Symmetry in Software Synthesis

ANDRÉS GOENS, Center for Advancing Electronics Dresden (cfaed), Chair for Compiler Construction, TU Dresden, Germany

SERGIO SICCHA, Lehr- und Forschungsgebiet Algebra, Lehrstuhl B für Mathematik,

RWTH Aachen, Germany

JERONIMO CASTRILLON, Center for Advancing Electronics Dresden (cfaed), Chair for Compiler Construction, TU Dresden, Germany

With the surge of multi- and many-core systems, much research has focused on algorithms for mapping and scheduling on these complex platforms. Large classes of these algorithms face scalability problems. This is why diverse methods are commonly used for reducing the search space. While most such approaches leverage the inherent symmetry of architectures and applications, they do it in a problem-specific and intuitive way. However, intuitive approaches become impractical with growing hardware complexity, like Network-on-Chip interconnect or heterogeneous cores. In this article, we present a formal framework that can determine the inherent local and global symmetry of architectures and applications algorithmically and leverage these for problems in software synthesis. Our approach is based on the mathematical theory of groups and a generalization called inverse semigroups. We evaluate our approach in two state-of-the-art mapping frameworks. Even for the platforms with a handful of cores of today and moderate-sized benchmarks, our approach consistently yields reductions of the overall execution time of algorithms. We obtain a speedup of more than $10\times$ for one use-case and saved 10% of time in another.

CCS Concepts: • **Mathematics of computing** → **Discrete mathematics; Permutations and combinations**; • **Computer systems organization** → **Multicore architectures; System on a chip**; • **Software and its engineering** → **Runtime environments; Abstraction, modeling and modularity**; Compilers;

Additional Key Words and Phrases: Scalability, automation, inverse-semigroups, group theory, symmetry, design-space exploration, metaheuristics, mapping, network-on-chip, clusters, heterogeneous

ACM Reference Format:

Andrés Goens, Sergio Siccha, and Jeronimo Castrillon. 2017. Symmetry in software synthesis. ACM Trans. Archit. Code Optim. 14, 2, Article 20 (July 2017), 26 pages.
DOI: <http://dx.doi.org/10.1145/3095747>

1. INTRODUCTION

Several multiprocessor and many-core systems are available in the market today. Systems like embedded heterogeneous platforms, for example, TI Keystone II [Biscondi et al. 2012] with multiple ARM and Digital Signal Processor (DSP) cores, the homogeneous many-cores provided by Adapteva [Gwennap 2011], or the ARM big.LITTLE platforms [Greenhalgh 2011] all show a trend that is clear: Hardware architectures are becoming larger and more complex, for example, with large heterogeneous systems

New article; not an extension of a conference paper.

This work was funded in part by the Center for Advancing Electronics Dresden (cfaed) and the Graduiertenkolleg Experimentelle und konstruktive Algebra (GK Euka).

Authors' addresses: A. Goens and J. Castrillon, Chair for Compiler Construction, TU Dresden, 01062 Dresden, Germany; emails: {andres.goens, jeronimo.castrillon}@tu-dresden.de; S. Siccha, Lehr- und Forschungsgebiet Algebra, Pontdriesch 10-16, 52062 Aachen, Germany; email: sergio@mathb.rwth-aachen.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 1544-3566/2017/07-ART20 \$15.00

DOI: <http://dx.doi.org/10.1145/3095747>

Fig. 1. Abstract view of a software synthesis flow.

with clusters of processing elements and Network-on-Chip (NoC) interconnect, like is the case of the Kalray MPPA-256 [de Dinechin et al. 2013].

For these complex architectures, it is overly complicated for a designer to reason about parallel execution. Thus, the research community has invested a lot of effort in automatically deriving efficient implementations from abstract task-based specifications, particularly in the embedded domain [Keinert et al. 2009; Thiele et al. 2007; Castrillon et al. 2011] (see Figure 1). The *software synthesis* process includes defining the mapping of tasks and communication to resources, scheduling in case of resource sharing, and the final code generation, among others. In the embedded domain, where predictability is usually as much an asset as performance, the mapping is usually static and computed at compile time. This class of approaches has a distinct advantage over distributed approaches since it leverages global knowledge of the architecture. As a result, it is possible to obtain significantly better and more reliable executions than with dynamic, distributed methods. This is not without cost, however; this software synthesis approach requires a design-space exploration (DSE) with a huge space of design options, which usually grows exponentially with application and architecture sizes. Recent research has also focused on hybrid solutions, that is, partially defined at compile time with flexibility left for the runtime [Quan and Pimentel 2015; Schwarzer et al. 2017].

Due to the large size of the design space, the DSE typically uses either custom heuristics or metaheuristics for (multiobjective) optimization (e.g., evolutionary algorithms [Deb 2001], particle swarm, or ant colony optimization [Dorigo et al. 2008]). In this context, optimality refers to different objectives, like minimum execution time, energy consumption, or the best load distribution to avoid thermal effects. Examples are evolutionary algorithms in the programming frameworks Sesame [Erbas et al. 2006] and DOL [Thiele et al. 2007], and simpler iterative heuristics in SystemCo-Designer [Keinert et al. 2009], in MAPS [Castrillon et al. 2013], and in Cheung et al. [2007] and Casale Brunet et al. [2013] among others.

To deal with the large design space, it is very common for engineers and researchers to intuitively harness symmetry properties of the problem. Approaches designed to work for homogeneous, bus-based architectures usually have a formulation that considers all cores to be identical. When core heterogeneity is involved, only the type of core is considered. In fact, the notions of homogeneous and heterogeneous architectures are just approximations for describing symmetry in an architecture. For architectures with NoC, the situation is more complex. In those cases, most methods involve a hand-tailored reduction for the particular problem, like in Varghese et al. [2015], or pessimistic heuristics, for example, in Singh et al. [2013a]. To illustrate the principle, consider the example for a simple homogeneous many-core with an NoC interconnect shown in Figure 2. Due to the homogeneous structures of the platform, it is intuitive to see that the two leftmost mappings of tasks to processing nodes should lead to basically the same execution behavior, if we neglect effects like process variation or aging. Conversely, we can expect the third, rightmost mapping to have a different execution

Fig. 2. Illustration of symmetries. The two left mappings are equivalent, while the one on the right is not.

behavior, since the communication paths are different than in the first two. In the presence of heterogeneous resources, such an analysis becomes less obvious. Even more so with complex network topologies, in architectures with hierarchical structure, when optimizing for different objectives, or when all these are combined.

In this article, we introduce a general mathematical framework that allows reasoning about such general symmetries. It is based on the theory of groups and inverse semigroups, mathematical methods that enable us to define and quantify the symmetries. With it, we can reduce the search space in an automated fashion, just as it is traditionally done by hand for specific problems. The main advantage of this approach (over, e.g., groups) is that it not only considers *global* symmetries, but also *local* ones. Additionally, the mathematical nature of the approach allows us to guarantee correctness and completeness of the described symmetry. In particular, since we can find **all** symmetries, we can optimize the amount of reuse in algorithms. We demonstrate how our proposed methods can be applied to task-based programming models and to several kinds of architectures, including bus-based and NoC-based architectures. We analyze algorithmic implementations to find symmetries and their complexity. Finally, we integrate our approach into two state-of-the-art DSE frameworks, showing a significant reduction in the number of candidates evaluated, even for today's moderate-sized systems.

The rest of the article is organized as follows. Sections 2 and 3 introduce the mathematical foundations for symmetry and the notion of equivalent objects, respectively, while discussing their application in a software-synthesis context. Section 4 deals with the algorithmic realization, while the approach is evaluated for two use-cases in Section 5. Finally, Sections 6 and 7 discuss related work and conclude the article.

2. SYMMETRIES AND INVERSE SEMIGROUPS

Symmetry is a concept that is well-known intuitively, but hard to make precise without a rigorous treatment. An intuitive definition of symmetry is the quality of a system to be transformed, in a controlled manner, without changing its properties. Modern hardware architectures exhibit a high degree of symmetry, even in heterogeneous platforms which usually feature several processors of the same type. Similarly, applications can sometimes feature a significant amount of symmetry, for example, in some cases of data-level parallelism when the mapping of two identical worker threads can be swapped without affecting performance. Both application and architecture symmetries combined induce symmetry in the mapping of computation and communication to hardware resources.

Fig. 3. A 2×2 NoC architecture and its rotation by 90° .

In this section, we present a formal mathematical treatment of these intuitively motivated properties. For this, we use the theory of groups, which is the conventional theory for describing symmetry in mathematics. Considering communication in NoC-based architectures, we also motivate the need for a generalization using so-called inverse semigroups.

2.1. Groups and Symmetries

Symmetry is a central concept in mathematics and has been studied extensively. Most commonly, symmetry is described using the mathematical field of group theory. Group theory studies symmetries of an object by analyzing the transformations that preserve the structure of the object. The set containing all these transformations is called a group. More precisely, if we refer to the object as $\mathbb{O}$, said transformations are called the *Group of Automorphisms* of $\mathbb{O}$. As an example, consider the architecture depicted in Figure 3. It represents a homogeneous multiprocessor system-on-chip (MPSoC) with four identical processing elements (PEs) connected by a 2×2 mesh NoC. Hence, $\mathbb{O} := \{\text{PE}_1, \dots, \text{PE}_4\}$. Mathematically, we regard the transformations of this architecture as functions $t : \mathbb{O} \rightarrow \mathbb{O}$. If we ignore communication, all four processing elements can be permuted arbitrarily, since they are identical. This means that if we change the names of the PEs in any way, we still have the same architecture in principle. However, if we consider communication, not every reordering of the PEs will respect the structure of the NoC. In fact, there are precisely eight symmetry transformations that, when applied to this architecture, preserve its structure. The simplest one is doing nothing: this is usually called the trivial or identity transformation. A more interesting transformation consists of rotating the PEs by 90° , as depicted in Figure 3. While the physical processors are not the same, they are all the same kind, and the intercommunication network also has the same structure. Thus, we consider this rotated architecture on the right of Figure 3 to be equivalent to the one on the left. Executing an application on this rotated architecture should yield the same results as on the original one, except for physical characteristics that cannot be predicted and accounted for at design time.

