

MASTER THESIS



RADBOUD UNIVERSITY NIJMEGEN

SAC and generating Cuda code
Transforming dimensional data using pragma

Author:

Niek Janssen, s4297091

December 5, 2021

Contents

1	Context	2
1.1	SAC	2
1.2	CUDA and SAC	3
1.3	CUDA limits and solutions	4
1.4	Mental model	6
2	Formalization of mappings	7
2.1	Preliminary definitions	7
2.2	Index spaces and mappings	8
2.3	Proving mapping equivalence	9
2.4	ShiftLB	11
2.5	CompressGrid	12
2.6	PruneGrid	15
2.7	SplitLast	17
2.8	FoldLast2	18
2.9	Permute	20
2.10	Padlast	21
3	Implementation	24
3.1	The SAC compiler	24
3.2	Mapping execution	26
3.3	Mapping generation using strategies	26
4	Listings	28

Chapter 1

Context

1.1 SAC

SAC (Single assignment C) [1] is a functional programming language build around multidimensional arrays.

4 Each such an array consists of the array data and the shape of the array. The ‘data’ part contains the array elements, while the ‘shape’ part contains information about the dimensionality and the lengths of those dimensions. The ‘shape’ can be accessed through the 8 built-in shape function, and is a one-dimensional array itself.

Some examples and facts about SAC arrays:

```
1 // Declare array x
2 x = [[1, 2, 3], [4, 5, 6]];
3
4 // The shape of x can be requested
5 // with the shape function
6 print(shape(x));
7 // [2, 3]
8
9 // Note that the dimensionality of x
10 // is equal to the length of shape(x)
11 print(length(shape(x)));
12 // 2
13
14 // A new array can also be generated
15 // using the built-in function genarray.
16 // genarray takes a shape and a default
17 // element
18 print(genarray([2, 3], 42);
19 // [[42, 42, 42], [42, 42, 42]]
20
21 // Array elements can be accessed using
22 // the index vector:
23 iv = [1,1];
24 print(x[iv]);
25 // 5
```

The with-loop

At the heart of the SAC programming language is the with-loop. A with-loop can be compared with the ‘map’ function in the map-reduce computation model, 4 although the with-loop is a bit more advanced. A with-loop creates a new array using a given shape, and fills it with data in the manner defined in it’s body. A with loop consists of roughly three parts: 8

- A default array x . The shape of this array will be taken as the shape of the result. If an index vector iv exists for which the with-loop body does not define a value, the value $x[iv]$ will be used 12 instead.
- An index space S , defining which indexes in array x have to be replaced with new values.
- A body, containing an expression which is used 16 to compute the new value at index vector iv .

A pair of an index space and a body is called a “partition.” A with loop can have multiple partitions:

```
1 a = with {
2     (index space) : { body expression };
3     (index space) : { body expression };
4 } : default array;
```

An index space is defined using a tuple of four shape vectors (L , U , T , W), and a variable identifier. In the body expression, this variable contains the current index vector. The four shape vectors constrain what indexes are mapped in a partition. Using iv as variable identifier, we can define the constraints on the index space: 20

- Lowerbound L (optional): $L \leq iv$, restrict to indexes greater or equal to L 28
- Upperbound U : $iv < U$, restrict to indexes strictly smaller than U
- Step T and Width W (optional): $iv \% T < W$, Every T elements, take W elements. If W is omitted, 32 $W = 1$.

Combining this all together, we may create a with-loop

like this:

```

1  a = with {
2      ([0, 1] <= iv < [9, 8]
3          step [2, 3] width [1, 2])
4          : 3;
5      ([1, 0] <= iv < [8, 9]
6          step [3, 2] width [2, 1])
7          : 7;
8  } : genarray ([9, 9], 0);

```

Which creates a 2-dimensional array with the following output:

```

4  0 3 3 0 3 3 0 3 0
   2 0 2 0 2 0 2 0 2
   2 3 2 0 2 3 2 3 2
   0 0 0 0 0 0 0 0 0
8  2 3 2 0 2 3 2 3 2
   2 0 2 0 2 0 2 0 2
   0 3 3 0 3 3 0 3 0
   2 0 2 0 2 0 2 0 2
12 0 3 3 0 3 3 0 3 0

```

With-loop execution


We can conceptualize the execution of a partition in a with-loop in a few consequent steps. Firstly, we have an index space descriptor as a tuple of the four shape vectors L, U, T, W as defined above. Secondly, we have an index generator that computes all index vectors inside this index space descriptor. We call this set of index vectors an index space. Thirdly, we have the partition body, which transforms each index into a value. These values are then put in the resulting array on the correct positions. Of course, this is only an abstract way of describing a with loop. SAC will optimize the execution of the with-loop to use the least amount of CPU-cycles and memory.

Index space descriptor

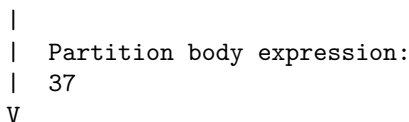
```

28 (lb, ub, step, width)

```



Index space
({[10, 10], [10, 11], ...})




Values

1.2 CUDA and SAC


The CUDA programming model is very similar to the programming model of a with-loop. Similarly to a with-loop, CUDA performs operations on elements of an index space. Where a with-loop uses a body with an expression, CUDA uses a function with arguments and a return value, but in effect a CUDA and a SAC operation uses the same concept: they execute a piece of code for each element of a certain index space.

In principle, because the programming models are so similar, it should be easy to utilize CUDA to execute with-loops on the GPU. Because the GPU architecture is in theory faster for these kinds of tasks than the CPU architecture, we should see significant performance improvements. We could just give the information about the index space to a CUDA index generator, wrap the expression body in a function, and let CUDA execute this function for each index vector.

Index space
(lb, ub, step, width)



Index space
({[10, 10], [10, 11], ...})



Values

The current implementation of the SAC compiler contains a mechanism to generate CUDA code, created by J. Guo [2]. This implementation follows the approach described above. Currently, roughly the following steps are implemented to make this work:

- Identify with-loops eligible for cudarizing: Not all with-loops can be implemented using cuda. CUDA cannot handle function calls to external functions for example, so all with-loops containing such function calls cannot be implemented using CUDA.
- Insert memory transfer primitives: GPU memory is separate from CPU memory, so necessary data has to be transferred to and from the GPU before and after the computations respectively.
- Create kernel functions: The creation of the functions wrapping the partition body expressions.

- The CUDA Index generator is given the index space information and it creates the kernel functions are run for each index in the index space.

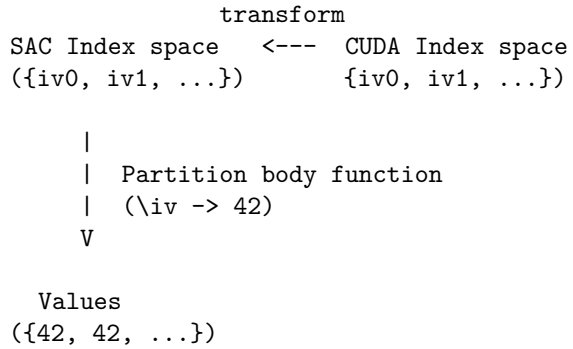

4 CUDA index spaces

A big problem with this approach is the limitations of the CUDA index spaces. The SAC with-loop index spaces have no constraints on dimensionality or dimension lengths other than they have to be finite. SAC index spaces also support a step and a width. CUDA index spaces for a GTX 1060 for example, can handle at most 6 dimensions, with the product of the lengths of the innermost 3 dimensions at most containing 1024 elements. Furthermore, these constraints can be different across the generations of hardware, and they do not support a step and width at all.

The current implementation deals with some of these issues, but not all of them. Furthermore, the current solution is very pragmatic and will have to be replaced with a more sophisticated mechanism with the flexibility to deal with index spaces of all shapes and sizes. In this thesis, we are going to design, implement and test this more sophisticated solution.

The general idea is that we are going to create transformations that transform the original SAC index space into an index space that fits into the CUDA architecture. Before we give the index space descriptor (the tuple of L, U, T, W) to the index generator, we will have to transform it to only an upperbound U that fits the CUDA index space requirements. Then, in the partition body function, we have to map the individual indexes of the transformed index space back to indexes from the original index space. After both transformations have been completed, exactly the same indexes should be generated as in the original CPU implementation of the index generator.

Because we have to accommodate for the broad spectrum of with-loops we have to support, it is nearly impossible to implement one single transformation to handle all cases. Instead, we are going to break the problem down into smaller pieces, and provide transformations for all of them. These transformations can then be combined together to fit each partition's needs.



Because with-loops come in many different shapes, sizes, steps and widths, we will need to have a very flexible way of defining these transformations. The easiest way to do that, is to create multiple mappings that can be combined together. In the next chapter, we will formalize the with-loop index generation and the mappings we can use to transform it.

1.3 CUDA limits and solutions

In this section, we will discuss the limitations and the mappings to handle these limitations.

CUDA index spaces do not support lowerbounds.

