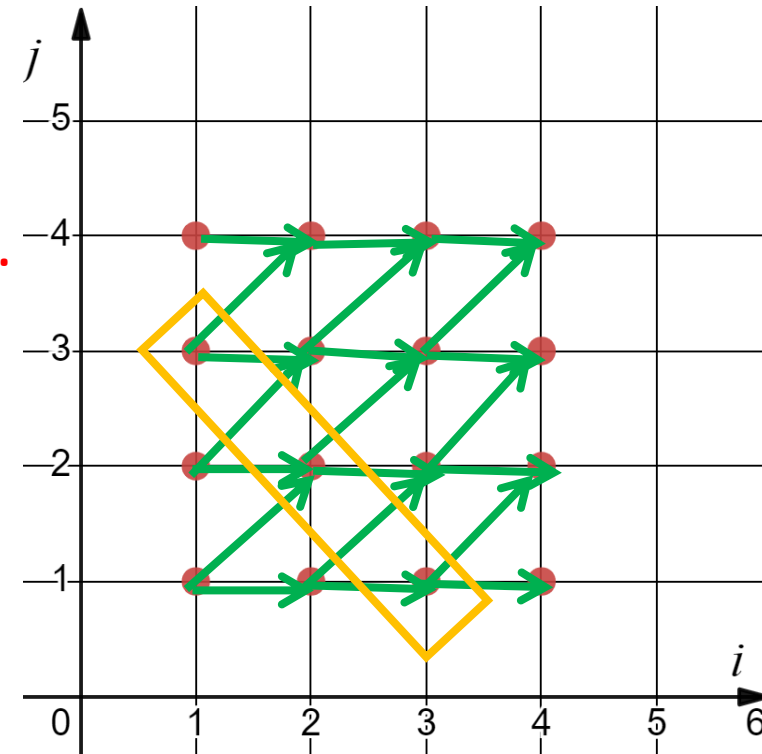
Example

Can we parallelize it?

Respecting dependences?

- Not along j . There is inner dependence.
- What about other axis?

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Iterations

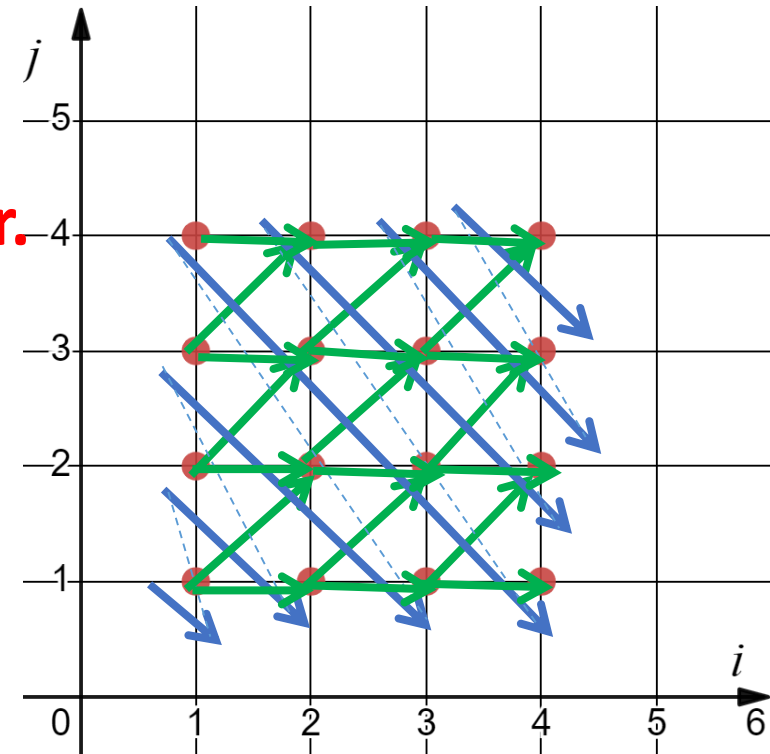
Execution order

Dependence

Example

Now we may find new execution order.
(Without breaking dependences)

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j] + A[i-1][j-1]
```



Iterations

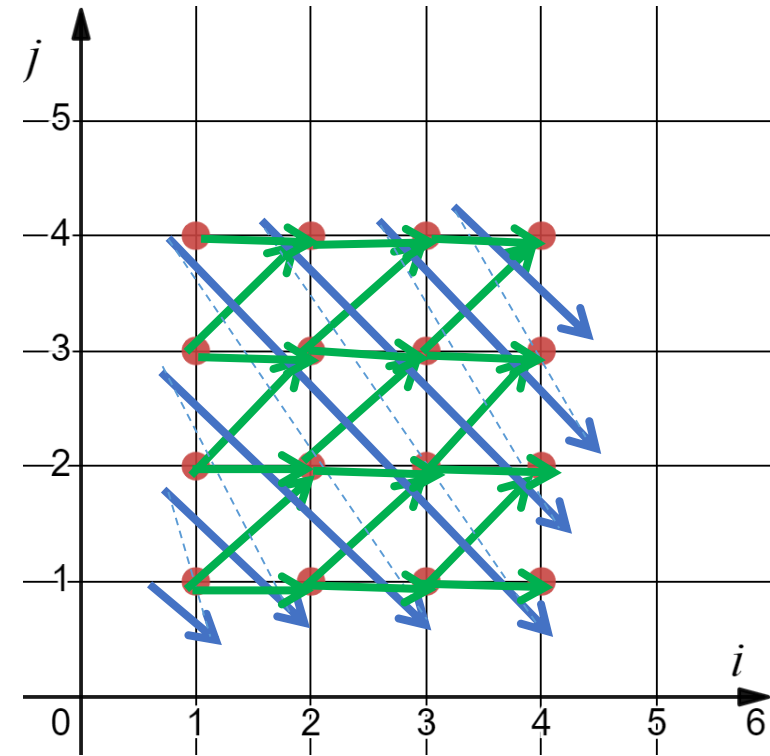
Execution order

Dependence

Example

Now we generate a new nested loop respecting new execution order.