We can repeat this rotation up to three times and get different transformations which preserve the structure of the whole architecture. After four rotations, the original architecture is obtained, which is equivalent to the identity transformation. Reflection along the vertical axis is also a valid transformation, as depicted in Figure 4(a). In this case, the identity transformation results from applying the reflection twice. Furthermore, combining one fixed reflection with the three rotations of 90° , 180° , 270° , respectively, yields three other reflections. In total, there are eight possible transformations that preserve the structure of the architecture. This group of transformations is well known as the *dihedral group* of order 8.

As depicted in Figure 4(b), swapping the processors PE_1 and PE_2 is not a symmetry transformation when communication is considered. In the figure, the thicker communication lines were not in the architecture before the transformation, whereas the dotted lines were. In other words, the reason that this transformation is not a symmetry is

Fig. 4. Symmetry transformations in a 2×2 NoC-based architecture.

Fig. 5. Partial functions define more symmetries in a 4×4 mesh.

that it does not preserve the communication costs between processors. For example, before applying the transformation, PE_1 is a direct neighbor of PE_3 . Afterward, it is not.

The transformations are thus captured in an abstract mathematical structure, a group, which we shall denote by G . The application of these transformations to the object $\mathbb{O}$ is called the *action* of the group G on $\mathbb{O}$, whereas the elements of $\mathbb{O}$ are usually called *points*. Note that G can also operate on another object $\mathbb{O}'$ containing, for example, subarchitectures or mappings as points, instead of single PEs. Groups are characterized by certain properties²:

- (1) Group elements, that is, the symmetry transformations, can be combined and give rise to other symmetry transformations. In the case of the architecture, this means applying symmetry transformations one after another yields an additional symmetry of the architecture.
- (2) There is always a unique symmetry transformation that does nothing; it is called the identity and denoted by $\text{Id} \in G$.
- (3) Every symmetry transformation can be undone. This means that for any symmetry transformation $t \in G$, there exists a unique symmetry transformation, $t^{-1} \in G$, such that combining both yields the identity transformation: $tt^{-1} = t^{-1}t = \text{Id}$. In our case, for example, rotating by 90° can be undone by rotating by 270° .

Informal names for these properties are *closedness under composition*, *existence of a unique identity*, and *existence of a unique inverse*.

2.2. Symmetry as Inverse Semigroups of Architectures

Consider the architecture depicted in Figure 5 as our object $\mathbb{O}$. It represents a 4×4 NoC mesh with 16 RISC PEs, each with a scratchpad memory. This architecture is inspired by the Paralella chip from Adapteva [Gwennap 2011]. However, this regular

¹We omitted associativity in the axioms of a group since function composition is associative.

mesh structure can be found in several modern multicores, for example, Balkind et al. [2016] and Ramey [2011]. For illustration purposes, we are not considering other factors like off-chip memories and peripherals in this example, but dealing with these is straightforward in this model, since they just change the symmetry group of the structure. The architecture from Figure 5 is similar to the architecture on the left of Figure 3: they both have n -by- n NoC topologies. In fact, by the formalism of symmetry groups, both architectures have the same symmetry. There are only eight transformations that preserve the structure of the whole architecture and they are the same as those of the 2×2 mesh: the identity, three rotations, and four reflections. The same is true for any such $n \times n$ mesh NoC.

However, intuitively, these are not the only symmetries of this architecture. To see this, compare the colored regions in the different depictions of the architecture in Figure 5. A computation using only the cores of the blue region to the left should have the same performance values as the one in the blue region to the right, with the transformation implied by the dark-blue arrow:

$$PE_1 \mapsto PE_2, PE_5 \mapsto PE_6, PE_9 \mapsto PE_{10}, PE_{13} \mapsto PE_{14} \quad (1)$$

The fundamental difference between this transformation and the aforementioned symmetries like rotations and reflections is that it is only defined on a subset of the 4×4 mesh and cannot be extended to a symmetry of the full mesh. In this context, we will denote symmetry transformations of an object that are not defined on the whole object as partial symmetries. Mathematically speaking, the partial symmetry transformations of this mesh do not form a group. They do not satisfy property 3 in its full rigor, as we shall see shortly. Instead of a group, all partial symmetries of an object constitute what is called an inverse semigroup, which can be seen as a generalization of a group [Lawson 1998]. With inverse semigroups of symmetry, we do not consider functions $t : \mathbb{O} \rightarrow \mathbb{O}$. Instead, we consider *partial functions*, or more precisely, *partial permutations* $t' : \mathbb{O} \rightarrow \mathbb{O}$, *partial* meaning that the domain and co-domain (image) of t' are only defined as subsets of $\mathbb{O}$. Thus, t' corresponds to a permutation $S \rightarrow S'$ from a subobject (subset) $S \subseteq \mathbb{O}$ to another (possibly different) subobject $S' \subseteq \mathbb{O}$. Due to this, when the co-domain of a transformation and the domain of another do not agree, the composition of the transformations is only defined on the subset where they match (which can be empty).

More precisely, we can equip the partial permutations with an extended composition rule such that they form an inverse semigroup: if $f : X \rightarrow Y$ is a partial permutation and $g : Y \rightarrow Z$ is a partial permutation, then their composition $g \circ f : X \rightarrow Z$ is another partial permutation. $f \circ g$ corresponds to a permutation from a subset of X to a subset of Z such that $(g \circ f)(x) = g(f(x))$ for all $x \in X$ where $f(x)$ and $g(f(x))$ are defined, that is, x is in the domain of f and $f(x)$ is in the domain of g .

For example, consider $X = Y = Z = \{1, \dots, 16\}$. We use this set instead of the PEs in the 4×4 NoC mesh for simplifying the notation. Consider also the partial permutations $f : 1 \mapsto 13, 5 \mapsto 9, 6 \mapsto 10, 7 \mapsto 11, 8 \mapsto 12, 12 \mapsto 8$, and $g : 1 \mapsto 2, 5 \mapsto 6, 9 \mapsto 10, 13 \mapsto 14$. These two transformations are depicted in Figure 6, f as the green mapping and g as the blue one. Then $gf = g \circ f$ is only defined for $\{1, 5\}$, since only $f(1)$ and $f(5)$ are in $\{1, 5, 9, 13\}$, the domain of g . Here, $gf(1) = (g \circ f)(1) = g(f(1)) = g(13) = 14$, and similarly, $gf(5) = g(f(5)) = g(9) = 10$, as can also be seen in Figure 6 by the orange mapping.

Groups have been studied much more in-depth than inverse semigroups have. With the extended composition rule, we generalize several concepts from groups to use them for inverse semigroups.

To understand the difference between groups and inverse semigroups at a more formal level, consider again the transformation $t = g$, depicted in Figure 6. We stated

Fig. 6. The product of partial permutations. The image of f and the domain of g overlap at 9 and 13.

earlier that general collections of partial permutations do not satisfy property (3) in its full rigor. Of course, g can be undone by a partial permutation g^{-1} . The resulting transformation gg^{-1} is only defined on $\{PE_1, PE_5, PE_9, PE_{13}\}$, though. We have $gg^{-1} \neq Id_{\mathbb{O}}$, that is, gg^{-1} is not the identity transformation on the whole of $\mathbb{O}$, which is not in line with property 3. Instead, the product $gg^{-1} =: i$ is called a *partial identity* on the subset $\{PE_1, PE_5, PE_9, PE_{13}\}$, that is, $i = Id_{\{PE_1, PE_5, PE_9, PE_{13}\}}$, the restriction of the (global) identity to this subset. In particular, it has the property that $ii = i$ and elements with this property are usually called *idempotents*. In groups, there is always exactly one idempotent, the identity. In our inverse semigroups, the idempotents are exactly all partial identities.

A transformation t in an inverse semigroup has an inverse element in a generalized sense, as illustrated earlier, which we will continue to denote as t^{-1} by abuse of notation. For our inverse semigroups, instead of the property 3, we get two new properties that are more general, albeit somewhat technical:

- (2') Partial identities on subsets of the architecture commute: if we have two idempotents $i = Id_I$, $f = Id_F$ for $I, F \subseteq \mathbb{O}$, then $if = fi = Id_{I \cap F}$.
- (3') For generalized inverse elements, we do not have $tt^{-1} = Id$, since tt^{-1} is not defined on the whole set, unlike Id . What does hold, however, are the equations

$$tt^{-1}t = t \text{ and } t^{-1}tt^{-1} = t^{-1}.$$

This rule actually already implies that the product tt^{-1} is an idempotent, as can be verified by a simple calculation.

In general, inverse semigroups need not satisfy property (2') either, but those studied in this article always do. In technical terms, thus, our semigroups are also *monoids*.

The inverse semigroup of symmetries of the 4×4 mesh, for example, includes the partial function $t = g$ from Equation (1). It yields an additional symmetry of the mesh, which is not embraced by the group of architecture symmetries. In general, inverse semigroups represent a richer set of symmetries by also considering symmetries between the substructures of an object.

2.3. Breaking Symmetries

In many occasions, the goal of software synthesis is not to optimize for latency, but, for example, power consumption or for avoiding thermal issues. Peripherals might be involved, or nonuniform access to main memory through the NoC. We have not dealt with this in the previous examples, but our methods are very capable of dealing with these scenarios. Indeed, the concept of symmetry in terms of inverse semigroups, as introduced here, provides us with a **language to describe and leverage the inherent symmetries** of the hardware (and software). The particular descriptions in this article

are mere examples of this. If a peripheral can only be accessed through a controller on one end of the NoC, then the symmetries studied become partial symmetries in a larger inverse semigroup. This larger inverse semigroup takes said peripheral into account and thus has less global symmetries.