To handle this, we introduce the *ShiftLB* mapping. This mapping shifts the index space so that the lowerbound is a vector of 0's. In the partition body function, the lowerbound is added to each individual index again. The cost of this mapping is expected to be very low.

Example index space transformation with
L = [1, 1]:

00	01	02	03	04	05		11	12	13	14	15
10	11	12	13	14	15		21	22	23	24	25
20	21	22	23	24	25		31	32	33	34	35
30	31	32	33	34	35	---->	41	42	43	44	45
40	41	42	43	44	45		51	52	53	54	55
50	51	52	53	54	55						

CUDA index spaces do not support step/width.

To handle this, we introduce two alternative mappings: *PruneGrid* and *CompressGrid*. These mappings can only be executed on index spaces without a lowerbound.

PruneGrid leaves the upperbound intact, but checks in each partition body function call whether the current index vector is inside the grid. If it is, the function can continue its execution, but if it's not it is returned. This can potentially result in a lot of unnecessary GPU

threads, but other than that the overhead is expected to be low.

Example index space transformation with

4 $T = [3, 2], W = [2, 1]:$
 (*.. means the function for this index vector will be returned*)

```

00 01 02 03 04 05      00 01 .. 03 04 ..
8 10 11 12 13 14 15     .. .. .. .. .. ..
20 21 22 23 24 25     20 21 .. 23 24 ..
30 31 32 33 34 35 ----> .. .. .. .. .. ..
40 41 42 43 44 45     40 41 .. 43 44 ..
12 50 51 52 53 54 55     .. .. .. .. .. ..

```

The step/width mechanism creates regularly spaced n-dimensional blocks with indexes for which the body has to be executed, separated by indexes for which the body is *not* executed. *CompressGrid* compresses the index space by placing all n-dimensional blocks of threads to be executed directly beside each other. As a result, no unnecessary GPU threads are created, but the overhead of shifting those blocks may be a bit higher. Note that no actual data is moved, just the indexes are modified. Additionally, *CompressGrid* gets an extra argument C specifying which dimensions should be compressed. This way, it is possible to compress some dimensions while pruning the others.

Example index space transformation with

$T = [3, 2], W = [2, 1], C = [1, 1]:$

```

28 00 01 02 03 04 05      00 01 03 04
10 11 12 13 14 15      20 21 23 24
20 21 22 23 24 25      40 41 43 44
30 31 32 33 34 35 ---->
32 40 41 42 43 44 45
50 51 52 53 54 55

```

In general, *PruneGrid* should be used on high density index space dimensions, meaning the majority of function calls should be executed. On low density index space dimensions, *CompressGrid* should be executed.

CUDA index spaces support only a limited number of dimensions, with a limited length.

To handle this, we introduce two inverse mappings: the *SplitLast* and *FoldLast2*. These two mappings can best be used in combination with the *Permute* mapping introduced a bit further below. *These mappings can only be executed on index spaces without a lowerbound, step or width.*

SplitLast splits the last (innermost) dimension into two, increasing the dimensionality by one. The mapping takes the length of the innermost dimension as an

argument l. Note that the length of the original dimension may be increased to make sure the upperbound is a multiple of the given length. All function executions for those extra index vectors will be returned, so they will not be executed, so some overhead is expected. In the partition body function, the original index is recomputed from the indexes on the two new dimensions. The index vectors are placed in the two new dimensions using the second to last dimension as mayor and the last dimension as minor dimensions.

Example index space transformation with

$l = 6:$

(*.. means the function for this index vector will be returned*)

```

00 01 02 03 04 05 06 07 08 09
  |
  v
00 01 02 03 04 05
06 07 08 09 .. ..

```

FoldLast2 folds the last (innermost) two dimensions into one, decreasing the dimensionality by one. It is the inverse mapping to *SplitLast*. In the partition body function, the original two indexes are recomputed from the index on the new dimension. The index vectors are placed in the new dimension using the original second to last dimension as mayor and the original last dimension as minor dimensions.

Example index space transformation with

$l = 6:$

```

00 01 02 03 04 05
10 11 12 13 14 15
  |
  v
00 01 02 03 04 05 10 11 12 13 14 15

```

CUDA index spaces spawn threads in warps of 32.

Having the total number of index vectors something else than a multiple of 32 will slow down the hardware thread creation process significantly. Because of this, we want to pad the number of threads to a multiple of 32. The *PadLast* mapping handles this for us. Note that the number 32 may differ per GPU device.

PadLast takes one parameter p. It adds indexes to the end of the last (innermost) dimension until its length is a multiple of p. In the partition body function, all threads added here are returned again. Unless this mapping is used to generate loads of extra indexes, the overhead should be reasonably small.

Example index space transformation with

p = 7:
 (... means the function for this index vector will be returned)

```

4 00 01 02 03 04 05      00 01 02 03 04 05 ..
  10 11 12 13 14 15      10 11 12 13 14 15 ..
  20 21 22 23 24 25      20 21 22 23 24 25 ..
  30 31 32 33 34 35 ----> 30 31 32 33 34 35 ..
8 40 41 42 43 44 45      40 41 42 43 44 45 ..
  50 51 52 53 54 55      50 51 52 53 54 55 ..

```

CUDA performs caching.

The GPU will be fastest when the indexes are ordered in a way that neighboring data is accessed by neighboring threads. Because SAC is a functional language, it may replace some operations on mappings with accessor functions. This means that the order of the dimensions in the data may be different to the order of the dimensions in the index space. *Permute* can be used to rectify that. Additionally, *Permute* can be used to setup the index space for a *SplitLast*, *FoldLast2* or *PadLast*.

Permute takes the permutation as argument P. It changes the order of dimensions by permuting L, U, T and W. In the partition body function, it changes the order of the dimensions of the iv back to the original permutation. The overhead of this mapping is exactly 0, as this can be fully handled in the compiler. No code is generated for this mapping.

Example index space transformation with

p = 7:

```

  00 01 02 03 04 05      00 10 20 30 40 50
  10 11 12 13 14 15      01 11 21 31 41 51
32 20 21 22 23 24 25      02 12 22 32 42 52
  30 31 32 33 34 35 ----> 03 13 23 33 43 53
  40 41 42 43 44 45      04 14 24 34 44 54
  50 51 52 53 54 55      05 15 25 35 45 55

```

1.4 Mental model

TODO: change or delete?

To help understanding what's going on, we will create a simple mental model of the steps we need to implement this. The mental model follows the structure of the figure above.

- On the left hand side, we have the SAC index-space and indexes
- On the right hand side, we have the CUDA thread-space and indexes
- When executing a mapped partition, we perform the following steps:

- Firstly, we map the index space itself, moving right in our mental model
- Secondly, we have CUDA generate the threads and indexes, moving down in our mental model
- Thirdly, we have to map the generated indexes back to their SAC counterparts, moving left in our mental model
- Lastly, we execute the with-loop body. We will not discuss this in further detail, so it is not included in our mental model.

Chapter 2

Formalization of mappings

2.1 Preliminary definitions

Before we jump into the definitions of index spaces and their mappings, we will first create some preliminary definitions and notations.

Definition 2.1 (Preliminary assumptions). (preliminaries)

- $\mathbb{B} = \{\top, \perp\}$ is the set of boolean values.
- \mathbb{N} is the set of natural numbers, including 0
- A^n denotes a vector of type A of length $n \in \mathbb{N}$.
- Usually, we distinguish capital variables from lowercase variables:
 - If we talk about capital U, L, \dots , they denote the *vectors*
 - If we talk about U_i, L_i, \dots , they denote the value of a vector in dimension i .
 - If we talk about lower u, f, \dots , they denote either a normal value not being a vector, or the variable is used as a shorthand for a vector variable when the dimension is apparent from context.
- If $f : A \rightarrow B$ is a function, then for any n we can apply that function on vectors $A^n \rightarrow B^n$ by applying it to each dimension of the vector:

$$f(I) = I', \text{ with } \forall d \in [0..n-1], I'_d = f(I_d)$$
- If $f : A \rightarrow \mathbb{B}$ is a function, then for any n , $f^\downarrow : A^n \rightarrow \mathbb{B}$ we will apply that same operation to each dimension of the vector. The result is then \top iff all results are \top :

$$f^\downarrow(N) = \begin{cases} \top & | \quad f(A) = \top^n \\ \perp & | \quad \textit{otherwise} \end{cases}$$

We will mostly work with vectors of natural or whole numbers. In line with the preliminary assumptions, these types are denoted as \mathbb{N}^n and \mathbb{Z}^n , for some $n \in \mathbb{N}$.

Definition 2.2 (Posets using product order). (poset)

We define two partial orderings (posets) on the set of all vectors of type \mathbb{Z}^n .

The first poset we define is the product order [3]. In our context, the product order for two vectors $I, I' \in \mathbb{N}^n$ will be defined as the poset (\mathbb{Z}^n, \leq) with for all $I, I' \in \mathbb{Z}^n$:

$$I \leq I' \iff \forall d \in [0..n-1], I_d \leq I'_d$$

The required proofs, such as transitivity, reflexivity and antisymmetry are given in [3].