```
for j' in [1, 2*N-1]:  
  for i' in [max(1, j'-N+1),  
            min(N-1, j'-1)]:  
S:  A[i'][(j'-i')] = A[i'+1][(j'-i')]  
      + A[i'-1][(j'-i')-1]
```

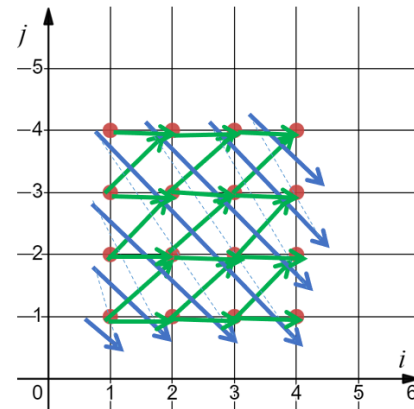
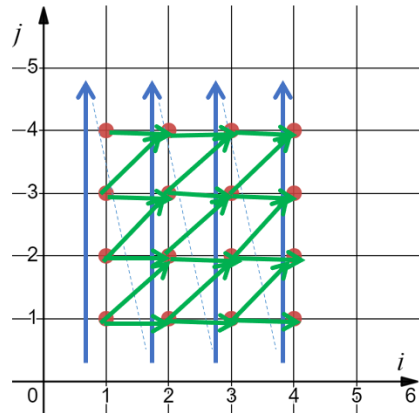


Iterations

Execution order

Dependence

Example

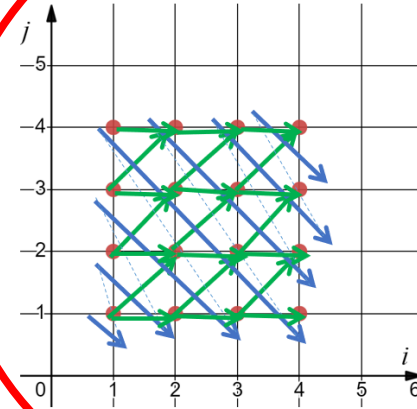
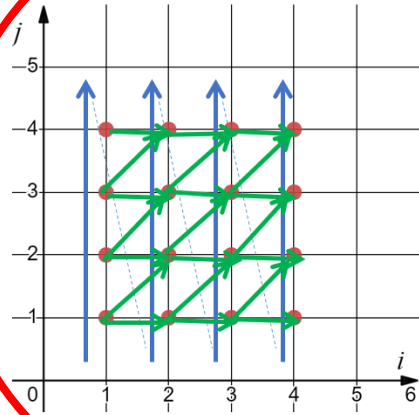


```
for i in [1, N]:  
  for j in [1, N]:  
    S: A[i][j] = A[i+1][j]  
        + A[i-1][j-1]
```

```
for j' in [1, 2*N-1]:  
  for i' in [max(1, j'-N+1),  
            min(N-1, j'-1)]:  
    S: A[i'][(j'-i')] = A[i'+1][(j'-i')]  
        + A[i'-1][(j'-i')-1]
```

loop skewing + loop interchange

Polyhedral Model



```
for i in [1, N]:  
  for j in [1, N]:  
    S: A[i][j] = A[i+1][j]  
        + A[i-1][j-1]
```

```
for j' in [1, 2*N-1]:  
  for i' in [max(1, j'-N+1),  
            min(N-1, j'-1)]:  
    S: A[i'][(j'-i')] = A[i'+1][(j'-i')]  
        + A[i'-1][(j'-i')-1]
```

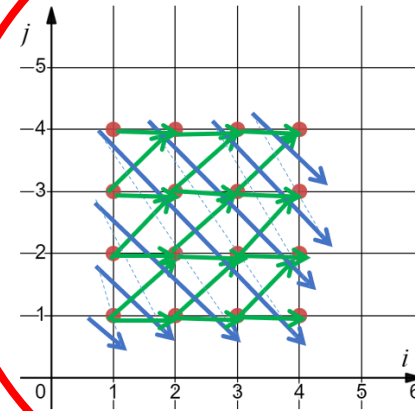
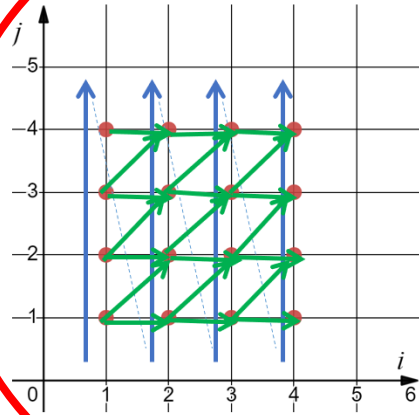
loop skewing + loop interchange

Polyhedral Model

Iterations – Domain – “The Polyhedron”

Execution order - Schedule

Dependence



```
for i in [1, N]:  
  for j in [1, N]:  
    S: A[i][j] = A[i+1][j]  
        + A[i-1][j-1]
```

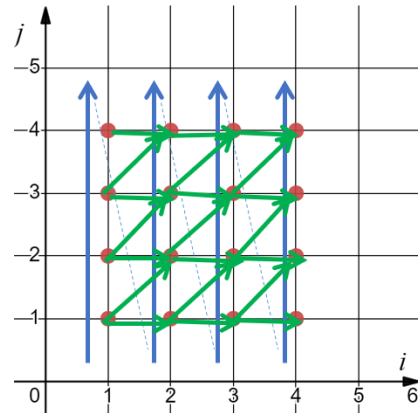
```
for j' in [1, 2*N-1]:  
  for i' in [max(1, j'-N+1),  
            min(N-1, j'-1)]:  
    S: A[i'][(j'-i')] = A[i'+1][(j'-i')]  
        + A[i'-1][(j'-i')-1]
```

loop skewing + loop interchange

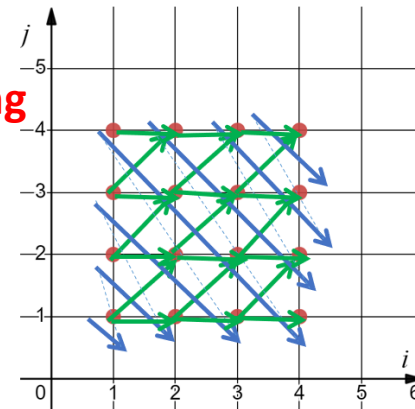
Polyhedral Compilation

extraction