3. EQUIVALENCE OF MAPPINGS

The formalism of groups and inverse semigroups can describe the symmetries of architectures and applications. In order to leverage them, however, one has to understand how this affects the mappings of software applications to hardware resources. In this section, we explain how the formal concepts of symmetry can be used to define when two mappings of computation and communication to hardware resources are equivalent.

3.1. Equivalence of Mappings

As motivated intuitively in the introduction (cf. Figure 2), different mappings of computation to hardware resources (PEs and communication) can be considered equivalent in many cases. For a given metric (e.g., latency or energy consumption), if two mappings of computation to hardware should yield the same results, only one has to be evaluated. Sometimes, on the other hand, we might explicitly want to use a different mapping, but retain the same performance guarantees: this might be interesting, for example, when optimizing thermal distribution within the chip, or to select an alternative mapping at runtime, depending on the available platform resources. Thus, these methods apply for a larger scale of the software synthesis process. They go beyond reducing the design space for exploration.

We collect equivalent elements of the objects, mappings in this case, in (unordered) sets. These sets are called *orbits* of the group/inverse semigroup *action*.³ To explain actions and orbits, we refer again to the 4×4 mesh architecture (cf. Figure 6). As was mentioned before, if we consider the **group** G of symmetries, it consists of the same eight transformations as the 2×2 mesh. A group action describes a way to apply a group element, or transformation, t to a point $o \in \mathbb{O}$. We denote the result of this application by $t \cdot o$. Here, we choose the action $t \cdot o := t(o)$, that is, simply apply the transformation function. Given an action, we can then define the orbit of a point $o \in \mathbb{O}$ to be

$$G \cdot o := \{t(o) \mid t \in G\}.$$

For example, consider the mesh itself as our object, $\mathbb{O} = \{1 \dots 16\}$, with our simplified notation. Consider then the orbit $G \cdot 1$ and let $t_1 = \varphi$ denote the anticlockwise rotation by 90° . Then, $13 \in G \cdot 1$, since $\varphi(1) = 13$ and since $4 = \varphi(16) = \varphi(\varphi(13))$ and $\varphi(4) = 1$, we get $G \cdot 1 = \{1, 13, 16, 4\}$, the set of all corner elements. Applying the other elements, that is, the reflections, yields no additional elements in the orbit.

If we consider the inverse semigroup of symmetry S , since all PEs are equivalent, the orbit of 1 under S is the whole architecture, $S \cdot 1 = \mathbb{O}$. If instead of a single PE, however, we consider the orbit of the subarchitecture $\{1, 2, 3, 5, 6, 7\}$, then, under the symmetry group, the orbit has four elements, namely $\{2, 3, 4, 6, 7, 8\}$, $\{9, 10, 11, 13, 14, 15\}$, $\{10, 11, 12, 14, 15, 16\}$, as depicted in Figure 7(a). In the case of the semigroup, on the other hand, another additional eight elements are contained in the orbit, as depicted in Figure 7(b).

A more interesting example is the orbit of a mapping. Let $T = \{T_1, \dots, T_s\}$ denote a set of tasks and $P = \{PE_1, \dots, PE_r\}$ a set of processing elements. In the previous section, we mentioned that symmetries can act both on the architecture, that is, on P ,

³In the literature, for inverse semigroups, these are usually called strong orbits, or the strongly connected components of the orbit graph, but we will use this simpler definition in this article.

Fig. 7. Examples of orbits on a 4×4 mesh NoC. Fig. 8. An example of an application: a stereo filter. Swapping the pipelines acting on the two channels, π , is a symmetry of this application.

and also on the tasks T . Let H denote the group of symmetries of the tasks and G the group of symmetries of the processors. As the object set, we define $\mathbb{O}$ to be the set of all task-to-processor mappings $m : T \rightarrow P$. The symmetries described by these two groups are independent of each other for a mapping, and we can think of them as two groups acting separately. This is called the *direct product* of G and H , and is written as $G \times H$.

Consider as an example the application depicted in Figure 8. It describes an audio filter, where a source task T_1 separates the stereo input into two channels. For each channel, an inverse Fourier transform (T_2, T_5), a filter on the frequency domain (T_3, T_6), and a Fourier transform (T_4, T_7) are applied and gathered at a sink T_8 . This performs the same computation on the pipeline T_2, T_3, T_4 as in T_5, T_6, T_7 , such that swapping both (complete) pipelines has no effect on the performance. We call this transformation on the application graph π ; see Figure 8. Let us assume that we want to map this audio filter to the $r = 4$ processors of the architecture from Figure 3. To simplify the notation, we identify a mapping $m : T \rightarrow P$ with a tuple of the form $(m(T_1), \dots, m(T_s))$. This way, the mapping $m_1 : T_1 \rightarrow \text{PE}_2, T_2 \rightarrow \text{PE}_3, T_3 \rightarrow \text{PE}_3, T_4 \rightarrow \text{PE}_3, T_5 \rightarrow \text{PE}_4, T_6 \rightarrow \text{PE}_4, T_7 \rightarrow \text{PE}_4, T_8 \rightarrow \text{PE}_1$ can be written as $m_1 = (\text{PE}_2, \text{PE}_3, \text{PE}_3, \text{PE}_3, \text{PE}_4, \text{PE}_4, \text{PE}_4, \text{PE}_1) = (2, 3, 3, 3, 4, 4, 4, 1)$, for example, again simplifying the notation.

To understand the action, consider the symmetry of the architecture τ , a reflection through the horizontal line (with respect to the depiction), which swaps PE_1 and PE_3 as well as PE_2 and PE_4 . The action of a symmetry of the architecture, that is, an element of G like π on $\mathbb{O}$ is straightforward. A permutation is just applied to each entry of a tuple m individually. Thus, as depicted in Figure 9(a), for $\tau \in G$, we have

$$\tau \cdot m_1 = (\tau(2), \tau(3), \tau(3), \tau(3), \tau(4), \tau(4), \tau(4), \tau(1)) = (4, 1, 1, 1, 2, 2, 2, 3).$$

A permutation of the tasks, in contrast to a symmetry of the architecture, permutes the entries of a tuple horizontally, that is, permutes the indices in the tuple. For example, the permutation π on the application (Figure 8). Swapping the tasks 2,3,4

Fig. 9. Actions of symmetries on the mapping m_1 . Dotted lines denote the changes in the mapping.

and 5,6,7 would swap the second, third, and fourth entries of the tuple, representing m_1 for its fifth, sixth, and seventh entries, respectively:

$$\pi \cdot m_1 = \pi \cdot (2, 3, 3, 3, 4, 4, 4, 1) = (2, 4, 4, 4, 3, 3, 3, 1).$$

This is illustrated in Figure 9(b). Note that all three pairs of tasks have to be swapped simultaneously. Using the action from just π , and the symmetry τ , the orbit of m_1 would be

$$\{(4, 1, 1, 1, 2, 2, 2, 3), (4, 2, 2, 2, 1, 1, 1, 3), (2, 4, 4, 4, 3, 3, 3, 1), (1, 3, 3, 3, 2, 2, 2, 4), \\ (3, 1, 1, 1, 4, 4, 4, 2), (1, 2, 2, 2, 3, 3, 3, 4), (3, 4, 4, 4, 1, 1, 1, 2), (2, 3, 3, 3, 4, 4, 4, 1)\}$$

These mappings constitute the equivalence class of $m_1 = (2, 3, 3, 3, 4, 4, 4, 1)$ and should, therefore, behave equally in any execution.

4. ALGORITHMIC CONSIDERATIONS

Up to this point, we have introduced formal concepts to quantify and identify symmetries in software synthesis precisely. These comprise groups, inverse semigroups, and orbits, for example, of a mapping or a subarchitecture. In order to leverage these concepts, we design an algorithm that, for a given architecture, finds the inverse semigroup of its symmetries. Then, we can use a standard algorithm to calculate the orbit of a given point at compile time. Both groups and inverse semigroups can be very large, and it is infeasible to store all their elements explicitly. The need for a compact representation, in turn, has led to the development of data structures, which can be used to work with these compact representations.

The computer algebra system GAP [The GAP Group 2016] implements many well-known algorithms in computational group theory. Currently, the GAP community is working toward making these also available in its parallel fork HPC-GAP [Behrends et al. 2016]. In this article, we use this computer-algebra system as a research vehicle to assess the effectiveness of symmetry considerations in this field, while we plan to use the parallel fork in future work.

This section describes the basic algorithmic concepts of our proposed methods, and discusses the overhead incurred by symmetry-related calculations and the potential for domain-specific optimizations.

4.1. A Compact Representation of Groups and Inverse Semigroups

For practical purposes, it is prohibitive to store or iterate over all elements of a group. Instead, a group G is usually given by a *generating set* S , written $G = \langle S \rangle$. We say the finite group G is *generated* by S if G consists of all products that can be written as words over S as alphabet. The advantage is that arbitrarily large groups can be written with small generating sets. An important example is the so-called *symmetric group*, written S_n , consisting of all permutations on n points. It has $n! = 1 \cdot \dots \cdot n$ elements, but can always be generated by the following two elements: a transposition $\tau : 1 \mapsto 2, 2 \mapsto 1$,

$3 \mapsto 3, \dots, n \mapsto n$, which only swaps the points 1 and 2 and the n -cycle $\pi : 1 \mapsto 2 \mapsto 3 \mapsto \dots \mapsto n-1 \mapsto n \mapsto 1$. This representation is obviously very compact, since it uses only two permutations, each moving up to n points to represent the whole group of $n!$ many permutations. It is a well-known fact that $S_n = \langle \tau, \pi \rangle$, that is, any permutation on n points can be written as a word in these two permutations. In general, there exist several generating sets for a given group or inverse semigroup, which may vary in their sizes and algorithmic properties. For a group, given by a generating set S of permutations on n points, many problems can be solved in polynomial time w.r.t. $|S|$ and n .