The second poset we define is the poset $(\mathbb{Z}^n, \leq^\downarrow)$ using the definition of \leq^\downarrow as described in definition 2.1. We still have to prove $(\mathbb{Z}^n, \leq^\downarrow)$ is a valid poset, by showing the properties *reflexivity*, *antisymmetry* and *transitivity* on it.

Proof. reflexivity: $I \leq^\downarrow I$. By definition 2.1, this is equal to $\forall d \in [0..n-1], I_d \leq I_d$. This is always true, as for every $x \in \mathbb{Z}$ holds $x \leq x$.

antisymmetry: if $I \leq^\downarrow I'$ and $I' \leq^\downarrow I$, then $I = I'$. We prove this by contradiction. Let's state that $I \neq I'$ and $I \leq I'$. Note that for $I \neq I'$ to hold, we need $n \geq 1$. This means that $\forall d \in [0..n-1], I_d < I'_d$. This contradicts the assumption that $I' \leq^\downarrow I$. \cancel{f}

transitivity: if $I \leq^\downarrow I''$ and $I'' \leq^\downarrow I'$ then $I \leq^\downarrow I'$. We know that $\forall d \in [0..n-1], I_d \leq I''_d \leq I'_d \rightarrow I_d \leq I'_d$, as it is the case for all $x, y, z \in \mathbb{Z}$. From these assumptions combined, we can conclude that $I \leq^\downarrow I'$. \blacksquare

Note. There is a difference between $I < I'$ and $I <^\downarrow I'$. $[0, 1] < [1, 1]$, but $[0, 1] \not<^\downarrow [1, 1]$. The same difference exists between $I \leq I' \iff I \leq^\downarrow I'$.

Note. $\mathbb{N}^n \subset \mathbb{Z}^n$, so this partial order works on natural numbers too.

2.2 Index spaces and mappings

Let us start with some definitions. Note that some definitions are further specified in definition 2.4.

Definition 2.3 (Index space transformations). (index-space-transform)

- Let $S \in \mathbb{S}$ be an index space. The exact definition of S and \mathbb{S} will be discussed in definition 2.4.
- Let $*S \in *\mathbb{S}$ be an index space descriptor. An index space descriptor is a tuple of values describing what values should be in the index space. The exact definition of $*S$ and $*\mathbb{S}$ will be discussed in definition 2.4.
- Let $I \in S$ be an index in index space S respectively.
- Let $gen : *\mathbb{S} \rightarrow \mathbb{S}$ be the index space generator function. It maps the index space descriptor to an index space.
- Let $tf_{\rightarrow}^s : *\mathbb{S} \rightarrow *\mathbb{S}$ be an index space mapping. It maps an index space descriptor from the SAC side (left) more towards the CUDA side (right).
- Let $tf_{\leftarrow}^i : S' \rightarrow S$ be an index mapping. It maps an index from the CUDA side (right) more towards the SAC side (left). The domain and codomain of tf_{\leftarrow}^i are dependent on the index space descriptor of the codomain and the gen and tf_{\rightarrow}^s functions: $*S' = tf_{\rightarrow}^s(*S)$, $S = gen(*S)$ and $S' = gen(*S')$.
- A partition mapping is a pair of $(tf_{\rightarrow}^s, tf_{\leftarrow}^i)$. In normal with-loop execution, we will first call tf_{\rightarrow}^s on the index space descriptor, then generate the index space using gen and finally transform all indexes back to the initial index space. The result of these function calls should be equal to calling gen directly:

$$gen(*S) = \{tf_{\leftarrow}^i(I) \mid I \in gen(tf_{\rightarrow}^s(*S))\}$$

Before we can formalize the definition of an index space further, we have to discuss the requirements. We already know that an index space descriptor is a set of integer vectors with the same length, defining the lowerbound, upperbound, step and width values. tf_{\rightarrow}^s can change these parameters when transforming an index space descriptor, as long as tf_{\leftarrow}^i can transform the resulting index space back into the original index space.

Because we need to recompute the original left-hand indexes from the mapped right-hand indexes, the mapped index space can never be smaller than the original index space on the left. However, some mappings may introduce extra elements in an index

space. An example of such a mapping is *PadLast*. *PadLast* $_{\rightarrow}^s$ will purposefully increase the size of the index space by increasing the upperbound. However, the extra partition body function calls for these index vectors should not actually be executed. We introduce the \perp element as a way to remove an index vector from an index space: a tf_{\leftarrow}^i can discard a partition body function call by mapping it's index to \perp . Because no indexes can be created out of thin air, we can assume that $tf_{\leftarrow}^i(\perp)$ is always equal to \perp as well.

We can now formally define our index spaces. We will differentiate between three types of index spaces. Grid index spaces, normalized index spaces, and dense index spaces. Grid index spaces are index spaces defined by the lowerbound, upperbound, step and width. Normalized index spaces are a subset of grid index spaces without a lowerbound, or where the lowerbound is zero. Dense index spaces are a subset of normalized index spaces without a step and width, or where the step equals the width.

Definition 2.4 (Index spaces). (index-space)

- Index space $S^n \in \mathbb{S}^n := \mathcal{P}(\mathbb{N}^n \cup \{\perp\})$, with $n \in \mathbb{N}$ and \mathcal{P} the power set.
- Grid index space $*S_g^n := (L, U, T, W) \in *\mathbb{S}_g^n \subset \mathbb{S}^n$
 - $L, U, T, W \in \mathbb{N}$, $L \leq U$ and $W \leq T$
 - $gen(*S_g^n) = \{I \mid I \in S^n \wedge (I = \perp \vee (L \leq I < \downarrow U \wedge (I - L) \overrightarrow{\%}_0 T < \downarrow W))\}$
- Normalized index sp. $S_n^n := (U, T, W) \in \mathbb{S}_n^n \subset \mathbb{S}_g^n$
 - $gen(*S_n^n) = gen((0, U, T, W))$
- Dense index space $S_d^n := (U) \in \mathbb{S}_d^n \subset \mathbb{S}_n^n$
 - $gen(*S_d^n) = gen((0, U, 1, 1))$
- For all $tf_{\leftarrow}^i : tf_{\leftarrow}^i(\perp) = \perp$

Lemma 2.5 (*gen* of normalized and dense index spaces). (*gen-norm-dense*)

Because *gen* of normalized and dense index spaces are just specialized versions of the *gen* of grid index spaces, we can simplify their formula's a bit.

- $gen(*S_n^n) = \{I \mid I \in S^n \wedge (I = \perp \vee (0 \leq I < \downarrow U \wedge (I - 0) \overrightarrow{\%}_0 T < \downarrow W))\}$
 \iff
 $gen(*S_n^n) = \{I \mid I \in S^n \wedge (I = \perp \vee (I < \downarrow U \wedge I \overrightarrow{\%}_0 T < \downarrow W))\}$
- $gen(*S_d^n) = \{I \mid I \in S^n \wedge (I = \perp \vee (0 \leq I < \downarrow U \wedge (I - 0) \overrightarrow{\%}_0 1 < \downarrow 1))\}$
 \iff

$$\text{gen}(*S^n) = \{I \mid I \in S^n \wedge (I = \perp \vee I <^{\downarrow} U)\}$$

Mapping correctness

As stated before, transforming the index space right and transforming the generated indexes back left should be equivalent to generating the indexes on the original index space:

$$\forall S^n \in \mathbb{S}^n : \text{gen}(S^n) = \{tf_{\leftarrow}^i(I) \mid I \in \text{gen}(tf_{\rightarrow}^s(S^n))\}$$

However, this predicate is not entirely accurate. Because sets cannot contain duplicate elements, there may be more than one element I in $S^{n'}$ with $tf_{\leftarrow}^i(I) = I$ with $I \in S^n$. When executing a with-loop partition, we do not want two partition body function calls with the same index. The one exception to this statement is \perp . Index invocations with index \perp won't be executed anyway, so there may be as many as we like.

To formalize this difference, we will distinguish between the *operative subset* and the *excess subset* of an index space.

Definition 2.6 (Operative and excess subsets). (operative-subset)

- The operative subset of an index space $\mathcal{O}(S^n)$ is the index space without \perp .

$$\mathcal{O}(S^n) = S^n \setminus \{\perp\}$$

- The excess subset of an index space $\mathcal{E}(S^n)$ is $\{\perp\}$.

$$\mathcal{E}(S^n) = \{\perp\}$$

- The operative subset of an index space S^n after a partition mapping ($\mathcal{O}_{tf}(S^n)$) is the subset of $S^{n'}$ that gets mapped to the operative subset of S^n by tf_{\leftarrow}^i .

$$\mathcal{O}_{tf}(S^n) = \{i \mid tf_{\leftarrow}^i(i) \in \mathcal{O}(S^n) \wedge i \in S^{n'}\}$$

- The excess subset of an index space S after a partition mapping ($\mathcal{E}_{tf}(S^n)$) is the subset of $S^{n'}$ that tf_{\leftarrow}^i maps to the excess subset of S^n , e.g. that tf_{\leftarrow}^i maps to \perp .

$$\mathcal{E}_{tf}(S^n) = \{i \mid tf_{\leftarrow}^i(i) = \perp \wedge i \in S^{n'}\}$$

Note that, by definition, for any well-defined tf and S :

- $\mathcal{O}(S^n) \cup \mathcal{E}(S^n) = S^n$
- $\mathcal{O}_{tf}(S^n) \cup \mathcal{E}_{tf}(S^n) = S^{n'}$
- $\mathcal{O}(S^n) \cap \mathcal{E}(S^n) = \emptyset$
- $\mathcal{O}_{tf}(S^n) \cap \mathcal{E}_{tf}(S^n) = \emptyset$

We can now fully formalize the requirements of a partition mapping. We will call these requirements *mapping equivalence*.