```
for i in [1, N]:  
  for j in [1, N]:  
S:  A[i][j] = A[i+1][j]  
      + A[i-1][j-1]
```



scheduling



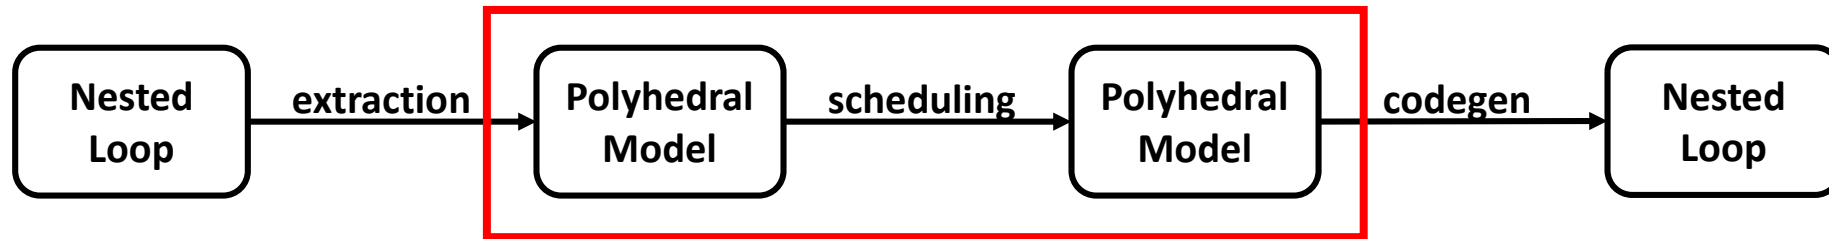
codegen