For our purposes, the most fundamental of these problems is the *membership test*. While computing a group or an inverse semigroup of symmetries, we collect the symmetries already found in a set S and consider the $G = \langle S \rangle$. Upon finding another symmetry t , we need to decide whether t is a “new” symmetry and needs to be added to S , or whether t can be written as a word in the symmetries in S collected so far. Deciding whether $t \in \langle S \rangle$ is called the *membership test*. Being able to solve membership tests efficiently is crucial when working with a group or an inverse semigroup given by generators. Usually, a data structure is computed, which facilitates fast *membership tests*. An algorithm to solve this for groups is the *Schreier-Sims algorithm*, first developed by C. Sims [1970]. It continues to be one of the main tools of computational group theory.

Since the inception of the Schreier-Sims algorithm, a tool-suite for working with arbitrary finite groups of permutations has been developed. These methods and many more are described by Seress in his book [Seress 2003]. State-of-the-art implementations of these algorithms are available in GAP [The GAP Group 2016] and Magma [Bosma et al. 1997], and mostly run in Monte Carlo nearly linear time $O^\sim(n \log n |S|)$.

4.2. Computing the Orbits

To compute the orbit $G \cdot m$, for both groups and semigroups, we can employ generating sets. Starting with a given point m , the simplest algorithm applies all generators of the group G to all points already encountered, repeatedly, until no new points are found. Then, we can be sure that the full orbit has been calculated. In general, for a group $G = \langle S \rangle$ and a point $m \in \mathbb{O}$, the orbit $G \cdot m$ can be computed with $\mathcal{O}(|G \cdot m| |S|)$ evaluations of the form $g \cdot m$ [Seress 2003]. In particular, the cardinality of the group is not relevant for the complexity of the algorithm; only the cardinalities of the orbit $G \cdot m$ and the generating set S are relevant, which are almost always significantly smaller. This is called the orbit algorithm.

In our case, the aim of using orbits of groups is only to identify whether two elements of the orbit are identical, for which we do not need to calculate the complete orbit. Instead, if a unique representative of every orbit can always be chosen, we can compare the orbits of two elements by comparing their unique representatives. This is sometimes called *canonization* or finding a *canonical representative*. Such a canonical representative can be found when there is an ordering on the points in $\mathbb{O}$, in which case the canonical representative of an orbit $G \cdot m$ can be chosen to be its minimal element. A special case of this (for the group S_n) yields the method employed in Thompson and Pimentel [2013].

To be able to calculate the orbit of a mapping, thus, we first need to compute a generating set of the automorphism group or semigroup. However, we only have to do this once for every architecture. If we have a generating set, we can use it in any subsequent orbit calculation.

4.3. Computing Symmetry Groups and Inverse Semigroups

Both architectures and the relationships between tasks can be modeled mathematically using graphs. We model an architecture by a multigraph that has a node for every PE, labeled with the PE type. For every different communication resource between two PEs,

Fig. 10. The architecture graph construction.

the graph has an edge annotated with that communication resource [Castrillon et al. 2012]. This graph is sometimes called the *architecture graph*. Figure 10 illustrates this on the example of the graph representation of a 4×4 homogeneous NoC-mesh architecture. The annotations are not shown for readability. We construct this kind of graph easily from architecture models. However, there are no widespread standard formal models for this [Goens et al. 2016], even though efforts have been made in this direction, like the SHIM standard from the Multicore Association [Gondo et al. 2014] or at a formal level [Pelcat et al. 2015].

On the application side, the situation is less straightforward. Most models of computation are based on a sequential, shared memory model, which means they are not well suited for separating the computation. Better suited for these are the task graphs commonly used for scheduling on multicore systems, or dataflow graphs from models of computation like Synchronous Data Flow (SDF) [Lee and Messerschmitt 1987] or Kahn Process Networks (KPN) [Kahn 1974].

The problem of calculating the group of automorphisms of a graph is an extensively studied subject. In practice, it is solvable efficiently for the majority of graph-classes and there are many specialized algorithms/frameworks like bliss [Junttila and Kaski 2011], conauto [López-Presa et al. 2011], saucy [Codenotti et al. 2013], nauty, and Traces [McKay and Piperno 2014] to name a few. For example, for a random graph of up to 10^3 nodes, they are able to determine its automorphism group in less than 0.01s [McKay and Piperno 2014]. All of these algorithms use modern variations of the Schreier-Sims algorithm to manage the graphs' automorphism groups. Computing the automorphism group of a graph is directly related to the Graph-Isomorphism-Problem, a problem in complexity theory famous for being apparently easier than NP-complete problems, while its complexity status is not yet settled. It has recently been claimed to be solvable in quasipolynomial time [Babai 2015] and is currently being peer reviewed.

The following fact is a reinterpretation of the definition of a graph's inverse semigroup of symmetries, which we use to calculate it.

FACT 4.1. *A partial permutation of the set of nodes of a graph is in its inverse semigroup of symmetries if and only if it induces an isomorphism of labeled graphs from the subgraph induced by its domain to the subgraph induced by its image.*

Based on this, we designed Algorithm 1 to find a generating set S of the semigroup of all such partial permutations. For two graphs $\Gamma = (V, E)$, $\Gamma' = (V', E')$ and a mapping $\varphi : W \rightarrow W'$ from a subset of vertices $W \subseteq V$, to a subset of vertices $W' \subseteq V'$, let $\Gamma|_W$ be the subgraph of Γ induced by the vertex set W , that is, $\Gamma|_W = (W, E|_W := \{(v, w) \in E \mid v, w \in W\})$. Then, we denote by $\varphi_\Gamma : \Gamma|_W \rightarrow \Gamma'|_{W'}$ the induced mapping on the induced

subgraphs, that is, $\varphi_\Gamma((v_1, v_2)) = (\varphi(v_1), \varphi(v_2))$, as is the natural way of extending the mapping to the edges. Note that this is not necessarily a partial automorphism of Γ . This is only the case if, additionally, it holds that $e \in E_{|W}$ if and only if $\varphi(e) \in E'_{|W}$.

ALGORITHM 1: Inverse Semigroup of Symmetries of a Graph

input: A graph $\Gamma = (V, E)$
output: A generating set of partial permutations S

```

1:  $S \leftarrow \emptyset$ 
2:  $A \leftarrow \langle S \rangle$  // the inv. semi. generated by  $S$ 
3:  $C \leftarrow \{\varphi \mid \varphi \text{ is a partial permutation of } V\}$ 
4: for all  $\varphi \in C$  do
5:   if  $\varphi_\Gamma$  is a partial automorphism then
6:     if  $\varphi \notin A$  then // membership test
7:        $S \leftarrow S \cup \{\varphi\}$ 
8:      $A \leftarrow \langle S \rangle$  // update the data structure
return  $S$ 
```

The correctness of Algorithm 1 follows from Fact 4.1 and its termination from the fact that the set of all partial permutations of a set is finite. Our implementation uses the computer algebra system GAP [The GAP Group 2016] and, specifically, its semigroups-package [Mitchell et al. 2016; East et al. 2015] for performing constructive recognition and membership testing in Algorithm 1, lines 2, 6, and 8.

It is easy to see that Algorithm 1 is not optimal, since it considers redundant generators. For example, consider two sets $V_1 \subsetneq V_2 \subseteq V$. For all partial functions $f \in \langle S \rangle$, which have V_2 as its domain, the partial function $f_{|V_1}$ is also in $\langle S \rangle$. As a first improvement, we developed a backtrack-based algorithm sketched in Algorithm 2. We utilize that if a partial permutation is not an isomorphism, neither are its extensions.

The recursive backtrack loops over the set of all partial permutations by organizing the partial permutations into a lattice and traversing it. The lattice, in which each node represents one partial permutation of V , is built up as follows: Let the root of the lattice be the empty partial permutation $\varepsilon : \{\} \rightarrow \{\}$. The successors of a node φ consist of all partial permutations φ' that extend φ by exactly one further point. For example, one successor of the node $\pi : 1 \mapsto 4, 3 \mapsto 2$ would be the node $\pi' : 1 \mapsto 4, 3 \mapsto 2, 4 \mapsto 1$. This way, if one node is not a partial automorphism, we can backtrack by throwing away all its successors, since they also cannot be partial automorphisms (line 2 in the backtrack function of Algorithm 2).

We can improve this algorithm by initializing S with partial automorphisms we can readily deduce from the NoC architecture, like the global symmetries from the group of automorphisms. There are two further technical improvements we need to make Algorithm 2 feasible. On the one hand, we need to make sure we do not perform the membership test for all restrictions of a partial automorphism φ , but only for φ itself. On the other hand, if we descend through the lattice by visiting all successors for every node, we would visit many nodes multiple times. For example, a partial permutation φ with domain $\{1, 3, 4\}$ would appear as a successor of both a node φ_1 with domain $\{1, 3\}$ and a node φ_2 with domain $\{1, 4\}$. We can make sure that each partial permutation is visited exactly once by adding new points to the domains only in ascending order.

4.4. Domain-Specific Optimizations

The standard methods from algorithmic group theory are designed to be general and include all corner cases. In the case of applications for software synthesis, this is not necessary. While we used existing standard algorithms or simple naive algorithms to

ALGORITHM 2: Inverse Semigroup of Symmetries of a Graph with Backtracking**input:** A graph $\Gamma = (V, E)$ **output:** A generating set of partial permutations S

- 1: $S \leftarrow \emptyset$
- 2: $A \leftarrow \langle S \rangle$ // the inv. semi. generated by S
- 3: $\varepsilon \leftarrow$ empty partial permutation
- 4: backtrack(ε)
- 5: **return** S

function: backtrack**input:** A partial permutation x

- 1: **if** x_G is not a partial automorphism **then**
- 2: **return**
- 3: **if** $x \notin A$ **then**
- 4: $S \leftarrow S \cup \{x\}$
- 5: $A \leftarrow \langle S \rangle$
- 6: **for all** $c \in \text{successors}(x)$ **do**
- 7: backtrack(c)

evaluate the applications, we can easily identify potential for domain-specific optimizations in these areas. In the future, we will investigate specially tailored algorithms that leverage assumptions about the possible symmetry groups. Using this domain-specific knowledge is bound to yield much better time-efficiency results.

