

# Faster MPSoC Task Mapping via Symmetry Reduction

Timo Nicolai

`timo.nicolai@mailbox.tu-dresden.de`

June 9, 2020

## 1. Inspirations

# Contents

1. Inspirations
2. Problem Statement

# Contents

1. Inspirations
2. Problem Statement
3. Symmetry Reduction

# Contents

1. Inspirations
2. Problem Statement
3. Symmetry Reduction
4. Representing and Extracting Symmetries

# Contents

1. Inspirations
2. Problem Statement
3. Symmetry Reduction
4. Representing and Extracting Symmetries
5. The TMOR Problem

# Contents

1. Inspirations
2. Problem Statement
3. Symmetry Reduction
4. Representing and Extracting Symmetries
5. The TMOR Problem
6. Experimental Results

# Contents

1. Inspirations
2. Problem Statement
3. Symmetry Reduction
4. Representing and Extracting Symmetries
5. The TMOR Problem
6. Experimental Results
7. Further Research




- Original idea: [Goens et al., 2017]


- Original idea: [Goens et al., 2017]
- Theoretical foundations: [Holt, 2005] and [East et al., 2019]

- Original idea: [Goens et al., 2017]
- Theoretical foundations: [Holt, 2005] and [East et al., 2019]
- Important optimizations: [Donaldson and Miller, 2009]

# Problem Statement




# Problem Statement




- Want to intelligently map tasks to processing elements

# Problem Statement



- Want to intelligently map tasks to processing elements
- Best choice depends on underlying optimality criteria

# Problem Statement



- Want to intelligently map tasks to processing elements
- Best choice depends on underlying optimality criteria
- Need to perform costly simulation!

# Symmetry Reduction

- One approach:
  - Generate promising mappings based on previous simulations
  - → Traverse search space “intelligently”




# Symmetry Reduction

- One approach:
  - Generate promising mappings based on previous simulations
  - → Traverse search space “intelligently”
- Another approach:
  - Partition search space *by (partial) symmetry*
  - Only work with representative of each partition
  - → “Collapse” search space


# Symmetry Reduction

- One approach:
  - Generate promising mappings based on previous simulations
  - → Traverse search space “intelligently”
- Another approach:
  - Partition search space *by (partial) symmetry*
  - Only work with representative of each partition
  - → “Collapse” search space
- Both approaches can be combined!


# Symmetry Reduction



(a)




(b)




(c)


# Symmetry Reduction



(a)




(b)




(c)


# Representing Symmetries: Automorphism Groups




# Representing Symmetries: Automorphism Groups



# Representing Partial Symmetries: Partial Automorphism Inverse Monoids



# Representing Partial Symmetries: Partial Automorphism Inverse Monoids



$$\Rightarrow [1\ 2][4\ 5][7\ 8]$$



# Extracting (Partial) Symmetries

- Describe MPSoC architecture as *architecture graph*



# Extracting (Partial) Symmetries

- Describe MPSoC architecture as *architecture graph*
- Determine Automorphism Group  $G$ :
  - Transform into “equivalent” vertex colored graph
  - Use *nauty* [McKay and Piperno, 2014]


# Extracting (Partial) Symmetries

- Describe MPSoC architecture as *architecture graph*
- Determine Automorphism Group  $G$ :
  - Transform into “equivalent” vertex colored graph
  - Use *nauty* [McKay and Piperno, 2014]
- Determine Partial Automorphism Inverse Monoid  $M$ :
  - Construct search tree of possible generators
  - Prune certain subtrees
  - Not efficient enough in practice


# Extracting (Partial) Symmetries



# Extracting (Partial) Symmetries



## Extracting (Partial) Symmetries



## Extracting (Partial) Symmetries

$\{(1\ 2), (3\ 4)\} \Rightarrow$  Base:  $[1, 3]$   
Strong Generating Set:  $\{(1\ 2), (3\ 4)\}$