```
for j' in [1, 2*N-1]:  
  for i' in [max(1, j'-N+1),  
             min(N-1, j'-1)]:  
S:  A[i'][(j'-i')] = A[i'+1][(j'-i')]  
      + A[i'-1][(j'-i')-1]
```


Polyhedral Compilation



Polyhedral Compilation



Our focus!

Scheduling can be done manually or by automatic algorithms, like Pluto¹.

We want to ensure its correctness.

¹ <https://pluto-compiler.sourceforge.net/>

Compilation correctness

- For compiler $Comp$, programs \mathcal{P}_s and \mathcal{P}_t that $Comp(\mathcal{P}_s) = \text{Some } \mathcal{P}_t$.
- If \mathcal{P}_t **refines** \mathcal{P}_s (written as $\mathcal{P}_t \sqsubseteq \mathcal{P}_s$), we say this compilation is correct.
 - It says, from the same beginning state, whenever \mathcal{P}_t terminates at some state, then \mathcal{P}_s is able to stop at the same final state.

Compilation correctness

- For compiler $Comp$, programs \mathcal{P}_s and \mathcal{P}_t that $Comp(\mathcal{P}_s) = \text{Some } \mathcal{P}_t$.
- If \mathcal{P}_t **refines** \mathcal{P}_s (written as $\mathcal{P}_t \sqsubseteq \mathcal{P}_s$), we say this compilation is correct.
 - **It says, from the same beginning state, whenever \mathcal{P}_t terminates at some state, then \mathcal{P}_s is able to stop at the same final state.**

- Two ways to guarantee correct compilation:

- Compiler proof: reasoning on $Comp$'s concrete definition to prove

$$\forall \mathcal{P}_s, \mathcal{P}_t. Comp(\mathcal{P}_s) = \text{Some } \mathcal{P}_t \implies \mathcal{P}_t \sqsubseteq \mathcal{P}_s.$$

- Verified validation: define a separate validator $Validate$ and prove

$$\forall \mathcal{P}_s, \mathcal{P}_t. Validate(\mathcal{P}_s, \mathcal{P}_t) = \text{true} \implies \mathcal{P}_t \sqsubseteq \mathcal{P}_s.$$

- And run $Validate$ after **each run** of $Comp$.

Compilation correctness

Why not directly verify the scheduling algorithm?

- On the one hand, it contains complex heuristic with heavy mathematics. Hard/Impractical to verify.
- On the other hand, it has simple validation algorithm due to its simple correctness criterion: not breaking dependence. (Also called *Bernstein's Condition*².)



- Two ways to guarantee correct compilation:
 - Compiler proof: reasoning on *Comp*'s concrete definition to prove

$$\forall \mathcal{P}_s, \mathcal{P}_t. \text{Comp}(\mathcal{P}_s) = \text{Some } \mathcal{P}_t \implies \mathcal{P}_t \sqsubseteq \mathcal{P}_s.$$

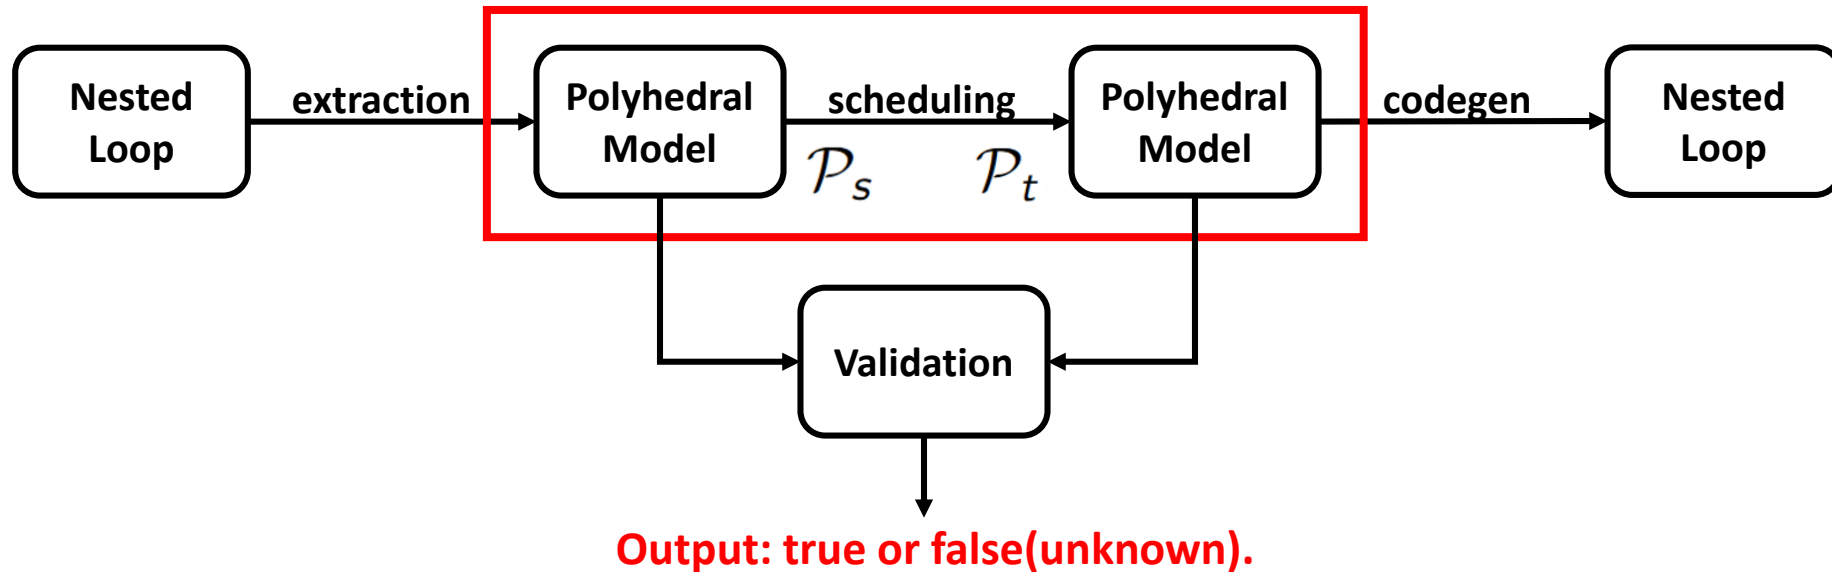
- Verified validation: define a separate validator *Validate* and prove

$$\forall \mathcal{P}_s, \mathcal{P}_t. \text{Validate}(\mathcal{P}_s, \mathcal{P}_t) = \text{true} \implies \mathcal{P}_t \sqsubseteq \mathcal{P}_s.$$

- And run *Validate* after each run of *Comp*.

² https://link.springer.com/referenceworkentry/10.1007/978-0-387-09766-4_521

Validating Polyhedral Scheduling



Our work:

Verified Validation for Polyhedral Scheduling

- We **implement and verify** a general validation function *Validate* for polyhedral scheduling in Coq proof assistant from scratch. It is parameterized by instruction language.
- We instantiate *Validate* with a variant of **CompCert**'s instruction language, showing its practicality.
- We adapt *Validate* so that it works with the **Pluto** compiler. It successfully validates all its available test cases with reasonable overhead.

General Verified Validation function for Polyhedral Scheduling

- We define a validation function *Validate* that checks **Bernstein's condition** within polyhedral model, and mechanized its correctness with the Coq proof assistant.
- It is parametrized by instruction language to be reusable.
- Our top lemma

Definition (correctness of the validator)

$$\begin{aligned} \text{Correct}(\text{Validate}) &\triangleq \forall \mathcal{P}_s, \mathcal{P}_t. \text{Validate}(\mathcal{P}_s, \mathcal{P}_t) = \text{true} \\ &\implies \mathcal{P}_t \sqsubseteq \mathcal{P}_s. \\ \mathcal{P}_t \sqsubseteq \mathcal{P}_s &\triangleq \forall \sigma, \sigma'. \\ &\quad \models \mathcal{P}_t, \sigma \Rightarrow \sigma' \implies \models \mathcal{P}_s, \sigma \Rightarrow \sigma'. \end{aligned}$$

Case study on CompCert



- What is CompCert³?
 - a high-assurance compiler for almost all of the C language (ISO C 2011), generating efficient code for the ARM, PowerPC, RISC-V and x86 processors.
 - It is not enough optimized⁴ than production compiler like clang and gcc. **So aggressive optimizers like polyhedral-based ones do help!**
- We successfully instantiate *Validate* with CompCert's semantics model, showing the possibility towards a fully verified polyhedral extension to CompCert.

³ <https://compcert.org/>

⁴ <https://doi.org/10.1145/3622799>