For example, by using structural properties of groups, a dynamic-programming approach can be used to substantially reduce the memory footprint of the methods [Lübeck and Neunhöffer 2001].

Another possible domain-specific optimization would rely on the fact that hardware architectures, while varied, mostly have a limited amount of complexity. It would thus be possible to create a classification of common groups in the architectures and applications. This works particularly well since the nature of the symmetry groups for architecture graphs is compositional. For example, the group of automorphisms of a heterogeneous bus-based architecture with k different PE types is the direct product of the collection of symmetric groups S_{n_i} , where the architecture has n_i PEs of type i , which is the group of automorphisms of a homogeneous bus-based architecture for the respective PE types.

Finally, Algorithm 2 still has to iterate over the whole semigroup of automorphisms, which, even for the 4×4 NoC, already contains around 1.2 million partial symmetries. Since the NoC architecture graphs have a very regular structure, we are confident that, in future work, we can leverage it and develop a specialized version for such NoC architectures. Even more so, it might be possible to classify the semigroups of automorphisms of such architectures by purely theoretical means, circumventing the need for computations at all.

4.5. Computational Overhead

As we have seen in this section, using the methods of group and inverse semigroup theory at diverse stages of the software synthesis process, will inevitably lead to computational overheads. This leads to a tradeoff: while knowing that two mappings are equivalent can have benefits in terms of computation time, this requires a computational effort. It is not evident that these methods will be computationally beneficial in every use-case. However, there are several reasons why such use-cases are bound to be common. Some parts of the calculation, like the symmetry group of the architecture, can be executed only once, at design time for a hardware architecture, for instance.

Fig. 11. An overview of the different symmetries in software synthesis.

They can subsequently be reused on every compilation that targets the architecture. Figure 11 illustrates this by giving an overview of the different components of symmetry in software synthesis and their time-point at which they can be leveraged. In many use cases, a complex and time-costly simulation is executed for every mapping. By leveraging the fact that symmetry calculations only depend on the structure of the problem, the complexity of the simulation can be increased with minimal penalty, for example, by using more level of detail in critical points, like cache simulators or flit-level NoC simulators instead of coarse system-level events. This has no effect on the symmetry calculations.

5. EVALUATION

In this section, we present two use-cases using the proposed symmetry principles for reducing design-space exploration. We use the principles to enhance existing tools and evaluate the usefulness of considering symmetry in software-synthesis, or more specifically, in design-space exploration.

5.1. Use-Case: Genetic Algorithms

The first use-case deals with the scalability of a metaheuristic based on genetic algorithms (GAs), which maps tasks to processing elements. Concretely, we consider a design-space exploration suite of the Sesame framework [Pimentel et al. 2006] and use symmetries to prune the search space.

Sesame is an open-source framework that tackles the problem of efficiently executing parallel applications on heterogeneous MPSoCs. By structuring programs as KPNs [Kahn 1974], Sesame gathers execution information in the form of process traces.

Using high-level hardware models, Sesame can use execution traces to drive an accurate, albeit efficient, simulation of the execution of the KPN application. The high-speed and high-fidelity of these trace-based simulations permit the exploration of the design space of static mappings by using so-called metaheuristics, like evolutionary algorithms [Erbaş et al. 2006]. Metaheuristics are a general approach to solving optimization problems with large design spaces that exhibit structure. They are based on the principle of iterating over different solutions in order to improve them and reach near-optimal results. In the case of evolutionary algorithms, the method for iteration is inspired in the biological process of evolution. Solutions are called individuals, and they are recombined to produce offspring, also introducing random mutations, over several generations.

Due to the iterative nature of metaheuristics, they rely on several simulations of the application using different mappings. These simulations are computationally expensive, especially so for larger applications. In most cases, the simulation time dominates the total exploration time.

The mutation and cross-over properties of GAs, a type of evolutionary algorithm, are likely to produce similar mappings in large populations. The DSE framework includes a

Fig. 12. A high-level overview of the use-cases with our symmetry improvements.

cache, which substitutes a simulation with a look-up when a mapping has already been simulated. To improve the execution time and scalability of the evolutionary algorithm, we propose extending this caching strategy. Figure 12(a) shows a high-level overview of the Sesame flow using GAs for mapping and our proposed solution. In our extended caching strategy, the simulation is only executed once for all equivalent mappings (see Section 3.1). This significantly reduces the search space, since it replaces the space of valid mappings with the space of equivalence classes of valid mappings. The python implementation of this GA uses a hash-map as a cache. Since we only need to ensure we do not simulate a mapping if we have simulated an equivalent one before, we implemented this by using a canonical representative (see Section 3.1) of the class as the key. This is significantly more efficient than having to calculate the whole orbit when orbits become large, that is, in the case when there is much symmetry.

5.2. Use Case: Subarchitectures

The second use case considers an iterative heuristic for minimizing the usage of system resources in the presence of application constraints. For this, we use an iterative heuristic from a modern, commercial multicore compiler from the SLX Tool Suite [Silexica 2016]. Similar algorithms have been presented in Cheung et al. [2007] and Castrillon et al. [2013].

Instead of finding the system with maximal performance, in several scenarios, application developers are interested in using the least resources possible in a system, while meeting some performance constraints. This is the case in hard or soft real-time applications, for example, in the multimedia and automotive domains. The iterative algorithm explores mapping configurations using limited resources. However, these resources might be heterogeneous, which makes the simple approach of the algorithm very limited. The basic algorithm works as follows:

- (1) Start by considering a subarchitecture consisting of a single PE.
- (2) With heuristics, calculate a mapping for close-to-optimal performance in the current subarchitecture.
- (3) Use a fast trace-driven simulation to estimate the performance of the mapping.
- (4) If the mapping is fast enough in the simulation, finish and return the mapping.
- (5) If the mapping is too slow, consider an additional, randomly chosen PE in the subarchitecture, as well as the corresponding communication resources.
- (6) Return to Step 2.

This algorithm can be very ineffective on heterogeneous platforms. The concrete subarchitecture considered depends on the choices of PEs taken in Step 5, which is too simplistic. Thus, the algorithm does not consider several configurations using different

Table I. Benchmarks and Symmetry Properties

Name	Short Description	No. of pr./ch.	Group	Group Size
sobel	Sobel filter on 40×40 image	5/15	C_2	2
matmult	10×10 matrix multiplication	5/6	$\{1\}$	1
mjpeg	Motion JPEG decoder (128×128)	12/15	S_4	$4! = 24$
mandelbrot	Mandelbrot set calculation (16 Jobs)	18/32	$\{1\}$	1
audio filter	Stereo filter from Figure 8	8/8	C_2	2

C_n denotes the cyclic group of order n , while “pr./ch.” means processes and channels, and $\{1\}$ denotes the trivial group.

PE types that could perform better and completely disregards the architecture topology. For comparison, we will refer to this variant of the algorithm as **simple**.

On the other hand, considering every possible subarchitecture in a **brute-force** approach scales prohibitively poorly: an architecture with n PEs has $2^n - 1$ possible (nonempty) subarchitectures, even disregarding communication. This is considerably larger than the n subarchitectures considered by iteratively adding PEs.

To deal with this issue, we propose to use the architecture symmetries to select the relevant subarchitectures to be considered in the iterative algorithm. Before mapping, we choose a single representative of every equivalence class of the architecture. We then modify the iterative algorithm. Figure 12(b) sketches the flow using the improved algorithm. It is clear from the figure that the symmetry reduction step can be pre-computed, since it only depends on properties of the target architecture. Instead of just increasing the size of the subarchitecture, our modified algorithm iterates over all nonequivalent subarchitectures of a given size first. If none of the nonequivalent subarchitectures of size n give enough performance, then subarchitectures of size $n + 1$ are considered. This modification improves the algorithm by letting it acknowledge the heterogeneity of the architecture in a reasonable time. The exact number of equivalence classes depends on the architecture and the inverse semigroup describing its symmetry. In this context, we differentiate two additional strategies: **groups**, which use the symmetry group to differentiate subarchitectures, and **inv. semi.**, which use the inverse semigroup instead.

5.3. Results

We test both example use cases with several benchmarks from the signal processing and multimedia domains. For this, we use current applications and real architectures. While we can only expect little gains in small applications mapped into architectures with a handful of cores, these serve to evaluate the symmetry methods with realistic parameters.

5.4. GAs

Table I gives a summary of the benchmarks used, and in particular, the number of processes of each benchmark application, the isomorphy type of the symmetry group, and its size. The symmetry groups of the applications were determined by hand, by inspecting the code of the benchmark. Automating this task in a compiler is highly nontrivial in general and is outside the scope of this article. We tested the applications by mapping on an accurate model of a state-of-the-art heterogeneous multicore system, the Texas Instruments (TI) Keystone II. It features four ARM Cortex-A15 processing elements and eight DSPs [Biscondi et al. 2012]. The model was adapted [Goens et al. 2016] from a model of the commercial state-of-the-art MPSoC compiler from the SLX Tool Suite and confirmed with measurements from hardware [Odendahl et al. 2013]. A schematic view of the architecture can be seen in Figure 13.

Fig. 13. Schematics of the architecture models.

Fig. 14. A comparison of benchmarks for the improved caching strategy in Sesame. The red bars on top represent the overhead, while the dotted gray line marks the baseline (without symmetries).