Definition 2.7 (Mapping equivalence). (mapping-equivalence)

A partition mapping tf is mapping equivalent iff:

- tf_{\rightarrow}^s is well-defined
 - tf_{\leftarrow}^i is well-defined
 - tf_{\leftarrow}^i is operational bijective:
- $$tf_{\leftarrow}^i : \mathcal{O}_{tf}(S^n) \rightarrow \mathcal{O}(S^n) \text{ is a bijection}$$

2.3 Proving mapping equivalence

In proving mapping equivalence for the different mappings, we will oftentimes use similar very similar techniques or lemma's. In this section we will specify a few of those in lemma's and definitions. We also talks about general approaches for some situations.

Notation of domains and codomains

Whenever we are discussing a mapping, either as a single transformation or as a pair of transformations, we define the following properties over the domains and codomains:

- The right hand side of the mapping (the codomain of a tf_{\rightarrow}^s and the domain of a tf_{\leftarrow}^i) will have the same variables as the left hand side of the mapping, with a prime ($'$) added. These can be S and \mathbb{S} and variations, I, L, U, T, W and lowercase variants, and any other helper variables we define in the process.
- The left hand side of the mapping (the domain of a tf_{\rightarrow}^s and the codomain of a tf_{\leftarrow}^i) will have the non-prime version of variables.
- For every $*S$, we assume there to be an S with $S = \text{gen}(*S)$. Likewise, for every S we assume there to be a $*S$.

Well-definedness of tf_{\rightarrow}^s

Lemma 2.8 (Well-definedness of tf_{\rightarrow}^s). (*tf-wds*)

For any $tf_{\rightarrow}^s : *S \rightarrow *S'$ we will have to prove well-definedness by proving that for any $*S \in *S : tf_{\rightarrow}^s(*S) \in *S'$. To see exactly what this entails, we have to expand the definition and constraints of $S \in \mathbb{S}$. For $*S_g^{n'} \in *S_g^n$, we have to prove that:

- $L', U', T', W' \in \mathbb{N}^n$
- $L' \leq U'$
- $T' \leq W'$

By definition of ${}^*S_n^n$ and ${}^*S_d^n$, we have to substitute $L = 0$ and $T = 1, W = 1$ respectively. This means we can simplify the proofs for any ${}^*S_n^n \in {}^*S_n^n$ and ${}^*S_d^n \in {}^*S_d^n$.

- For ${}^*S_n^n \in {}^*S_n^n$ it is sufficient to prove:
 - $U', T', W' \in \mathbb{N}^n$
 - $T' \leq W'$

We can omit $0 \leq U'$, as this holds for any $U' \in \mathbb{N}^n$.

- For ${}^*S_d^n \in {}^*S_d^n$ it is sufficient to prove:
 - $U' \in \mathbb{N}^n$

Well-definedness of tf_{\leftarrow}^i

Lemma 2.9 (Well-definedness of tf_{\leftarrow}^i). (tf-wdi)

For any $tf_{\leftarrow}^i : S' \rightarrow S$ we will have to prove well-definedness by proving that for any $I' \in S' : tf_{\leftarrow}^i(I') \in S$. To see what exactly this means, we have to expand the definition of gen for that particular index space. For any tf_{\leftarrow}^i , we have to prove:

Given I' with $I' = \perp \vee (L' \leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W')$,

Then $I = \perp \vee (L \leq I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W)$, with $I = tf_{\leftarrow}^i(I')$ should hold.

Note that this automatically holds for $I' = \perp$, because by definition, $tf_{\leftarrow}^i(\perp) = \perp$ for any tf_{\leftarrow}^i . Additionally, we can use the simplifications from lemma 2.5 to make simpler proofs for codomains in S_n^n and S_d^n :

- For $tf_{\leftarrow}^i : S_g^{n'} \rightarrow S_n^n$:

Given I' with $L' \leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W'$,

Then $I = \perp \vee (L \leq I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W)$, with $I = tf_{\leftarrow}^i(I')$ should hold.

- For $tf_{\leftarrow}^i : S_g^{n'} \rightarrow S_n^n$:

Given I' with $L' \leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W'$,

Then $I = \perp \vee (I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W)$, with $I = tf_{\leftarrow}^i(I')$ should hold.

- For $tf_{\leftarrow}^i : S_g^{n'} \rightarrow S_n^n$:

Given I' with $L' \leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W'$,

Then $I = \perp \vee I < \downarrow U$, with $I = tf_{\leftarrow}^i(I')$ should hold.

Operational bijectivity of tf_{\leftarrow}^i

Lemma 2.10 (Operational bijectivity of tf_{\leftarrow}^i). (tf-opbi)

For any well-defined $tf_{\leftarrow}^i : S' \rightarrow S$ we need prove that $tf_{\leftarrow}^i : \mathcal{O}_{tf}(S) \rightarrow \mathcal{O}(S)$ is a bijection. We do this by showing **(A)** surjectivity and **(B)** injectivity.

(A) For surjectivity, we need to prove that for each element I in $\mathcal{O}(S)$, there is an element I' in $\mathcal{O}_{tf}(S)$ with $tf_{\leftarrow}^i(I') = I$:

Given an $I \in \mathcal{O}(S)$:
 $\exists I' \in \mathcal{O}_{tf}(S), tf_{\leftarrow}^i(I') = I$

However, by definition, we know that $\mathcal{O}_{tf}(S)$ is exactly the subset of S' with $tf_{\leftarrow}^i(I') \in S$. Because of this, it is equivalent to prove \S :

Given an $I \in \mathcal{O}(S)$:
 $\exists I' \in S', tf_{\leftarrow}^i(I') = I$

(B) For injectivity, we need to prove that if two indexes $I', I'' \in \mathcal{O}_{tf}(S)$ both map to the same element I in S , then they are the same element:

Given an $I', I'' \in \mathcal{O}_{tf}(S)$:
 $tf_{\leftarrow}^i(I') = tf_{\leftarrow}^i(I'') \rightarrow I' = I''$

Again, we can make use of the definition of $\mathcal{O}_{tf}(S)$ and rewrite the proposition as:

Given an $I \in \mathcal{O}(S)$ and $I', I'' \in S'$:
 $tf_{\leftarrow}^i(I') = I \wedge tf_{\leftarrow}^i(I'') = I \rightarrow I' = I''$

Separating dimensions

In many cases, it is easier to talk about or prove something for one single dimension instead of for all dimensions at once. This is possible when the dimensions are independent from each other in the current context. To enable us to easily talk about those individual dimensions, we may *separate dimensions*.

Definition 2.11 (Separating dimensions). (sep-dim)

When separating dimensions, we prove or show a property on a single abstract dimension. If we can prove it for any dimension, we know it has to be true for all dimensions.

Whenever we separate dimensions, we will implicitly introduce the following variables:

- The variable d contains the current abstract dimension we are talking about
- The variables $i, l, u, t, w, i', l', u', t', w'$ are the values of their capital counterparts in dimension d
- $tf_{\leftarrow}^i(I)$ are the values of $tf_{\leftarrow}^i(I)$ in dimension d ,

- $tf_{\leftarrow d}^i(i)$ are the values of $tf_{\leftarrow}^i(I)$ in dimension d , iff this function is only dependent on the values on dimension d ,

4 **Definition 2.12** (Dimension dependency). (dim-dep)

In the context of a certain mapping, two dimensions d and d' with $d \neq d'$ of their index spaces can be either dependent or independent. Which one of them is applicable can be derived from the definitions of tf_{\rightarrow}^s and tf_{\leftarrow}^i .

For some dimension d holds:

- If the definition of tf_{\rightarrow}^s or tf_{\leftarrow}^i for dimension d possibly refers to some $L_{d'}, U_{d'}, T_{d'}, W_{d'}, I_{d'}$ with $d \neq d'$, then dimensions d and d' are dependent.
- The opposite is automatically also true: if the definition of tf_{\rightarrow}^s or tf_{\leftarrow}^i for some dimension d' possibly refers to L_d, U_d, T_d, W_d, I_d with $d \neq d'$, dimensions d and d' are dependent.
- If dimensions d and d' are not dependent, they are dependent.
- If there is no d' so that d and d' are dependent, we speak of d as being independent.