Case study on Pluto compiler

- Loop optimizers like polyhedral-based ones are error prone⁵ indeed!
So formal methods do help.
- We evaluate on Pluto, one of the famous polyhedral compiler.
 - Pluto: ACM SIGPLAN **PLDI** Most Influential Paper **award** in 2018

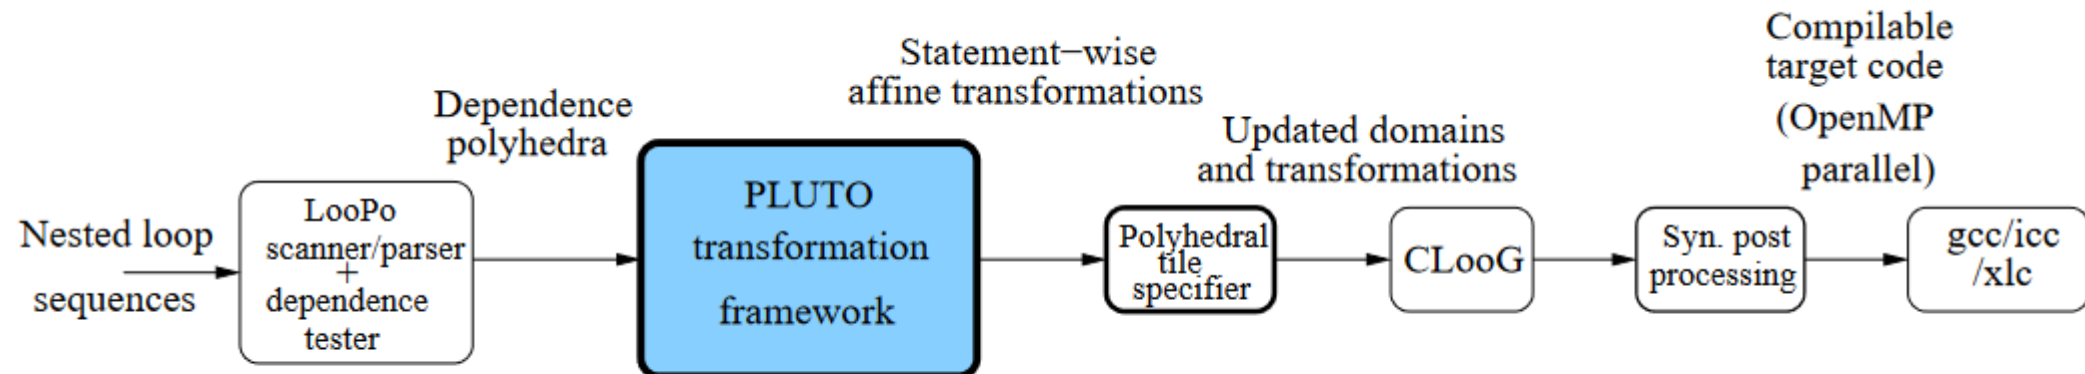
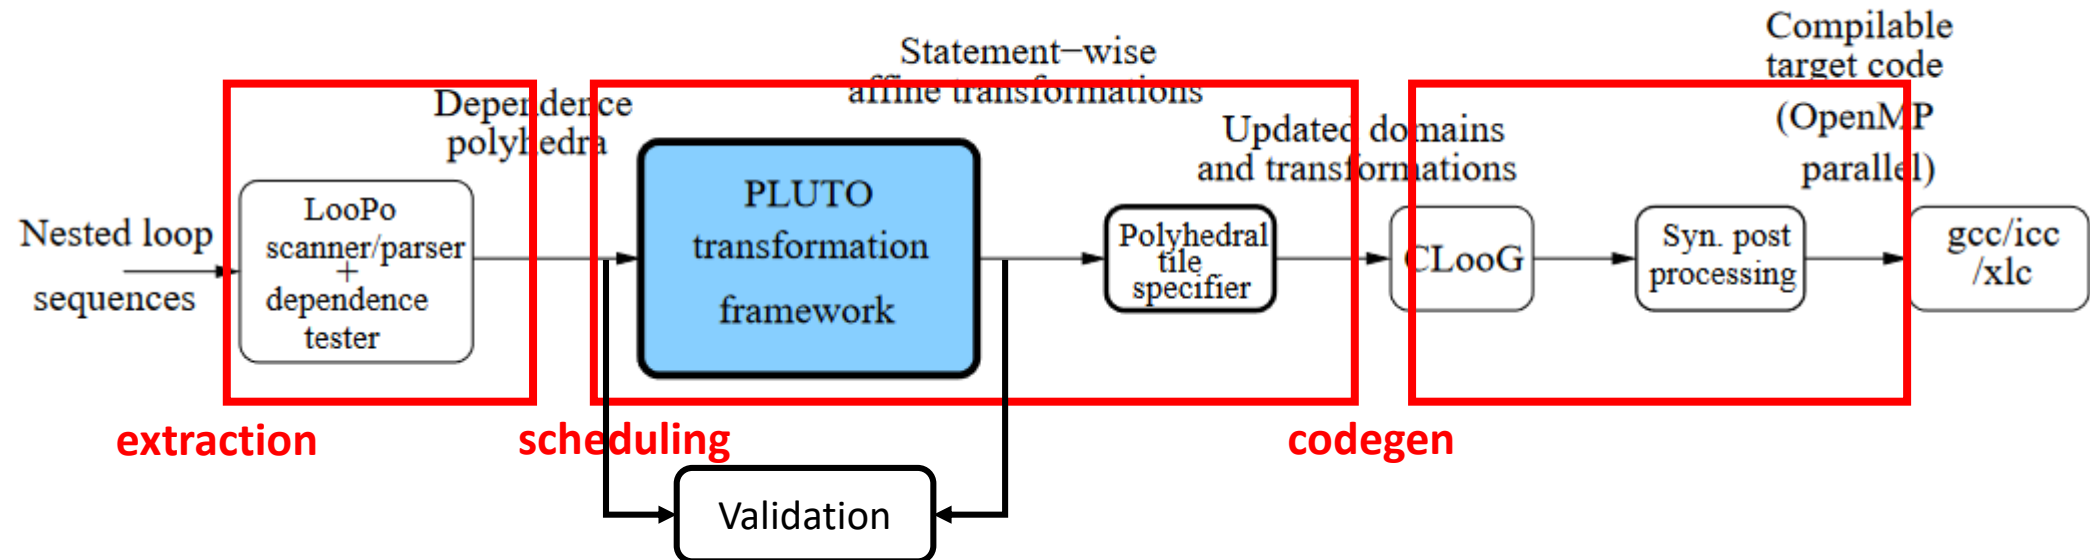


Figure from <https://www.csa.iisc.ac.in/~udayb/publications/uday-thesis.pdf>, Page 98

⁵ <https://doi.org/10.1145/3591295>

Case study on Pluto compiler

- Loop optimizers like polyhedral-based ones are error prone⁵ indeed!
So formal methods do help.
- We evaluate on Pluto, one of the famous polyhedral compiler.
 - Pluto: ACM SIGPLAN **PLDI** Most Influential Paper **award** in 2018



⁵ <https://doi.org/10.1145/3591295>

Case study on Pluto compiler

- Loop optimizers like polyhedral-based ones are error prone⁵ indeed!
So formal methods do help.
- We evaluate on Pluto, one of the famous polyhedral compiler.
 - Pluto: ACM SIGPLAN **PLDI** Most Influential Paper **award** in 2018
- Work and work well with **Pluto compiler**: successfully validate its all 62 available test cases.
 - Not just a research prototype: work as an executable on practical polyhedral compiler.
 - Soundness v.s. completeness:
 - Formal verification proves the algorithm **soundness**: guarantee **refinement** whenever it outputs **true**.
 - The evaluation shows validation's algorithm **completeness**: if refinement establishes on inputs, *the algorithm* tries its best to output **true** rather than **unknown!**
 - Also, the algorithm has reasonable overhead.

⁵ <https://doi.org/10.1145/3591295>

Case st

- Loop opti
- So formal
- We evalu
 - Pluto: A
- Work and available
 - Not jus compile
 - Soundr
 - For out
 - The *the*
 - Also, th

| Test | Time of Pluto (ms) | Time of Validation (ms,ms) | Result |
|-----------------|--------------------|----------------------------|--------|
| covcol | 3.5 | 434.6, 320.7 | EQ |
| dsyr2k | 2.6 | 106.0, 83.4 | EQ |
| fdtd-2d | 46.4 | 1615.5, 1296.3 | EQ |
| gemver | 7.0 | 247.9, 240.4 | EQ |
| lu | 6.1 | 410.6, 331.2 | EQ |
| mvt | 2.2 | 70.2, 56.3 | EQ |
| ssymm | 40.7 | 726.0, 551.2 | EQ |
| tce | 568.6 | 4442.0, 4422.5 | EQ |
| adi | 77.5 | 2531.7, 2377.8 | EQ |
| corcol | 5.5 | 442.5, 362.1 | EQ |
| dct | 21.8 | 879.4, 739.4 | EQ |
| dsyrk | 1.8 | 96.8, 78.9 | EQ |
| floyd | 12.1 | 502.6, 421.7 | EQ |
| jacobi-1d-imper | 3.8 | 184.0, 167.8 | EQ |
| matmul-init | 2.9 | 257.8, 192.4 | EQ |
| pca | 202.5 | 2923.6, 2679.5 | EQ |
| strmm | 1.9 | 141.4, 110.8 | EQ |
| tmm | 1.6 | 109.7, 89.6 | EQ |
| advect3d | 1023.1 | 579.1, 498.1 | EQ |
| corcol3 | 13.6 | 851.3, 733.4 | EQ |
| doitgen | 10.4 | 1069.2, 837.4 | EQ |
| fdtd-1d | 6.0 | 268.7, 229.9 | EQ |
| jacobi-2d-imper | 17.7 | 619.5, 543.5 | EQ |
| matmul | 3.2 | 157.1, 125.5 | EQ |
| seidel | 24.5 | 818.1, 725.5 | EQ |
| strsm | 6.4 | 209.3, 161.2 | EQ |
| trisolv | 5.1 | 338.9, 248.8 | EQ |
| 1dloop-invar | 0.3 | 6.7, 6.0 | EQ |
| costfunc | 0.8 | 47.4, 35.0 | EQ |
| fusion1 | 0.9 | 15.3, 13.9 | EQ |
| ... | ... | ... | ... |

ed!

all 62

dral

enever it

ishes on inputs,

Table 1: Evaluation results on Pluto's test suits

Coq development

- Around 18000 lines of Coq, 1000 lines of OCaml. Open source at <https://github.com/verif-scop/PolCert>.
- Our work bases on Verified Polyhedron Library⁶ and is syntactically compatible to PolyGen⁷ (Verified Polyhedral Code Generation) in POPL'21.

⁶ <https://ieeexplore.ieee.org/document/8750763>

⁷ <https://dl.acm.org/doi/10.1145/3434321>

Future work

- Complete verified polyhedral compilation.
 - Verified extractor.
 - Engineering in CompCert's driver & frontend.
 - Apply optimistic approach⁸ to deal with polyhedral model's heavy assumptions, like integer overflow⁹.
- Support validation for other polyhedral transformations, like index set split (as a pre-phrase), tiling (as a post-phrase), layout transformation (as an orthogonal phrase).
- Support vectorization, parallelization, GPU compilation ...

⁸ <https://dl.acm.org/doi/10.5555/3049832.3049864>

⁹ <https://inria.hal.science/hal-00655485>

Q&A

Our work:

Verified Validation for Polyhedral Scheduling

- We **implement and verify** a general validation function *Validate* for polyhedral scheduling in Coq proof assistant from scratch. It is parameterized by instruction language.
- We instantiate *Validate* with a variant of **CompCert**'s instruction language, showing its practicality.
- We adapt *Validate* so that it works with the **Pluto** compiler. It successfully validates all its available test cases with reasonable overhead.

Original code π_{cov} for covariance matrix calculation, 1.84s