The TI Keystone II has a symmetry group of size $4! \cdot 8! = 967,680$. It is isomorphic to the direct product of the symmetric groups on 4 and 8 points, usually denoted by $S_4 \times S_8$. For this example, the symmetry group already yields all symmetries, as motivated in Section 2.2. This is the case since the simple interconnect topology (bus-based) does not break the symmetry, only the different core types do.

We executed the DSE framework with 5, 20, 50, and 200 generations for all benchmarks, with five different (fixed) random seeds, repeating each such configuration five times. It uses an evolution-strategy called $\mu + \lambda$, which we used with fixed populations of size 20 and with 20 children. We measured the overall wall-clock time of the DSE using our improved cache strategy and the conventional cache. Additionally, we measured explicitly the time spent in symmetry-related calculations in the improved cache. The results are summarized in Figure 14. It shows the normalized average of the total execution time for all benchmarks, separating the time spent on the DSE from the overhead of the symmetry calculations. The difference to the baseline of 1 is the percentage of time, which was saved by avoiding redundant calculations. The plot shows the overhead is negligible in all cases; indeed, it represents between 0.1% and 0.2% of the total execution time across the different benchmarks. The plot also shows that the improved cache results in a net improvement of the execution time, in spite being a prototype implementation. As expected, the effects of symmetry were more visible in explorations with a larger number of generations (and thus, of simulations).

Additionally, we compared the number of cache hits in the original, unmodified cache and our symmetry-improved version, as can be seen in Figure 15. We see that a large majority of hits happen because exactly the same mapping is being calculated, instead

Fig. 15. Percentage of cached mappings, which were found using symmetries.

of an equivalent one. This hints to the fact that the variation in the space is small, and exacerbates the problem that the space itself is not being reduced directly. Indeed, removing the symmetries in the cache still leaves them in the chromosome models, which, technically speaking, is not a reduction of the search space, but rather just a runtime optimization of the DSE. It is also a tentative explanation why the results for the mjpeg and mandelbrot benchmarks are worse than the others, since these two have more processes and thus, larger design spaces. Thus, an in-depth improvement would include symmetry in the chromosome models, letting them characterize only the representatives of the orbits, instead of all mappings. In particular, since our strategy just improved the caching, the Pareto fronts of best mappings resulting from exploration stay unchanged. By factoring out the symmetry at the chromosome level, we should also be able to produce better mappings in less time.

5.5. Subarchitectures

For the second use-case, we used a different architecture. The Parallella[®] architecture, commercialized by Adapteva, includes 16 Epiphany [Olofsson et al. 2014] cores, each with its scratchpad memory, and a shared DRAM. We use this architecture to showcase the effects of considering communication in NoC-based architectures and inverse semigroups. The architecture model is part of the commercial SLX Toolsuite[®] MP-SoC compiler and is designed for the high-speed, high-accuracy requirements of this use-case. As application, we use an mjpeg decoder, described in Sheng et al. [2012].

The 4×4 mesh structure in the Parallella architecture has $2^{16} - 1 = 65,535$ possible subarchitectures (excluding the empty one). In contrast, the architecture has only 8,547 equivalence classes of (nonempty) subarchitectures considering the symmetry group, and these are, in turn, reduced to 6,803 if we consider the inverse semigroup. In most use-cases, not all subarchitecture sizes would be used. Instead, only architectures up to a particular size would be considered.

We evaluated the four variants of the strategy for choosing subarchitectures described in Section 5.2: simple, groups, inverse semigroups (inv. semi.), and brute-force. The results obtained with all four variants of the algorithm can be seen in Figure 16. The plot depicts, for every subarchitecture size, the maximal number of trials required to obtain the best mapping for that subarchitecture size. Several bars for the “simple” strategy cannot be seen in the plot since, in most cases, this strategy did not find the best mapping for a given subarchitecture size. In those cases, the best mapping found with the “simple” strategy was worse by an average margin of 2%. However, for smaller subarchitecture sizes, which are usually the ones we are interested in, the error was mostly around 7% to 10% in the experiments.

The benefit of using inverse semigroups over groups decreases for larger subarchitecture sizes. This is a consequence of the nature of the inverse semigroup of symmetries, which is composed of *partial* permutations. When the subarchitecture

Fig. 16. Maximal number of trials for obtaining the best mapping using the different strategies. The missing bars on the “simple” strategy represent failures of the strategy to find the best mapping.

(a) Comparison for a fixed size (4).

(b) Quality of the solutions.

Fig. 17. Comparison of strategies for resource-constrained mapping.

becomes large enough, many partial symmetries do not apply anymore, only *global* symmetries do (those already captured in the concept of the symmetry group).

A different perspective of the results can be obtained from Figure 17(a). It shows a series of subarchitecture configurations, where every point represents an evaluation for a different subarchitecture. All subarchitectures have four PEs, and the evaluations are sorted in descending order by the simulated execution time. Consider the first trial in Figure 17(a). All three strategies find similar mappings of $9.6 \cdot 10^{11}$ cycles. The inverse semigroups strategy finds the better mappings with $8.8 \cdot 10^{11}$ cycles after 50 trials, whereas the group-based strategy needs over 160. The simple strategy never finds these better mappings. Assuming this application had a requirement of running under $9 \cdot 10^{11}$ cycles, the simple strategy would need more than four PEs to achieve this. In fact, the simple strategy only found an application running under $9 \cdot 10^{11}$ cycles after considering six PEs. Figure 17(b) shows, for different runtime requirements, the number of PEs in the best subarchitecture found with the different strategies. For every execution time target (x -axis), all four methods were executed to find a mapping that satisfies the target time, using the least possible amount of resources. The relative sizes of the dots represent the relative number of trials to find that sub-architecture and mapping, while the y -axis shows the amount of resources (here, PEs) used, which indicates the quality of the result. In the example marked in the figure, with $6.5 \cdot 10^{11}$ cycles, instead of the 6 PEs required with the mappings found with the other strategies, the “simple” strategy requires 14 PEs. This shows how the quality of the solutions can be improved significantly with our proposed methods.

Even for the modestly sized Parallella architecture, the brute-force algorithm requires at least 10 times as many iterations as the other algorithms.

Finally, Figure 18 depicts concrete mappings from Figure 17(a). The processes and the FIFO buffers of the mjpeg application are depicted as a graph, embedded on the

Fig. 18. Examples of mappings of the mjpeg application to subarchitectures of size 4 of the Parallella, colored to help identify the processes. Symmetries become less evident with growing complexity.

Table II. Summary of Gains and Overheads from Symmetry in Our Use-Cases

Use-case	Benchmark	Relative overhead (avg.)		Relative time-reduction
		Design time	Runtime	
Genetic algorithms	sobel	≈ 0	$(0.11 \pm 0.04)\%$	$(2.38 \pm 2.05)\%$
	matmult	≈ 0	$(0.095 \pm 0.04)\%$	$(10.2 \pm 3.6)\%$
	mjpeg	≈ 0	$(0.14 \pm 0.04)\%$	$(0.71 \pm 0.24)\%$
	mandelbrot	≈ 0	$(0.21 \pm 0.07)\%$	$(0.1 \pm 0.2)\%$
	audio filter	≈ 0	$(0.12 \pm 0.05)\%$	$(2.96 \pm 3.73)\%$
Subarchitectures	groups	≈ 0	$3 \cdot 10^{-7}\%$	87.0%
	inv. semi	0.05%	$3 \cdot 10^{-5}\%$	89.6%

All relative numbers are normalized with the (average) baseline benchmarks without symmetries.

architecture. The location of the processes in the figure represents the PE to which they are mapped. The mapping labeled with “Result simple” in Figure 18 is precisely the single red dot from the “simple” strategy in Figure 17(a). The mapping labeled as “Best result” is one of the mappings with the lowest execution time found by the “inv. semi.” strategy in Figure 17(a). Notice that the communication seems to be sub-optimal, but from the critical path for this execution, in this particular benchmark, this had no effect: moving the “Best result” mapping to contiguous PEs yielded the same execution time in the simulation. The other two mappings show two subarchitectures, which were pruned by our method; the one labeled with “Pruned (equiv. to simple)” is equivalent to the first mapping, “simple,” that is, they both produced the same results, which was identified by our method. The situation is similar for the second and the last depicted mappings. Note how some transformations cannot be given simple geometrical interpretations, as “shapes.” While the “best” results are equivalent through a reflection, no simple geometrical interpretation can be given for the “simple” case. This figure serves to show concretely how the limits of the intuitive symmetry concepts can be quickly reached and why an automated approach like ours is necessary.

Finding the complete (generating set of the) inverse semigroup took about 30min on an Intel Xeon E5-4617 running at 2.90 GHz, with Algorithm 2, and using the 12 generators obtained, it took slightly over 3sec to calculate all representative sets. This time is negligible compared to the synthesis time of the 58,732 mappings for subarchitectures that were pruned, which would represent over 34 days of computation on the system used to evaluate this use-case. Additionally, this computation was done once and can be used with basically no overhead for any software synthesis to the same architecture.

As a final summary, Table II gives an overview of the gains from symmetry and overheads in the different benchmarks for the two use-cases. We see that, even for the

small architectures and applications of today, prototype implementations already yield significant improvements.

6. RELATED WORK

Besides our two use-cases, a large body of research exists that has addressed hardware and software models for design-space exploration. In particular, the mapping problem and diverse heuristics have been intensively studied before, as can be seen in Singh et al. [2013b]. Besides works that optimize for resources under runtime constraints or multiobjective metaheuristics, the use-cases studied in this article, several other different objectives for mapping have been studied, usually without considering symmetries. For example, the authors in Khdr et al. [2015] consider thermal-aware application mappings, inspired by dark silicon, but their approach does not leverage the problem's symmetries. Similarly, mappings aimed at achieving reliability via redundancy [Chen et al. 2016], robustness [Hempel et al. 2017], or motivated by security [Weichslgartner et al. 2016] have been studied as well.