Lemma 2.13 (Proofs with separated dimensions). (sep-dim-prf)

24 *In the proofs outlined by lemma's 2.8, 2.9 and ??, it is always possible to separate dimensions as in definition 2.11. The properties in question are:*

- Well definedness of tf_{\rightarrow}^s
- Well definedness of tf_{\leftarrow}^i
- Operational injectivity of tf_{\leftarrow}^i
- Operational surjectivity of tf_{\leftarrow}^i

32 *Proof.* For the well definedness of tf_{\rightarrow}^s and tf_{\leftarrow}^i , this is trivial. Both the definition of S_g^n and $I \in S_g^n$ are defined so that the dimensions are completely independent of eachother.

36 For the operational injectivity and surjectivity of tf_{\leftarrow}^i , this follows from how $I \in S_g^n$ is defined. It is defined as $\forall d \in [0..n-1] : I_d \in S_{g,d}^n \iff I \in S_g^n$. This means that if you can prove for all dimensions that they are surjective and injective, you can prove that tf_{\leftarrow}^i is surjective and injective. Of course, some extra steps may be required if the dimensions are not independent of eachother in tf_{\leftarrow}^i . ■

(Partially) Identity mappings

44 In some cases, mappings leave certain dimensions completely intact, behaving like an identity mapping for those dimensions. To avoid repeating the same proof over and over again, we will prove mapping equivalence for such an identity dimension.

Lemma 2.14 (Mapping equivalence for identity dimensions). (tf-idd)

For any dimension d that is separated using definition 2.11 and $tf_{\rightarrow d}^s((l, u, s, w)) = (l, u, s, w), tf_{\leftarrow d}^i(i) = i$, tf is mapping equivalent.

Proof. Using definition 2.7 we now have to prove **(A)** tf_{\rightarrow}^s is well defined, **(B)** tf_{\leftarrow}^i is well defined and **(C)** tf_{\leftarrow}^i is operationally bijective.

(A) We know that $l, u, s, w \in \mathbb{N}$. By definition, $tf_{\rightarrow}^s((l, u, s, w)) = (l, u, s, w)$. Using lemma §2.8, we know that tf_{\rightarrow}^s is well defined.

(B) We know that $tf_{\rightarrow}^s(*S_g^n) = *S_g^n$. We also know that $tf_{\leftarrow}^i(i) = i$. This makes tf_{\leftarrow}^i an identity function, which is well defined.

(C) Using lemma 2.10, we have to prove **(C.1)** $\exists i' \in S_g^n, tf_{\leftarrow d}^i(i') = i$ for any $i \in S_g^n$ *indimension* d , and **(C.2)** $tf_{\leftarrow}^i(i') = i \wedge tf_{\leftarrow}^i(i'') = i \rightarrow i' = i''$ for any $i, i', i'' \in S_g^n$.

(C.1) By definition, if $i' = i$ then $tf_{\leftarrow d}^i(i') = i$.

(C.2) By definition, $tf_{\leftarrow d}^i(i') = i \iff i' = i$. This also holds for i'' . Now we can prove $i' = i = i'' \implies i' = i''$. ■

2.4 ShiftLB

ShiftLB removes the lowerbound of an index space by shifting the index space so that the lowerbound becomes 0. This means that the mapping transforms any grid index space to a normalized index space.

Definition 2.15 (*ShiftLB*). (def-shiftlb)

- $ShiftLB_{\rightarrow}^s : *S_g^n \rightarrow *S_g^n$

$ShiftLB_{\rightarrow}^s((L, U, T, W)) = (U', T', W')$ with:

- $U' = U - L$
- $T' = T$
- $W' = W$

- $ShiftLB_{\leftarrow}^i : S_n^{n'} \rightarrow S_n^n$

$ShiftLB_{\leftarrow}^i(I) = I + L$

Lemma 2.16 (*ShiftLB*_→^s is well-defined). (shift-wds)

Given any $*S_g^n \in *S_g^n$. We now have to prove that $ShiftLB_{\rightarrow}^s(*S_g^n) \in S_n^n$.

Proof. Using lemma 2.8 for $*S_n^n$, we have to prove **(A)** $U', T', W' \in \mathbb{N}$, **(B)** $W' \leq T'$.

(A) As $L \leq U$ holds, we know that $U - L$ is a vector of natural numbers ≥ 0 . This means that $U' \in \mathbb{N}$.

$T' \in \mathbb{N}^n$ and $W' \in \mathbb{N}^n$ holds too, as $T \in \mathbb{N}^n$ and $W \in \mathbb{N}^n$.

ⓑ $W' \leq T'$ holds, as $W \leq T$ holds.

4 **Lemma 2.17** (*ShiftLB $_{\leftarrow}^i$ is well-defined*). (*shift-wdi*)

Given any $I' \in S_n^{n'}$. We now have to prove that $ShiftLB_{\leftarrow}^i(I') \in S_g^n$.

Proof. Using lemma 2.9 for codomain S_g^n , we can now prove for any given I' and $I = ShiftLB_{\leftarrow}^i(I')$:

$$\begin{aligned} L' &\leq I' <\downarrow U' \wedge (I' - L') \% T' <\downarrow W' \rightarrow \\ I &= \perp \vee (L \leq I <\downarrow U \wedge (I - L) \% T <\downarrow W) \end{aligned}$$

\iff

$$\begin{aligned} 0 &\leq I' <\downarrow U - L \wedge (I' - 0) \% T <\downarrow W \rightarrow \\ L &\leq (I' + L) <\downarrow U \wedge (I' + L - L) \% T <\downarrow W \end{aligned}$$

Inserted definitions of *shiftLB*, pruned left or clause

\iff

$$\begin{aligned} 0 &\leq I' <\downarrow U - L \wedge I' \% T <\downarrow W \rightarrow \\ L &\leq (I' + L) <\downarrow U \wedge I' \% T <\downarrow W \end{aligned}$$

Some simplifications of formulas

\iff

$$\begin{aligned} 0 &\leq I' <\downarrow U - L \wedge I' \% T <\downarrow W \rightarrow \\ L &\leq (I' + L) <\downarrow U \end{aligned}$$

Right 'and' clause is true from our assumption, subtracted L from formula

\iff

$$\begin{aligned} 0 &\leq I' <\downarrow U - L \wedge I' \% T <\downarrow W \rightarrow \\ 0 &\leq I' <\downarrow U - L \end{aligned}$$

28 **Lemma 2.18** (*ShiftLB $_{\leftarrow}^i$ is operational bijective*). (*shift-opbi*)

Take the operational parts of $S_g^n : \mathcal{O}(S_g^n)$ and $S_n^{n'} : \mathcal{O}_{ShiftLB_{\leftarrow}^i}(S_g^n)$. We now have to prove that $ShiftLB_{\leftarrow}^i : \mathcal{O}_{ShiftLB_{\leftarrow}^i}(S_g^n) \rightarrow \mathcal{O}(S_g^n)$ is bijective.

Proof. Using lemma 2.10, we now have to prove ⓐ surjectivity and ⓑ injectivity.

ⓐ Given $I \in \mathcal{O}(S_g^n)$. We now have to prove that there is some $I' \in S_n^{n'}$ with $tfi(I') = I$.

We define this I' as $I - L$. We now have to prove that ⓑ $tff_{\leftarrow}^i(I') = I$ and ⓐ $I' \in S_n^{n'}$.

$$\begin{aligned} \textcircled{1} &tff_{\leftarrow}^i(I') \\ &= I' + L \\ &= I - L + L \\ &= I \end{aligned}$$

ⓐ We know that $I \in \mathcal{O}(S_g^n)$. We can now prove:

$$\begin{aligned} L &\leq I <\downarrow U \wedge (I - L) \% T <\downarrow W \rightarrow \\ L' &\leq I' <\downarrow U' \wedge (I' - L') \% T' <\downarrow W' \end{aligned}$$

\iff

$$\begin{aligned} L &\leq I <\downarrow U \wedge (I - L) \% T <\downarrow W \rightarrow \\ 0 &\leq (I - L) <\downarrow (U - L) \wedge (I - L - 0) \% T <\downarrow W \end{aligned}$$

Inserted definitions of prime values

\iff

$$\begin{aligned} L &\leq I <\downarrow U \wedge (I - L) \% T <\downarrow W \rightarrow \\ L &\leq I <\downarrow U \wedge (I - L) \% T <\downarrow W \end{aligned}$$

Right 'and' clause is true from our assumption, added L to formula

ⓑ Given $I \in \mathcal{O}(S_g^n)$ and $I', I'' \in S_n^{n'}$. We now have to prove that if $tff_{\leftarrow}^i(I') = I$ and $tff_{\leftarrow}^i(I'') = I$, then $I' = I''$. We prove this by contradiction. Assume there are two $I' \neq I''$ with $tff_{\leftarrow}^i(I') = I$ and $tff_{\leftarrow}^i(I'') = I$:

$$I' \neq I'' \wedge tff_{\leftarrow}^i(I') = I \wedge tff_{\leftarrow}^i(I'') = I$$

\iff

$$I' \neq I'' \wedge I' + L = I \wedge I'' + L = I$$

\iff

$$I' \neq I'' \wedge I' = I - L \wedge I'' = I - L$$

\iff

$$I' \neq I'' \wedge I' = I'' \quad \blacksquare$$

Theorem 2.19 (*ShiftLB is mapping-equivalent*). (*shift-me*)

By definition 2.7, *ShiftLB* is mapping equivalent iff:

- ⓐ *ShiftLB $_{\rightarrow}^s$ is well defined*
- ⓑ *ShiftLB $_{\leftarrow}^i$ is well defined*
- ⓒ *ShiftLB $_{\leftarrow}^i$ is operationally bijective*

Proof. ⓐ is proven in lemma 2.16

ⓑ is proven in lemma 2.17

ⓒ is proven in lemma 2.18 ■

2.5 CompressGrid

CompressGrid removes the step and width from a normalized index space by removing the space between the indexes in the index space. This means that the mapping transforms any normalized index space into a dense index space.