```
1 for (j1 = 1; j1 <= M; j1++) {
2   for (j2 = j1; j2 <= M; j2++) {
3     for (i = 1; i <= N; i++) {
4       I0: symmat[j1][j2] += data[i][j1] * data[i][j2];
5     }
6     I1: symmat[j2][j1] = symmat[j1][j2];
7   }
8 }
```

Optimized code π'_{cov} for covariance matrix calculation, 0.43s, with loop distribution and loop interchange

```
1 for (i = 1; i <= N; i++) {
2   for (j1 = 1; j1 <= M; j1++) {
3     for (j2 = j1; j2 <= M; j2++) {
4       symmat[j1][j2] += data[i][j1] * data[i][j2];
5     }
6   }
7 }
8 for (j1 = 1; j1 <= M; j1++) {
9   for (j2 = j1; j2 <= M; j2++) {
10    symmat[j2][j1] = symmat[j1][j2];
11  }
12 }
```

See at <https://github.com/verif-scop/speed-up>.

Polyhedral compilation does high-level structural transformations and only impose a few properties of the underlying instruction language (called \mathbb{I}). The validation function given in this work is parameterized by \mathbb{I} .

\mathbb{I} allows user define the syntax, types, state, semantics of the language, how it initializes, and its the non-alias proposition. It demands user to provide a verified *Checker* function to validate the consistency between the read and write access function and an instruction's semantics, and prove that any two instances that satisfy Bernstein's conditions are permutable.

We assume Pluto's instruction language satisfy this abstraction.

Module Type $\mathbb{II} \triangleq$

$\mathbb{T}, \mathbb{I}, \mathbb{S} : \text{Type}$

$\models : \mathbb{I} \rightarrow \text{List}(\mathbb{Z}) \rightarrow \text{Memory Cells}$
 $\rightarrow \text{Memory Cells} \rightarrow \mathbb{S} \rightarrow \mathbb{S} \rightarrow \text{Prop}$

$\text{Compat} : \text{List}(\text{Identifier}) \rightarrow \mathbb{S} \rightarrow \text{Prop}$

$\text{Consistent} : \text{List}(\text{Identifier} \times \mathbb{T}) \rightarrow \text{List}(\mathbb{Z}) \rightarrow \mathbb{S} \rightarrow \text{Prop}$

$\text{NonAlias} : \mathbb{S} \rightarrow \text{Prop}$

$\text{NonAliasPsr} :$

$\forall \mathbb{I}, \sigma, \sigma'. \text{NonAlias}(\sigma) \wedge p \models \mathbb{I}, \sigma \xrightarrow{\cdot} \sigma' \implies \text{NonAlias}(\sigma').$

$\text{Checker} : \mathbb{I} \rightarrow \text{Access Functions}$

$\rightarrow \text{Access Functions} \rightarrow \text{Bool}$

$\text{Correct}(\text{Checker}) :$

$\forall \mathbb{I}, \mathcal{W}, \mathcal{R}. \text{Checker}(\mathbb{I}, \mathcal{W}, \mathcal{R}) = \text{true}$

$\implies (\forall \sigma, \sigma', p, \Delta_r, \Delta_w. p \models \mathbb{I}, \sigma \xrightarrow{\Delta_r, \Delta_w} \sigma' \implies \Delta_r \subseteq \mathcal{R}(p) \wedge \Delta_w \subseteq \mathcal{W}(p)).$

$\text{BCPermut} :$

$\forall \mathbb{I}_1, \mathbb{I}_2, p_1, p_2, \sigma, \sigma', \sigma'', \Delta_r, \Delta_w, \Delta'_r, \Delta'_w.$

$(p_1 \models \mathbb{I}_1, \sigma \xrightarrow{\Delta_r, \Delta_w} \sigma' \wedge p_2 \models \mathbb{I}_2, \sigma' \xrightarrow{\Delta'_r, \Delta'_w} \sigma''$
 $\wedge \Delta_r \cap \Delta'_w = \emptyset \wedge \Delta_w \cap \Delta'_r = \emptyset \wedge \Delta_w \cap \Delta'_w = \emptyset)$

$\implies \exists \sigma^*. p_2 \models \mathbb{I}_2, \sigma \xrightarrow{\Delta'_r, \Delta'_w} \sigma^* \wedge p_1 \models \mathbb{I}_1, \sigma^* \xrightarrow{\Delta_r, \Delta_w} \sigma''.$

Figure 1: Definition of Instruction Language Module \mathbb{II}

Case study: Towards verified polyhedral compilation for CompCert

We instantiate the validation function with CompCert C's type, state and subset of its instruction language, and implement *Checker* with symbolic execution. All verified. Only differences are, affine expression is evaluated in \mathbb{Z} (no overflow), and multi-dimensional array access is sugared.

| | | | |
|-------------------------|----------------|-------|---|
| (Base Type) | τ_{\perp} | ::= | int32s |
| (Type) | τ | \in | Base Type \times List(\mathbb{Z}) |
| (Value) | v | ::= | I32(n) ... |
| (Iterator) | i | \in | \mathbb{N} |
| (Unop) | op_1 | ::= | - ... |
| (Binop) | op_2 | ::= | + * ... |
| (May Affine Expression) | ε | ::= | z i $op_1 \varepsilon$ $\varepsilon_1 op_2 \varepsilon_2$ |
| (Access Expression) | ϵ | \in | Identifier \times Base Type \times List(May Affine Expression) |
| (Expression) | e | ::= | v i ϵ $op_1 e$ $e_1 op_2 e_2$ |
| (Base Instruction) | I | ::= | skip $\epsilon := e$ |

Polyhedral Compilation

- Extraction -> Scheduling -> Codegen

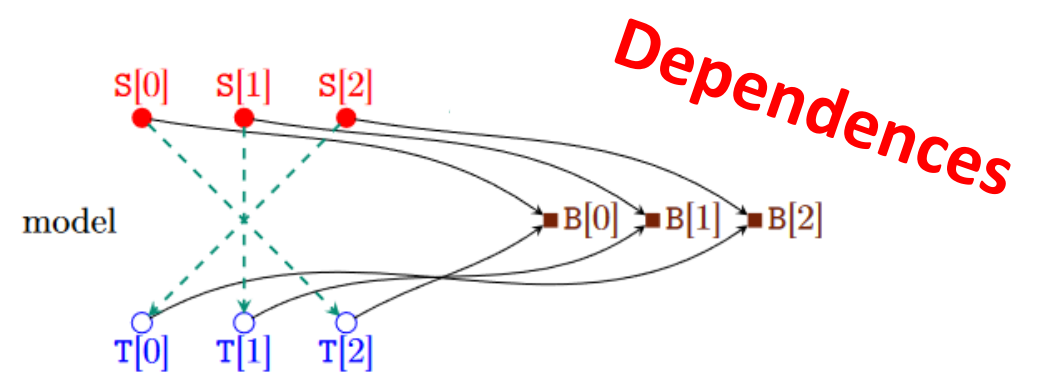
loop reverse + loop fusion

~>