# The TMOR Problem I: Mapping Orbits

- Represent mappings by  $k$ -tuples  $t = (t_1, t_2, \dots, t_k)$



# The TMOR Problem I: Mapping Orbits

- Represent mappings by  $k$ -tuples  $t = (t_1, t_2, \dots, t_k)$
- Define *orbit*  $t^G = \{((g(t_1), g(t_2), \dots, g(t_k)) \mid g \in G\}$

# The TMOR Problem I: Mapping Orbits


- Represent mappings by  $k$ -tuples  $t = (t_1, t_2, \dots, t_k)$
- Define *orbit*  $t^G = \{((g(t_1), g(t_2), \dots, g(t_k)) \mid g \in G\}$
- Orbits **partition** search space

# The TMOR Problem I: Mapping Orbits

- Represent mappings by  $k$ -tuples  $t = (t_1, t_2, \dots, t_k)$
- Define *orbit*  $t^G = \{((g(t_1), g(t_2), \dots, g(t_k)) \mid g \in G\}$
- Orbits **partition** search space
- Reduce search space to set of *canonical orbit representatives*

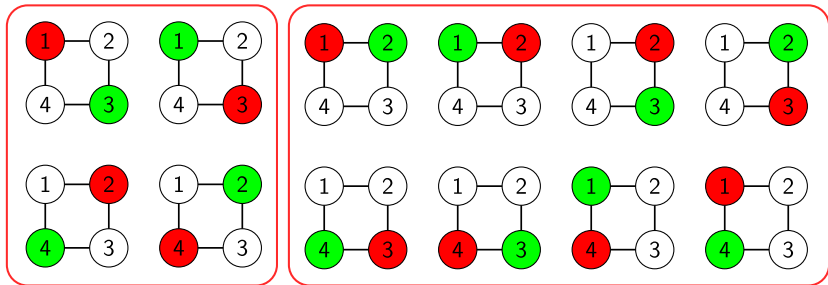
# The TMOR Problem I: Mapping Orbits

$$G = \langle \{(1\ 2\ 3\ 4), (2\ 4)\} \rangle$$



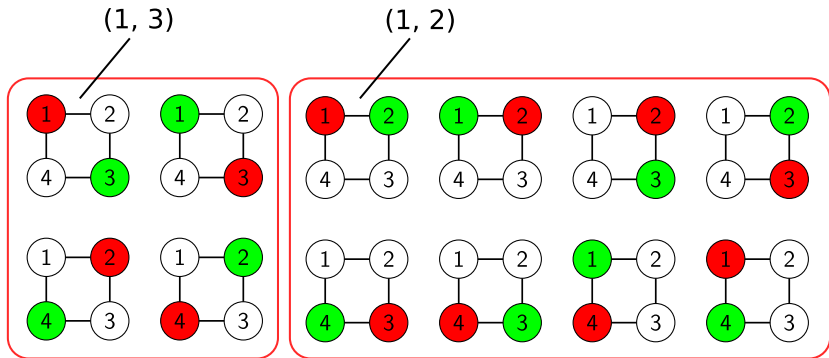
# The TMOR Problem I: Mapping Orbits

$$G = \langle \{(1\ 2\ 3\ 4), (2\ 4)\} \rangle$$



# The TMOR Problem I: Mapping Orbits

$$G = \langle \{(1\ 2\ 3\ 4), (2\ 4)\} \rangle$$



# The TMOR Problem I: Mapping Orbits

- Determining all orbits usually too costly

# The TMOR Problem I: Mapping Orbits

- Determining all orbits usually too costly
- Iteratively determine and hash canonical representatives



# The TMOR Problem I: Mapping Orbits

- Determining all orbits usually too costly
- Iteratively determine and hash canonical representatives

```
1: procedure ORBIT_IDENTIFIER( $t$ , reprs)
2:   repr  $\leftarrow$  TMOR( $t$ )
3:
4:   if repr  $\notin$  reprs then
5:     Append repr to reprs
6:   end if
7:
8:   return index of repr in reprs
9: end procedure
```