CompressGrid takes an extra argument $C \in \mathbb{B}^n$, defining which dimensions should be compressed.

Definition 2.20 (*CompressGrid*). (*def-compress*)

- $CompressGrid_{\rightarrow}^s : \mathbb{B}^n \rightarrow *S_n^n \rightarrow *S_n^n$

$$CompressGrid_{\rightarrow}^s(C)(U, T, W) = (U', T', W')$$

For definitions of U', T', W' we separate dimensions as in definition 2.11.

$$u' = \begin{cases} \left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) & | c = \top \\ u & | c = \perp \end{cases}$$

$$t' = \begin{cases} 1 & | c = \top \\ t & | c = \perp \end{cases}$$

$$w' = \begin{cases} 1 & | c = \top \\ w & | c = \perp \end{cases}$$

- $CompressGrid_{\leftarrow}^i : S_d^{n'} \rightarrow S_n^n$

$$CompressGrid_{\leftarrow}^i(I') = I:$$

For the definition of I we separate dimensions as in definition 2.11.

$$i = \begin{cases} \left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w & | c = \top \\ i' & | c = \perp \end{cases}$$

Lemma 2.21 ($CompressGrid_{\rightarrow}^s$ is well defined). (*compress-wds*)

Given any $*S_n^n \in *S_n^n$. We now have to prove that $CompressGrid_{\rightarrow}^s(*S_n^n) \in *S_n^n$.

Proof. Using lemma 2.8 for $*S_d^n$, we now only have to prove that $U' \in \mathbb{N}^n$. We separate dimensions as in definition 2.11. We also distinct two cases: $c = \top$ and $c = \perp$.

- $c = \top$:

$$u, t, w \in \mathbb{N} \rightarrow \left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \in \mathbb{N}$$

\iff

$$u, t, w \in \mathbb{N} \rightarrow \left\lfloor \frac{u}{t} \right\rfloor, u, t, w \in \mathbb{N}$$

For u, t, w , this immediately holds from the assumption. For $\left\lfloor \frac{u}{t} \right\rfloor$ this holds as well, as the floor makes the result an integer again. As u and t are at least zero, $\frac{u}{t}$ is at least zero as well.

- $c = \perp$:

We know $u, t, w \in \mathbb{N}$. We also know $u' = u, t' = t, w' = w$. So, we can conclude that $u', t', w' \in \mathbb{N}$ as well.

Lemma 2.22 ($CompressGrid_{\leftarrow}^i$ is well defined). (*compress-wdi*)

Given any $I' \in S_d^{n'}$. We now have to prove that $CompressGrid_{\leftarrow}^i(I') \in S_n^n$.

Proof. Using lemma ?? for codomain S_n^n , we can now prove for any given I' and $I = CompressGrid_{\leftarrow}^i(I')$: $I' \in S_d^{n'} \rightarrow I \in S_n^n$. We separate dimensions as in definition 2.11.

$$\begin{aligned} i' < u' \wedge i' \% t' < w' &\rightarrow \\ i = \perp \vee (i < u \wedge i \% t < w) & \end{aligned}$$

\iff

$$\begin{aligned} i' < u' \wedge i' \% t' < w' &\rightarrow \\ i < u \wedge i \% t < w & \end{aligned}$$

Prune left hand side of the 'or'

We now distinct two cases: $c = \top$ and $c = \perp$.

- $c = \top$:

$$\begin{aligned} i' < u' \wedge i' \% t' < w' &\rightarrow \\ i < u \wedge i \% t < w & \end{aligned}$$

\iff

$$\begin{aligned} i' < \left(\left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \right) \wedge i' \% 1 < 1 &\rightarrow \\ \left(\left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w \right) < u \wedge & \\ \left(\left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w \right) \% t < w & \end{aligned}$$

We now split the proof into the left hand side and the right hand side of the 'and.'

- Left hand side:

$$i' < \left(\left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \right) \rightarrow$$

$$\left(\left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w \right) < u$$

\iff

$$i' < \left(\left\lfloor \frac{u - u \% t}{t} \right\rfloor * w + \min(u \% t, w) \right) \rightarrow$$

$$\left\lfloor \frac{i'}{w} \right\rfloor * t \leq u - u \% t \wedge i' \% w < u \% t$$

$$\text{Because } \left\lfloor \frac{u}{t} \right\rfloor = \left\lfloor \frac{u - u \% t}{t} \right\rfloor \text{ and } ab < c +$$

$$d \leftarrow a \leq c \wedge b < d$$

\iff

$$i' - i' \% w \leq \left\lfloor \frac{u - u \% t}{t} \right\rfloor * w \wedge i' \% w < \min(u \% t, w) \rightarrow \left\lfloor \frac{i' - i' \% w}{w} \right\rfloor * t \leq u - u \% t \wedge i' \% w < u \% t$$

Similarly, because $\left\lfloor \frac{i'}{t} \right\rfloor = \left\lfloor \frac{i' - i' \% t}{t} \right\rfloor$ and $ab < c + d \leftarrow a \leq c \wedge b < d$

Now, we can split this proof in two:

$$* i' - i' \% w \leq \left\lfloor \frac{u - u \% t}{t} \right\rfloor * w \rightarrow \left\lfloor \frac{i' - i' \% w}{w} \right\rfloor * t \leq u - u \% t$$

\iff

$$i' - i' \% w \leq \left\lfloor \frac{u - u \% t}{t} \right\rfloor * w \rightarrow$$

$$i' - i' \% w \leq \frac{u - u \% t}{t} * w$$

Divide by t , multiply by w

\iff

$$i' - i' \% w \leq \left\lfloor \frac{u - u \% t}{t} \right\rfloor * w \rightarrow$$

$$\left\lfloor \frac{i' - i' \% w}{w} \right\rfloor \leq \frac{u - u \% t}{t}$$

Divide by t

$$\text{This holds, as } \frac{u - u \% t}{t} = \left\lfloor \frac{u - u \% t}{t} \right\rfloor.$$

$$* i' \% w < \min(u \% t, w) \rightarrow i' \% w < u \% t$$

\iff

$$i' \% w < u \% t \rightarrow i' \% w < u \% t$$

– Right hand side:

$$i' < \left(\left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \right) \wedge i' \% 1 < 1 \rightarrow$$

$$\left(\left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w \right) \% t < w$$

\iff

$$i' < \left(\left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \right) \wedge i' \% 1 < 1 \rightarrow$$

$$\left(\left\lfloor \frac{i'}{w} \right\rfloor * t \right) \% t + (i' \% w) \% t < w$$

Because $a \% t + b \% t < w \rightarrow (a+b) \% t < w$

\iff

$$i' < \left(\left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w) \right) \wedge i' \% 1 < 1 \rightarrow 0 + (i' \% w) \% t < w$$

Because $(x * t) \% t = 0$, if $x \in \mathbb{N}$

This holds, by definition of $\%$.

• $c = \perp$:

$$i' < u' \wedge i' \% t' < w' \rightarrow i < u \wedge i \% t < w$$

\iff

$$i' < u \wedge i' \% t < w \rightarrow i < u \wedge i' \% t < w$$

Expand definitions of $\text{CompressGrid}_{\rightarrow}^s$ for u', t', w' and $\text{CompressGrid}_{\leftarrow}^i$ for i

■

Lemma 2.23 ($\text{CompressGrid}_{\leftarrow}^i$ is operational bijective). (*compress-opbi*)

Take the operational parts of $S_n^n : \mathcal{O}(S_n^n)$ and $S_n^{n'} : \mathcal{O}_{\text{CompressGrid}_{\leftarrow}^i}(S_n^n)$. We now have to prove that $\text{CompressGrid}_{\leftarrow}^i : \mathcal{O}_{\text{CompressGrid}_{\leftarrow}^i}(S_n^n) \rightarrow \mathcal{O}(S_n^n)$ is bijective.

Proof. We separate dimensions as in definition 2.11. If $c = \perp$, we have an identity mapping and we can use lemma 2.14. For the rest of the proof, we can now assume that $c = \top$.

Using lemma 2.10, we now have to prove $\textcircled{\text{A}}$ surjectivity and $\textcircled{\text{B}}$ injectivity.

$\textcircled{\text{A}}$ Given an $i \in \mathcal{O}(S_n^n)_d$ and an $i' = \left\lfloor \frac{i}{t} \right\rfloor * w + i \% t$. We

now have to prove that $\textcircled{\text{A.1}}$ $\text{CompressGrid}_{\leftarrow d}^i(i') = i$ and $\textcircled{\text{A.2}}$ $i' \in S_n^{n'}_d$.