```
for i in [0, 3]:  
S:   B[i] = A[i] + 1  
for i in [0, 3]:  
T:   C[i] = B[2-i] * 3
```

```
for i in [0, 3]:  
S:   B[i] = A[i] + 1  
T:   C[2-i] = B[i] * 3
```

Polyhedral Model



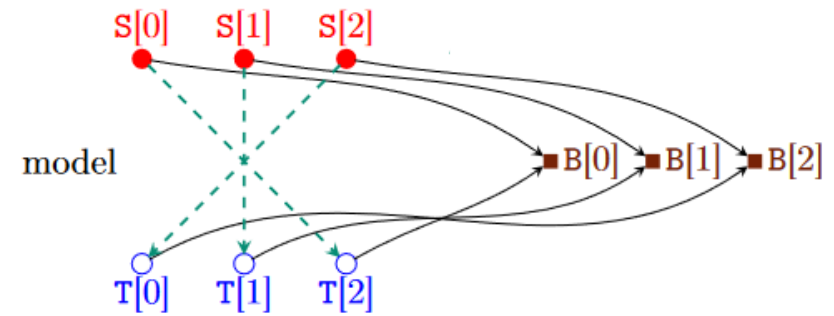
- An alternative integer-programming-based representation for nested loop. A **polyhedral program** is a multiset of **polyhedral instructions**, which consists of at least three parts: the **instruction**, the **domain**, and the **schedule**.
- “Extraction”