In general, most authors consider the symmetries of the problem in an intuitive fashion [Roloff et al. 2015; Varghese et al. 2015; Thompson and Pimentel 2013; Kreutz et al. 2005; Hannig and Teich 2001; Singh et al. 2013a]. The prevailing theme in all these approaches is that the methods are not general. They are restricted to a certain architecture topology or class of topologies considered by the authors or particular to the application [Cohen 1988]. For example, the authors in Singh et al. [2013a] define a model for synthesizing NoC-based MPSoCs on reconfigurable hardware. In their model, they define an architecture composed of tile-types and with a property “hop_distance,” which they use to calculate communication costs in a heuristic by using the value of the largest hop_distance in a pessimistic fashion. While this model seems generic at first, it indeed refers only to a restricted family of architectures, since it includes several assumptions about the platform. It assumes the platform has only one basic communication resource and can be parameterized solely by the hop_distance and the types of tiles. This leaves out exotic NoC topologies, with clustering and heterogeneous memories, or whatever the future might hold. Additionally, while the heuristic might be useful in many use-cases, for those where it is not, the approach has to be redesigned from start to account for the symmetry in a more precise fashion. Similarly, in Thompson and Pimentel [2013], the authors explicitly design a method of removing the symmetry as presented in this article, albeit for a very particular use-case: fully symmetric homogeneous architectures. Finally, the approaches in Weichslgartner et al. [2016] and Schwarzer et al. [2017] do recognize and leverage some symmetries explicitly, namely those which are induced by geometric transformations (rotations, reflections, and translations). While this is a more explicit exposition of symmetries, it is limited to a very particular architecture family (regular NoC architectures) and will fail to find many symmetries, all those that do not correspond to this intuition; this is especially problematic when the notions of (Euclidean) geometric shapes break down outside of this very particular type of architecture.

On the other hand, works on high-level design-space exploration [Palermo et al. 2005; Liu et al. 2012; Quan and Pimentel 2014] do not consider symmetries in the design space per se, which implicitly implies that the designer has to come up with clever models for every particular problem. Research efforts have also concentrated directly in the reduction of the design space [Wang and Shan 2004]. However, these works achieve a reduction based on properties of the data, not with domain-specific knowledge obtained from application and hardware symmetries.

Group theory has many other applications outside of pure mathematics. It is ubiquitous in the fields of crystallography and particle physics, for example. In Goens and Castrillon [2015], we considered this issue and primitive version of this framework

based on group theory as well. The theory of inverse semigroups seems to find applications mostly in other fields of mathematics [Lawson 1998]. A similar approach, based on the abstract concept of groupoids, has found other applications, for example, in theoretical physics [Weinstein 1996] or cell biology [Stewart et al. 2003]. However, groupoids do not allow arbitrary multiplications between elements, which makes it complex to generalize most algorithms from group theory in an efficient fashion. We believe this makes our approach using inverse semigroups superior. To the best of our knowledge, this work is the first one to address the full symmetries of arbitrary applications and hardware architectures in system-level design in a structured and general fashion.

7. CONCLUSIONS AND OUTLOOK

In this article, we have shown how symmetries can be effectively leveraged in software synthesis. We introduced a rigorous approach using group theory, as well as its generalization in the form of inverse semigroups. We built a framework for automatically identifying symmetries and analyzed algorithms for doing so, as well as the potential for domain-specific optimizations.

As a proof of concept, we applied it to two examples of design-space exploration, concentrating our evaluation on current architectures and applications. Even for the modestly sized problems of today, we found that the methods provide an advantage concerning scalability by reducing the design space. The results yielded time-improvements in virtually all use-cases, with up to 10% time savings in one use-case, and a speedup of over $10\times$ in the other.

Our exploration of applying the methods of group and inverse semigroup theory to software synthesis revealed a very promising field of application. We have outlined diverse domain-specific improvements that could improve these methods. Additionally, the methods presented are generic and can thus be applied to a wide variety of problems in neighboring domains. For example, careful investigation of other programming models (e.g., OpenMP 4.X) can be used for identifying application symmetries in the proposed formalism and improving their execution. Similarly, these methods can be used on runtime systems for selecting equivalent variants of an execution, to accommodate multiple applications or deal with thermal issues.

ACKNOWLEDGMENTS

We thank Silexica for making their SLX Tool Suite available to us. We thank Andy Pimentel and Simon Polstra for help with Sesame and Akash Kumar, Alice Niemeyer, and Sebastian Posur for interesting discussions, which helped shape this work.

REFERENCES

- László Babai. 2015. Graph isomorphism in quasipolynomial time. *arXiv preprint arXiv:1512.03547* (2015).
- J. Balkind, M. McKeown, Y. Fu, T. Nguyen, Y. Zhou, A. Lavrov, M. Shahrad, A. Fuchs, S. Payne, X. Liang, M. Matl, and D. Wentzlaff. 2016. OpenPiton: An open source manycore research framework. In *Proceedings of the 21st International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'16)*. ACM, New York, NY, 217–232. DOI : <http://dx.doi.org/10.1145/2872362.2872414>
- Reimer Behrends, Kevin Hammond, Vladimir Janjic, Alexander Kononov, Steve Linton, Hans-Wolfgang Loidl, Patrick Maier, and Phil Trinder. 2016. HPC-GAP: Engineering a 21st-century high-performance computer algebra system. *Concurrency and Computation: Practice and Experience* 28 (2016), 3606–3636. DOI : <http://dx.doi.org/10.1002/cpe.3746>
- Eric Biscondi, Tom Flanagan, Frank Fruth, Zhihong Lin, and Filip Moerman. 2012. Maximizing Multi-core Efficiency with Navigator Runtime. White Paper. (Feb. 2012). Retrieved from www.ti.com/lit/wp/spry190/spry190.pdf.
- Wieb Bosma, John Cannon, and Catherine Playoust. 1997. The Magma algebra system: The user language. *Journal of Symbolic Computing* 24, 3–4 (1997), 235–265. DOI : <http://dx.doi.org/10.1006/jsc.1996.0125>

- Simone Casale Brunet, Marco Mattavelli, Claudio Alberti, and Jorn W Janneck. 2013. Design space exploration of high level stream programs on parallel architectures: A focus on the buffer size minimization and optimization problem. In *Proceedings of the 8th International Symposium on Image and Signal Processing and Analysis*.
- Jeronimo Castrillon, Rainer Leupers, and Gerd Ascheid. 2013. MAPS: Mapping concurrent dataflow applications to heterogeneous MPSoCs. *IEEE Transactions on Industrial Informatics* 9, 1 (Feb. 2013), 527–545. DOI : <http://dx.doi.org/10.1109/TII.2011.2173941>
- Jeronimo Castrillon, Weihua Sheng, and Rainer Leupers. 2011. Trends in embedded software synthesis. In *Proceedings of the International Conference on Embedded Computer Systems (SAMOS'11)*. IEEE, 347–354.
- Jeronimo Castrillon, Andreas Tretter, Rainer Leupers, and Gerd Ascheid. 2012. Communication-aware mapping of KPN applications onto heterogeneous MPSoCs. In *Proceedings of the 49th Annual Conference on Design Automation (DAC'12)*.
- Kuan-Hsun Chen, Jian-Jia Chen, Florian Kriebel, Semeen Rehman, Muhammad Shafique, and Jörg Henkel. 2016. Task mapping for redundant multithreading in multi-cores with reliability and performance heterogeneity. *IEEE Transactions on Computers* 65, 11 (2016), 3441–3455.
- Eric Cheung, Harry Hsieh, and Felice Balarin. 2007. Automatic buffer sizing for rate-constrained KPN applications on multiprocessor system-on-chip. In *Proceedings of the 2007 IEEE International High Level Design Validation and Test Workshop*. IEEE Computer Society, Washington, D.C., 37–44. DOI : <http://dx.doi.org/10.1109/HLDVT.2007.4392782>
- Paolo Codenotti, Hadi Katebi, Karem A. Sakallah, and Igor L. Markov. 2013. Conflict analysis and branching heuristics in the search for graph automorphisms. In *Proceedings of the 2013 IEEE 25th International Conference on Tools with Artificial Intelligence (ICTAI'13)*. IEEE, 907–914.
- Harvey A. Cohen. 1988. Symmetry considerations applied to hardware convolvers for image filtering. In *Proceedings of the 1988 IEEE International Conference on Systems, Man, and Cybernetics*, Vol. 2. IEEE, 1128–1131.
- Benoît Dupont de Dinechin, Renaud Ayrignac, Pierre-Edouard Beaucamps, Patrice Couvert, Benoit Ganne, Pierre Guironnet de Massas, François Jacquet, Samuel Jones, Nicolas Morey Chaisemartin, Frédéric Riss, and others. 2013. A clustered manycore processor architecture for embedded and accelerated applications. In *HPEC*. 1–6.
- K. Deb. 2001. *Multi-objective Optimization Using Evolutionary Algorithms*. Vol. 16. John Wiley & Sons.
- Marco Dorigo, Mauro Birattari, Christian Blum, Maurice Clerc, Thomas Stützle, and Alan Winfield. 2008. *Proceedings of the 5th International Conference on Ant Colony Optimization and Swarm Intelligence (ANTS'08)*. Vol. 5217. Springer.
- J. East, A. Egri-Nagy, J. D. Mitchell, and Y. Péresse. 2015. Computing finite semigroups. *arXiv preprint arXiv:1510.01868* (2015).
- C. Erbas, S. Cerav-Erbas, and A. D. Pimentel. 2006. Multiobjective optimization and evolutionary algorithms for the application mapping problem in multiprocessor system-on-chip design. *IEEE Transactions on Evolutionary Computation* 10, 3 (June 2006), 358–374. DOI : <http://dx.doi.org/10.1109/TEVC.2005.860766>
- Andrés Goens and Jeronimo Castrillon. 2015. Analysis of process traces for mapping dynamic KPN applications to MPSoCs. In *Proceedings of the IFIP International Embedded Systems Symposium (IESS)*. Foz do Iguaçu, Brazil.
- A. Goens, R. Khasanov, J. Castrillon, S. Polstra, and A. Pimentel. 2016. Why comparing system-level MPSoC mapping approaches is difficult: A case study. In *Proceedings of the IEEE 10th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc-16)*.
- Masaki Gondo, Fumio Arakawa, and Masato Eda. 2014. Establishing a standard interface between multi-manycore and software tools-SHIM. In *COOL Chips XVII, 2014 IEEE*. IEEE, 1–3.
- P. Greenhalgh. 2011. Big.LITTLE processing with arm cortex-a15 & cortex-a7. *ARM White Paper* (2011), 1–8.
- Linley Gwennap. 2011. Adapteva: More flops, less watts. *Microprocessor Report* 6, 13 (2011), 11–02.
- Frank Hannig and Jürgen Teich. 2001. Design space exploration for massively parallel processor arrays. In *Parallel Computing Technologies*. Springer, 51–65.
- G. Hempel, A. Goens, J. Asmus, J. Castrillon, and I. Sbalzarini. 2017. Robust mapping of process networks to many-core systems using bio-inspired design centering. In *Proceedings of the 20th International Workshop on Software and Compilers for Embedded Systems (SCOPES'17)*.
- Tommi Junttila and Petteri Kaski. 2011. Conflict propagation and component recursion for canonical labeling. In *Theory and Practice of Algorithms in (Computer) Systems*. Springer, 151–162.
- Gilles Kahn. 1974. The semantics of a simple language for parallel programming. In *Information Processing'74: Proceedings of the IFIP Congress*, Vol. 74. 471–475.