$\textcircled{\text{A.1}}$ $\text{CompressGrid}_{\leftarrow d}^i(i')$

$$= \left\lfloor \frac{\left\lfloor \frac{i}{t} \right\rfloor * w + i \% t}{w} \right\rfloor * t + i' \% w$$

Expanded definitions of $\text{CompressGrid}_{\leftarrow}^i$ and i'

$$= \left\lfloor \frac{\left\lfloor \frac{i}{t} \right\rfloor * w}{w} \right\rfloor * t + i' \% w$$

$\lfloor x \rfloor * w$ is a multiple of w , and by definition of $i \in S_n^n$, $i \% t < w$, so we can remove the $i \% t$.

$$= \left\lfloor \frac{i}{t} \right\rfloor * t + i' \% w$$

$$= \left\lfloor \frac{i}{t} \right\rfloor * t + \left(\left\lfloor \frac{i}{t} \right\rfloor * w + i \% t \right) \% w$$

Expanded definition of i'

$$= \left\lfloor \frac{i}{t} \right\rfloor * t + (i \% t) \% w$$

4 $\lfloor x \rfloor * w$ is a multiple of w

$$= \left\lfloor \frac{i}{t} \right\rfloor * t + i \% t$$

By definition of $i \in S_n^n$, $i \% t < w$

$$= i$$

8 (A.2) To prove:

$$i < u \wedge i \% t < w \rightarrow i' < u' \wedge i' \% t' < w'$$

We will now assume $i < u \wedge i \% t < w$ to be true, and prove the right hand side of the implication in two parts: $i' < u'$, and $i' \% t' < w'$.

- $i' < u'$

\iff

$$\left\lfloor \frac{i}{t} \right\rfloor * w + i \% t < \left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w)$$

16 Expand definitions of i' and u'

\iff

We will split into three cases:

$$- \left\lfloor \frac{i}{t} \right\rfloor * w > \left\lfloor \frac{u}{t} \right\rfloor * w$$

20 This case is not possible, as by definition $i < u$.

$$- \left\lfloor \frac{i}{t} \right\rfloor * w = \left\lfloor \frac{u}{t} \right\rfloor * w$$

24 Because of this equality and the fact that $i < u$, we know that $i \% t < u \% t$. We also know that $i \% t < w$. So we can conclude that $i \% t < \min(u \% t, w)$, and thus that:

$$\left\lfloor \frac{i}{t} \right\rfloor * w + i \% t < \left\lfloor \frac{u}{t} \right\rfloor * w + \min(u \% t, w)$$

$$- \left\lfloor \frac{i}{t} \right\rfloor * w < \left\lfloor \frac{u}{t} \right\rfloor * w$$

Because of this equality, we know that the difference between the two is at least w , so it is sufficient to prove that:

$$32 i \% t < \min(u \% t, w) + w$$

This is trivially true, as by definition $i \% t < w$ is true.

- $i' \% t' < w'$

36 By definition, $t' = 1$ and $w' = 1$, so this true for any i' .

(B) Given $I \in S_n^n$. We now have to prove that if $\text{CompressGrid}_{\leftarrow d}^i(i') = i$ and $\text{CompressGrid}_{\leftarrow d}^i(i'') = i$, then $i = i''$. We prove this by contradiction. Assume there are two $i' \neq i''$ with $\text{CompressGrid}_{\leftarrow d}^i(i') = i$ and $\text{CompressGrid}_{\leftarrow d}^i(i'') = i$:

$$i' \neq i'' \wedge \text{compressgrid}_d(i') = i \wedge \text{CompressGrid}_{\leftarrow d}^i(i'') = i$$

\iff

$$i' \neq i'' \wedge \left\lfloor \frac{i'}{w} \right\rfloor * t + i' \% w = i \wedge \left\lfloor \frac{i''}{w} \right\rfloor * t + i'' \% w = i$$

Let's try to take $i' = i''$, and try to add/subtract to i' so that $i' \neq i''$:

- If you want the left hand side of the addition to stay the same, the maximum you can add/subtract to i' is $w - 1$
- If you want the right hand side of the addition to stay the same, you'd have to add/subtract a multiple of t . Remember that $t \geq w$.

This means that it's impossible to change the value of i' with either the left hand side or the right hand side of the addition to change. Furthermore, it is also impossible to change both sides and get them to even out:

- The left hand side of the addition only changes with multiples of w
- The right hand side of the addition changes with a maximum of $w - 1$, as by definition $i \% t < w$.

Theorem 2.24 (CompressGrid is mapping-equivalent). (*compress-me*)

By definition 2.7, CompressGrid is mapping equivalent iff:

- (A) $\text{CompressGrid}_{\rightarrow}^*$ is well defined
- (B) $\text{CompressGrid}_{\leftarrow}^*$ is well defined
- (C) $\text{CompressGrid}_{\leftarrow}^*$ is operationally bijective

Proof. (A) is proven in lemma 2.21

(B) is proven in lemma 2.22

(C) is proven in lemma 2.23

2.6 PruneGrid

PruneGrid removes the step and width from a normalized index space by mapping all indexes outside of the grid to \perp . This means that the mapping transforms any normalized index space into a dense index space.

Definition 2.25 (*PruneGrid*). (def-prune) ■

- $PruneGrid_{\rightarrow}^s : *S_n^n \rightarrow *S_d^n$

$$PruneGrid_{\rightarrow}^s(U, T, W) = (U') \text{ with:}$$

- $U' = U$

- $PruneGrid_{\leftarrow}^i : S_d^{n'} \rightarrow S_n^n$

$$PruneGrid_{\leftarrow}^i(I') = \begin{cases} I' & | & I' \% T <^{\downarrow} W \\ \perp & | & I' \% T \not<^{\downarrow} W \end{cases}$$

Lemma 2.26 (*PruneGrid*_→^s is well defined). (*prune-wds*)

Given any $*S_n^n \in *S_n^n$. We now have to prove that $PruneLB_{\rightarrow}^s(*S_n^n) \in *S_d^n$.

Proof. Using lemma 2.8, we just have to prove $U' \in \mathbb{N}^n$. We know $U \in \mathbb{N}^n$. We also know $U' = U$. So, we can conclude that $U' \in \mathbb{N}^n$ as well. ■

Lemma 2.27 (*PruneGrid*_←ⁱ is well defined). (*prune-wdi*)

Given any $I \in S_d^{n'}$. We now have to prove that $PruneGrid_{\leftarrow}^i(I') \in S_n^n$.

Proof. Using lemma 2.9, we can now prove for any given I' and $I = PruneGrid_{\leftarrow}^i(I')$:

$$\begin{aligned} I' <^{\downarrow} U' \wedge I' \xrightarrow{\%} T' <^{\downarrow} W' \rightarrow \\ I = \perp \vee (I <^{\downarrow} U \wedge I \% T <^{\downarrow} W) \end{aligned}$$

We now distinct two cases:

- $I' \% T <^{\downarrow} W$:

$$\begin{aligned} I' <^{\downarrow} U \wedge I' \xrightarrow{\%} T <^{\downarrow} W \rightarrow \\ I = \perp \vee (I' <^{\downarrow} U \wedge I' \% T <^{\downarrow} W) \end{aligned}$$

Insert definitions of $PruneGrid_{\rightarrow}^s$ and $PruneGrid_{\leftarrow}^i$

$$\iff$$

$$\begin{aligned} I' <^{\downarrow} U \wedge I' \xrightarrow{\%} T <^{\downarrow} W \rightarrow \\ I' <^{\downarrow} U \wedge I' \% T <^{\downarrow} W \end{aligned}$$

Choose right ‘or’ clause

- $I' \% T \not<^{\downarrow} W$:

$$\perp = \perp \vee (\perp <^{\downarrow} U \wedge \perp \% T <^{\downarrow} W)$$

Insert definition of $PruneGrid_{\leftarrow}^i$, prune assumption

$$\iff$$

$$\perp = \perp$$

Choose left ‘or’ clause

Lemma 2.28 (*PruneGrid*_←ⁱ is operational bijective). (*prune-opbi*)

Take the operational parts of $S_n^n : \mathcal{O}(S_n^n)$ and $S_d^{n'} : \mathcal{O}_{PruneGrid_{\leftarrow}^i}(S_n^n)$. We now have to prove that $PruneGrid_{\leftarrow}^i : \mathcal{O}_{PruneGrid_{\leftarrow}^i}(S_n^n) \rightarrow \mathcal{O}(S_n^n)$ is bijective.