## The TMOR Problem II: Finding Canonical Orbit Representatives

- First approach: explicit orbit enumeration

```
1: procedure TMOR_ORBIT( $t, G = \langle S \rangle$ )
2:   orbit  $\leftarrow \{ \}$ 
3:
4:   while orbit is changing do
5:     for  $t' \in$  orbit,  $g$  in  $S$  do
6:       orbit  $\leftarrow$  orbit  $\cup \{g(t')\}$ 
7:     end for
8:   end while
9:
10:  return min(orbit)
11: end procedure
```

## The TMOR Problem II: Finding Canonical Orbit Representatives

- Second approach: group enumeration

```
1: procedure TMOR_ITERATE( $t, G$ )
2:   repr  $\leftarrow t$ 
3:
4:   for  $g \in G$  do
5:     if  $g(t) < \text{repr}$  then
6:       repr  $\leftarrow g(t)$ 
7:     end if
8:   end for
9:
10:  return repr
11: end procedure
```

## The TMOR Problem II: Finding Canonical Orbit Representatives

- Third approach: local search

```
1: procedure TMOR_LOCAL_SEARCH( $t, G = \langle S \rangle$ )
2:   repr  $\leftarrow t$ 
3:
4:   while repr is changing do
5:     repr  $\leftarrow \min(\{g(\text{repr}) \mid g \in S\})$ 
6:   end while
7:
8:   return repr
9: end procedure
```

## The TMOR Problem III: Optimization by Decomposition

- Many architecture graphs are *separable* or *hierarchical*

## The TMOR Problem III: Optimization by Decomposition


- Many architecture graphs are *separable* or *hierarchical*
- Separable: any heterogeneous architecture

# The TMOR Problem III: Optimization by Decomposition

- Many architecture graphs are *separable* or *hierarchical*
- Separable: any heterogeneous architecture
- Hierarchical: HAEC

# The TMOR Problem III: Optimization by Decomposition


- Many architecture graphs are *separable* or *hierarchical*
- Separable: any heterogeneous architecture
- Hierarchical: **HAEC**





# The TMOR Problem III: Optimization by Decomposition

- Many architecture graphs are *separable* or *hierarchical*
- Separable: any heterogeneous architecture
- Hierarchical: HAEC, **Kalray**



# The TMOR Problem III: Optimization by Decomposition

- Automorphism groups of such graphs decompose into:

# The TMOR Problem III: Optimization by Decomposition

- Automorphism groups of such graphs decompose into:

- Separable architectures: *direct products*

$$G = G_1 \times G_2, |G| = |G_1| \cdot |G_2| \cdots |G_n|$$

- Hierarchical architectures: *wreath products*

$$G = G_{\text{proto}} \wr G_{\text{super}}, |G| = |G_{\text{proto}}|^{\deg(G_{\text{super}})} \cdot |G_{\text{super}}|$$

# The TMOR Problem III: Optimization by Decomposition

- Automorphism groups of such graphs decompose into:
  - Separable architectures: *direct products*  
 $G = G_1 \times G_2, |G| = |G_1| \cdot |G_2| \cdots |G_n|$
  - Hierarchical architectures: *wreath products*  
 $G = G_{\text{proto}} \wr G_{\text{super}}, |G| = |G_{\text{proto}}|^{\deg(G_{\text{super}})} \cdot |G_{\text{super}}|$
- Solve TMOR problem separately for components

# The TMOR Problem III: Optimization by Decomposition

- Automorphism groups of such graphs decompose into:
  - Separable architectures: *direct products*  
 $G = G_1 \times G_2, |G| = |G_1| \cdot |G_2| \cdots |G_n|$
  - Hierarchical architectures: *wreath products*  
 $G = G_{\text{proto}} \wr G_{\text{super}}, |G| = |G_{\text{proto}}|^{\deg(G_{\text{super}})} \cdot |G_{\text{super}}|$
- Solve TMOR problem separately for components
- Decomposition specified by user or detected automatically