```
for i in [0, 3]:           ~>
S:   B[i] = A[i] + 1      { (S, 0<=i<3, [i] ~> [0, i]),
for i in [0, 3]:          (T, 0<=i<3, [i] ~> [1, i]) }
T:   C[i] = B[2-i] * 3
```

*Lexicographically
ordered!*

Execution order: $S[0]; S[1]; S[2]; T[0]; T[1]; T[2]$

Polyhedral Model



Our work's focus!

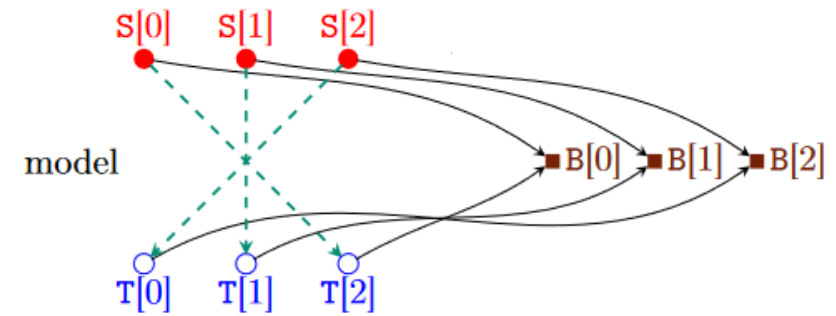
- “Scheduling”: do reordering transformation on polyhedral model, guided by dependence analysis.
- Well-designed automatic algorithms, like Pluto², serve for this purpose.
- Manual scheduling is also possible.
- Recall the correctness criterion: Bernstein’s condition.

$$\{ (S, 0 \leq i < 3, [i] \sim > [0, i]), \sim > \{ (S, 0 \leq i < 3, [i] \sim > \dots), \\ (T, 0 \leq i < 3, [i] \sim > [1, i]) \} \quad (T, 0 \leq i < 3, [i] \sim > \dots) \}$$

Old Execution order: $S[0]; S[1]; S[2]; T[0]; T[1]; T[2]$

² <https://pluto-compiler.sourceforge.net/>

Polyhedral Model



Our work's focus!

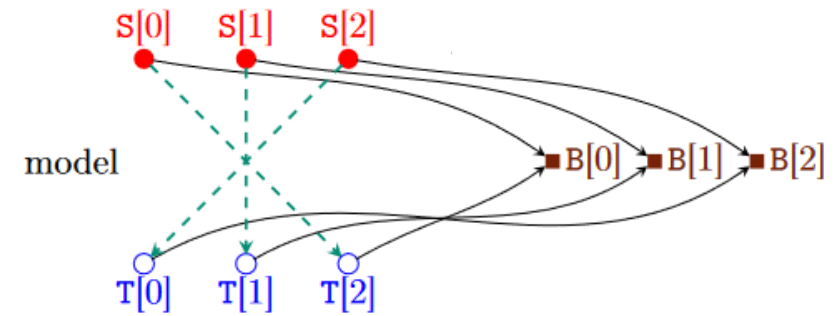
- “Scheduling”: do reordering transformation on polyhedral model, guided by dependence analysis.
- Well-designed automatic algorithms, like Pluto², serve for this purpose.
- Manual scheduling is also possible.

**Average distance of
dependence's satisfaction
seems too long. Cache locality.**

Old Execution order: S [0] ; S [1] ; S [2] ; T [0] ; T [1] ; T [2]

² <https://pluto-compiler.sourceforge.net/>

Polyhedral Model



Our work's focus!

- “Scheduling”: do reordering transformation on polyhedral model, guided by dependence analysis.
- Well-designed automatic algorithms, like Pluto², serve for this purpose.
- Manual scheduling is also possible.

Average distance of dependence's satisfaction seems to long.

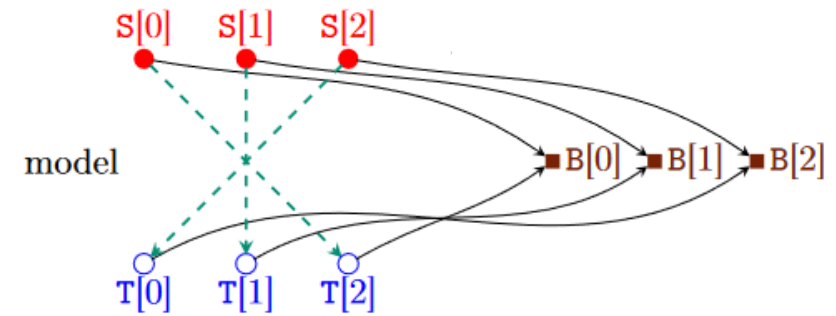
Try to minimize the distance without breaking dependences!

Old Execution order: `S [0] ; S [1] ; S [2] ; T [0] ; T [1] ; T [2]`

New execution order: `S [0] ; T [2] ; S [1] ; T [1] ; S [2] ; T [0]`

² <https://pluto-compiler.sourceforge.net/>

Polyhedral Model



Our work's focus!

- “Scheduling”: do reordering transformation on polyhedral model, guided by dependence analysis.
- Well-designed automatic algorithms, like Pluto², serve for this purpose.
- Manual scheduling is also possible.
- Recall the correctness criterion: Bernstein’s condition.

$$\{ (S, 0 \leq i < 3, [i] \sim > [0, i]), (T, 0 \leq i < 3, [i] \sim > [1, i]) \} \sim > \{ (S, 0 \leq i < 3, [i] \sim > [i, 0]), (T, 0 \leq i < 3, [i] \sim > [2-i, 1]) \}$$

Old Execution order: S [0] ; S [1] ; S [2] ; T [0] ; T [1] ; T [2]

New execution order: S [0] ; T [2] ; S [1] ; T [1] ; S [2] ; T [0]

² <https://pluto-compiler.sourceforge.net/>

Polyhedral Model

- “Codegen”: recover the imperative control structure from the optimized polyhedral model.

```
{ (S, 0<=i<3, [i] ~> [i, 0]),  
  (T, 0<=i<3, [i] ~> [2-i, 1]) } ~>  
for i in [0, 3]:  
S:   B[i] = A[i] + 1  
T:   C[2-i] = B[i] * 3
```