- J. Keinert, T. Schlichter, J. Falk, J. Gladigau, C. Haubelt, J. Teich, M. Meredith, and others. 2009. SystemCoDesigner- An automatic ESL synthesis approach by design space exploration and behavioral synthesis for streaming applications. *ACM TODAES* 14, 1 (2009), 1.
- Heba Khdr, Santiago Pagani, Muhammad Shafique, and Jörg Henkel. 2015. Thermal constrained resource management for mixed ILP-TLP workloads in dark silicon chips. In *Proceedings of the 52nd Annual Design Automation Conference*. ACM, 190–195.
- Márcio Kreutz, César A. Marcon, Luigi Carro, Flávio Wagner, and Altamiro A. Susin. 2005. Design space exploration comparing homogeneous and heterogeneous network-on-chip architectures. In *Proceedings of the 18th Annual Symposium on Integrated Circuits and System Design*. ACM, 190–195.
- Mark V. Lawson. 1998. *Inverse Semigroups: The Theory of Partial Symmetries*. World Scientific.
- E. Lee and D. Messerschmitt. 1987. Synchronous data flow. *Proceedings of the IEEE* 75, 9 (1987), 1235–1245.
- Hung-Yi Liu, Michele Petracca, and Luca P. Carloni. 2012. Compositional system-level design exploration with planning of high-level synthesis. In *Proceedings of the Conference on Design, Automation and Test in Europe*. EDA Consortium, 641–646.
- José Luis López-Presa, Antonio Fernández Anta, and Luis Núñez Chiroque. 2011. Conauto-2.0: Fast isomorphism testing and automorphism group computation. *arXiv preprint arXiv:1108.1060* (2011).
- Frank Lübeck and Max Neunhöffer. 2001. Enumerating large orbits and direct condensation. *Experimental Mathematics* 10, 2 (2001), 197–205.
- Brendan D. McKay and Adolfo Piperno. 2014. Practical graph isomorphism, II. *Journal of Symbolic Computation* 60, 0 (2014), 94–112. DOI: <http://dx.doi.org/10.1016/j.jsc.2013.09.003>
- J. D. Mitchell, M. Delgado, J. East, A. Egri-Nagy, N. Ham, J. Jonusas, M. Pfeiffer, B. Steinberg, J. Smith, M. Torpey, and W. Wilson. 2016. Semigroups, Version 2.8.0. Retrieved from <https://gap-packages.github.io/Semigroups>.
- Maximilian Odendahl, Jeronimo Castrillon, Vitaliy Volevach, Rainer Leupers, and Gerd Ascheid. 2013. Split-cost communication model for improved MPSoC application mapping. In *Proceedings of the 2013 International Symposium on System on Chip (SoC'13)*. IEEE, 1–8. DOI: <http://dx.doi.org/10.1109/ISSoC.2013.6675280>
- A. Olofsson, T. Nordström, and Z. Ul-Abdin. 2014. Kickstarting high-performance energy-efficient manycore architectures with Epiphany. In *2014 48th Asilomar Conference on Signals, Systems and Computers*. IEEE, 1719–1726.
- Gianluca Palermo, Cristina Silvano, and Vittorio Zaccaria. 2005. Multi-objective design space exploration of embedded systems. *Journal of Embedded Computing* 1, 3 (2005), 305–316.
- M. Pelcat, K. Desnos, L. Maggiani, Y. Liu, J. Heulot, J. F. Nezan, and S. Bhattacharyya. 2015. *Models of Architecture*. Research Report PREESM/2015-12TR01, 2015. IETR/INSA Rennes ; Scuola Superiore Sant'Anna, Pisa ; Institut Pascal, Clermont Ferrand ; University of Maryland, College Park ; Tampere University of Technology, Tampere. <https://hal.archives-ouvertes.fr/hal-01244470>
- A. Pimentel, C. Erbas, and S. Polstra. 2006. A systematic approach to exploring embedded system architectures at multiple abstraction levels. *IEEE Transactions on Computers* 55, 2 (2006), 99–112.
- W. Quan and A. Pimentel. 2014. Towards exploring vast MPSoC mapping design spaces using a bias-elitist evolutionary approach. In *Proceedings of the 2014 17th Euromicro Conference on DSD*. IEEE, 655–658.
- Wei Quan and Andy D. Pimentel. 2015. A hybrid task mapping algorithm for heterogeneous MPSoCs. *ACM Transactions on Embedded Computing Systems (TECS)* 14, 1 (2015), 14.
- Carl Ramey. 2011. TILE-Gx100 ManyCore Processor: Acceleration Interfaces and Architecture. Presented at HotChips 23. (Aug. 2011).
- Sascha Roloff, David Schafhauser, Frank Hannig, and Jürgen Teich. 2015. Execution-driven parallel simulation of PGAS applications on heterogeneous tiled architectures. In *Proceedings of the 2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. IEEE, 1–6.
- T. Schwarzer, A. Weichslgartner, M. Gla, S. Wildermann, P. Brand, and J. Teich. 2017. Symmetry-eliminating Design Space Exploration for Hybrid Application Mapping on Many-Core Architectures. Retrieved from <https://cris.fau.de/converis/publicweb/Publication/1061555>.
- Ákos Seress. 2003. *Permutation Group Algorithms*. Vol. 152. Cambridge University Press.
- Weihua Sheng, Artur Wiebe, Anastasia Stulova, Rainer Leupers, Bart Kienhuis, Johan Walters, and Gerd Ascheid. 2012. FIFO exploration in mapping streaming applications onto the TI OMAP3530 platform: Case study and optimizations. In *Proceedings of the 2012 IEEE 6th International Symposium on Embedded Multicore SoCs*. IEEE, 51–58.
- Software Solutions GmbH Silexica. 2016. SLXMapper. Retrieved from <http://www.silexica.com>.
- Charles C. Sims. 1970. Computational methods in the study of permutation groups. In *Computational Problems in Abstract Algebra*. 169–183.

- A. Singh, M. Shafique, A. Kumar, and J. Henkel. 2013b. Mapping on multi/many-core systems: Survey of current and emerging trends. In *Proceedings of the 50th Annual Design Automation Conference*. ACM, 1.
- Amit Kumar Singh, Akash Kumar, and Thambipillai Srikanthan. 2013a. Accelerating throughput-aware runtime mapping for heterogeneous MPSoCs. *ACM TODAES* 18, 1 (2013), 9.
- Ian Stewart, Martin Golubitsky, and Marcus Pivato. 2003. Symmetry groupoids and patterns of synchrony in coupled cell networks. *SIAM Journal on Applied Dynamical Systems* 2, 4 (2003), 609–646.
- The GAP Group 2016. *GAP – Groups, Algorithms, and Programming, Version 4.8.5*. The GAP Group. Retrieved from <http://www.gap-system.org>.
- Lothar Thiele, Iuliana Bacivarov, Wolfgang Haid, and Kai Huang. 2007. Mapping applications to tiled multiprocessor embedded systems. In *Proceedings of the 7th International Conference on Application of Concurrency to System Design (ACSD'07)*. IEEE Computer Society, Washington, D.C., 29–40. DOI : <http://dx.doi.org/10.1109/ACSD.2007.53>
- Mark Thompson and Andy D. Pimentel. 2013. Exploiting domain knowledge in system-level MPSoC design space exploration. *Journal of Systems Architecture* 59, 7 (2013), 351–360.
- Anish Varghese, Bob Edwards, Gaurav Mitra, and Alistair P Rendell. 2015. Programming the Adapteva Epiphany 64-core network-on-chip coprocessor. *International Journal of High Performance Computing Applications* (2015), 1094342015599238.
- G. Gary Wang and Songqing Shan. 2004. Design space reduction for multi-objective optimization and robust design optimization problems. *SAE SP* 113, 5 (2004), 37–46.
- A. Weichslgartner, S. Wildermann, J. Götzfried, Felix Freiling, M. Glaß, and J. Teich. 2016. Design-time/run-time mapping of security-critical applications in heterogeneous MPSoCs. In *Proceedings of the 19th International Workshop on Software and Compilers for Embedded Systems*. ACM, 153–162.
- A. Weinstein. 1996. Groupoids: Unifying internal and external symmetry. *Notices of the AMS* 43, 7 (1996), 744–752.

Received November 2016; revised May 2017; accepted May 2017