# The TMOR Problem III: Optimization by Decomposition


- Automorphism groups of such graphs decompose into:
  - Separable architectures: *direct products*  
 $G = G_1 \times G_2, |G| = |G_1| \cdot |G_2| \cdots |G_n|$
  - Hierarchical architectures: *wreath products*  
 $G = G_{\text{proto}} \wr G_{\text{super}}, |G| = |G_{\text{proto}}|^{\deg(G_{\text{super}})} \cdot |G_{\text{super}}|$
- Solve TMOR problem separately for components
- Decomposition specified by user or detected automatically
- Idea based on [Donaldson and Miller, 2009]

# Experimental Results I: Architectures

- Experiments run for: **Exynos**

# Experimental Results I: Architectures


- Experiments run for: Exynos , **Parallella**






# Experimental Results I: Architectures

- Experiments run for: Exynos , Parallella , **HAEC**




# Experimental Results I: Architectures

- Experiments run for: Exynos , Parallella , HAEC , **Kalray**



# Experimental Results I: Architectures


- Experiments run for: Exynos , Parallella , HAEC , Kalray



# Experimental Results II: TMOR

## Exynos TMOR


- Iterate (mpsym)      Iterate (GAP)
- Local search (mpsym)      Local search (GAP)
- Orbits (mpsym)      Orbits (GAP)



# Experimental Results II: TMOR

## Parallella TMOR


- Iterate (mpsym)      Iterate (GAP)
- Local search (mpsym)      Local search (GAP)
- Orbits (mpsym)      Orbits (GAP)



# Experimental Results II: TMOR

## HAEC TMOR


- Iterate (mpsym)
- Iterate\* (mpsym)
- Local search (mpsym)
- Local search\* (mpsym)
- Orbits (mpsym)
- Orbits\* (mpsym)



# Experimental Results II: TMOR

## Kalray TMOR

Iterate (mpsym) Local search (mpsym) Orbits (mpsym)



## Experimental Results II: TMOR

- Lessons learned:
  - Performance of orbit enumeration depends strongly on  $k$



# Experimental Results II: TMOR

- Lessons learned:
  - Performance of orbit enumeration depends strongly on  $k$
  - Local search can be fast and accurate

# Experimental Results II: TMOR

- Lessons learned:
  - Performance of orbit enumeration depends strongly on  $k$
  - Local search can be fast and accurate
  - Decomposition can be very powerful

# Experimental Results II: TMOR

- Lessons learned:
  - Performance of orbit enumeration depends strongly on  $k$
  - Local search can be fast and accurate
  - Decomposition can be very powerful
  - `mpsymb` outperforms GAP

## Further Research

- Heuristic local search

## Further Research

- Heuristic local search
- Partial Automorphism discovery

## Further Research

- Heuristic local search
- Partial Automorphism discovery
- Inverse Monoid enumeration

## Further Research

- Heuristic local search
- Partial Automorphism discovery
- Inverse Monoid enumeration
- Interfacing GAP and C++

# The End

Thank you for your attention!



# References



Donaldson, A. F. and Miller, A. (2009).  
On the constructive orbit problem.  
*Ann Math Atrif Intell*, 57:1–35.



East, J., Egri-Nagy, A., Mitchell, J., and Péresse, Y. (2019).  
Computing finite semigroups.  
*Journal of Symbolic Computation*, 92:110–155.



Goens, A., Siccha, S., and Castrillon, J. (2017).  
Symmetry in software synthesis.  
*ACM Trans. Archit. Code Optim.*, 14(2).



Holt, D. F. (2005).  
*Handbook of Computational Group Theory*.  
CRC Press.



McKay, B. D. and Piperno, A. (2014).  
Practical graph isomorphism, ii.  
*Journal of Symbolic Computation*, 60:94 – 112.



Nicolai, T. (2020).  
mpsym.  
<https://github.com/Time0o/mpsym>.