Proof. Using lemma 2.10, we now have to prove **(A)** surjectivity and **(B)** injectivity. 8

(A) Given an $I \in \mathcal{O}(S_n^n)$. We take $I' = I$. We now have to prove that **(A.1)** $PruneGrid_{\leftarrow}^i(I') = I$ and **(A.2)** $I' \in S_d^{n'}$. 12

We know that $I \in \mathcal{O}(S_n^n)$. This means, by definition, that $I \leq U \wedge I \% T <^{\downarrow} W$ holds. From this, we can conclude that:

$$\text{(A.1)} \quad PruneGrid_{\leftarrow}^i(I') = I' = I$$

$$\text{(A.2)} \quad U' = U \implies I' <^{\downarrow} U' \iff I' \in S_d^{n'}$$

(B) Given $I \in \mathcal{O}(S_n^n)$ and $I', I'' \in S_d^{n'}$. We now have to prove that if $PruneGrid_{\leftarrow}^i(I') = I$ and $PruneGrid_{\leftarrow}^i(I'') = I$, then $I' = I''$. We prove this by contradiction. Assume there are two $I' \neq I''$ with $PruneGrid(I') = I$ and $PruneGrid_{\leftarrow}^i(I'') = I$:

$$I' \neq I'' \wedge PruneGrid(I') = I \wedge PruneGrid_{\leftarrow}^i(I'') = I$$

$$\iff$$

$I' \neq I'' \wedge I' = I \wedge I'' = I$ By definition of $PruneGrid_{\leftarrow}^i$. $PruneGrid_{\leftarrow}^i(\perp) = \perp$, and $I \in \mathcal{O}(S_n^n)$, so I' and I'' cannot be \perp . 28

This is a contradiction by transitivity of $=$. ■

Theorem 2.29 (*PruneGrid* is mapping-equivalent). (*prune-me*)

By definition 2.7, *PruneGrid* is mapping equivalent iff: 32

- **(A)** $PruneGrid_{\rightarrow}^s$ is well defined
- **(B)** $PruneGrid_{\leftarrow}^i$ is well defined
- **(C)** $PruneGrid_{\leftarrow}^i$ is operationally bijective

Proof. **(A)** is proven in lemma 2.26 36

(B) is proven in lemma 2.27

(C) is proven in lemma 2.28 ■

2.7 SplitLast

SplitLast increases the dimensionality by splitting the last dimension in two. This means that the product of the lengths of the two new dimensions is equal to the length of the original dimension.

SplitLast takes an extra argument $l \in \mathbb{N}$. We assume U_M to be a multiple of l and $l > 0$.

Definition 2.30 (*SplitLast*). (def-split)

- We define the major and minor dimensions to be: $M = n - 1$ the major dimension $m = n$ the minor dimension

$$SplitLast_{\rightarrow}^s : \mathbb{N} \rightarrow *S_d^n \rightarrow *S_d^{n+1'}$$

$$SplitLast_{\rightarrow}^s(l)(U) = U'$$

For the definition of U' we separate dimensions as in definition 2.11.

$$u' = \begin{cases} u & | d < M \\ \frac{U_M}{l} & | d = M \\ l & | d = m \end{cases}$$

$$SplitLast_{\leftarrow}^i : S_d^{n+1'} \rightarrow S_d^n$$

$$SplitLast_{\leftarrow}^i(l)(I') = I$$

For the definition of I we separate dimensions as in definition 2.11.

$$i = \begin{cases} i' & | d < M \\ l * I'_M + I'_m & | d = M \end{cases}$$

Lemma 2.31 (*SplitLast_{\rightarrow}^s* is well defined). (*split-wds*)

Given any $*S_d^n \in S_d^n$. We now have to prove that $SplitLast_{\rightarrow}^s(*S_d^n) \in *S_d^{n+1'}$.

Proof. Using lemma 2.8 for S_d^n , we now only have to prove that $U' \in \mathbb{N}$. We will separate dimensions as in definition 2.11.

We distinct three cases:

- $d < M$

In this case, $u' = u$, and we know that $u \in \mathbb{N}$, as $*S_d^n \in S_d^n$.

- $d = M$

In this case, $u' = \frac{U_M}{l}$. By definition, U_M is a multiple of l , so $\frac{U_M}{l} \in \mathbb{N}$.

- $d = m$

In this case, $u' = l$, which by definition comes from \mathbb{N} . ■

Lemma 2.32 (*SplitLast_{\leftarrow}^i* is well defined). (*split-wdi*)

Given any $I' \in *S_d^{n+1'}$. We now have to prove that $SplitLast_{\leftarrow}^i(I') \in S_d^n$.

Proof. Using lemma 2.9 for codomain S_d^n , we can now prove for any given I' and $I = SplitLast_{\leftarrow}^i(I')$: $I' \in S_d^{n+1'} \rightarrow I \in S_d^n$. We separate dimensions as in definition 2.11\$. For $d < M$, we have an identity mapping and we can use lemma 2.14\$. For the rest of the proof, we can now assume that $d = M$.

$$I' < U' \rightarrow i < u$$

\Leftarrow

$$I'_M < U'_M \wedge I'_m < U'_m \rightarrow i < u$$

We are only interested in dimensions M and m

\Leftrightarrow

$$I'_M < \frac{u}{l} \wedge I'_m < l \rightarrow l * I'_M + I'_m < u$$

Expand definitions of i and U'

\Leftrightarrow

$$I'_M \leq \frac{u}{l} - 1 \wedge I'_m < l \rightarrow l * I'_M + I'_m < u$$

As all parts are integers, we can rewrite the leftmost comparison

\Leftrightarrow

$$l * I'_M \leq u - l \wedge I'_m < l \rightarrow l * I'_M + I'_m < u$$

Multiply leftmost comparison by l

This is true, as $l * I'_M + I'_m$ combined is less then $u - l + l$. ■

Lemma 2.33 (*SplitLast_{\leftarrow}^i* is operational bijective). (*split-opbi*)

Take the operational parts of $S_d^n : \mathcal{O}(S_d^n)$ and $S_d^{n+1'} : \mathcal{O}_{SplitLast_{\leftarrow}^i}(S_d^n)$. We now have to prove that $SplitLast_{\leftarrow}^i : \mathcal{O}_{SplitLast_{\leftarrow}^i}(S_d^n) \rightarrow \mathcal{O}(S_d^n)$ is bijective.

Proof. We separate dimensions as in definition 2.11. If $d < M$, we have an identity mapping and we will use lemma 2.14. For the rest of this proof, we can now assume that $d = M$.

Using lemma ??, we now have to prove \textcircled{A} surjectivity and \textcircled{B} injectivity.

Ⓐ Given an $i \in \mathcal{O}(S_d^n)_M$, $I'_M = \left\lfloor \frac{i}{l} \right\rfloor$, $I'_m = i \% l$ and $I'_d = I_d$ for any $d < M$. We now have to prove that

Ⓐ.1 $SplitLast^i_{\leftarrow M}(I') = i$ and Ⓐ.2 $I' \in S_d^{n+1'}$.

4 Ⓐ.1 $SplitLast^i_{\leftarrow M}(I')$

$$= l * \left\lfloor \frac{i}{l} \right\rfloor + i \% l$$

Expanded definitions of $SplitLast^i_{\leftarrow}$ and I'

$$= (i - i \% l) + i \% l$$

8 Cancelled out the multiplication and division of l , reimplemented the flooring as a modulo operation

$$= i$$

12 Ⓐ.2 For dimensions $d < M$ this is trivial, as this is the identity mapping. We now only have to discuss dimensions M and m .

- Dimension M : $I_M < U_M \rightarrow I'_M < U'_M$

16 We assume $I_M < U_M$ to be true, and prove $I'_M < U'_M$.

$$I'_M < U'_M$$

\iff

$$20 \left\lfloor \frac{I_M}{l} \right\rfloor < \frac{U_M}{l}$$

Expanded definitions of I'_M and U'_M

\iff

$$\frac{I_M}{l} < \frac{U_M}{l}$$

$$24 \text{ Because } \left\lfloor \frac{I_M}{l} \right\rfloor < \frac{I_M}{l}$$

$$I_M < U_M$$

$$\text{As } l > 0$$

Ⓑ Given an $I \in S_d^n$ and $I', I'' \in S_d^{n+1'}$. We now have to prove that if $SplitLast^i_{\leftarrow d}(I') = i$ and $SplitLast^i_{\leftarrow d}(I'') = i$, then $I'_M = I''_M$ and $I'_m = I''_m$. We prove this by contradiction. Assume there are I', I'' with $I'_M \neq I''_M$ or $I'_m \neq I''_m$, $SplitLast^i_{\leftarrow d}(I') = i$ and $SplitLast^i_{\leftarrow d}(I'') = i$:

$$(I'_M \neq I''_M \vee I'_m \neq I''_m) \wedge SplitLast^i_{\leftarrow d}(I') = i \wedge SplitLast^i_{\leftarrow d}(I'') = i$$

\iff

$$36 (I'_M \neq I''_M \vee I'_m \neq I''_m) \wedge l * I'_M + I'_m = i \wedge l * I''_M + I''_m = i$$

Expanded definition of $SplitLast^i_{\leftarrow}$

Let's compute I''_M and I''_m from I'_M and I'_m . If we want to have $l * I''_M + I''_m$ retain its value, we have to either increase $l * I''_M$ and decrease I''_m by the same amount, or the other way around. However, by definition,

$0 \leq I''_m < l$, and $l * I''_M$ can only be increased or decreased by multiples of l . Because of this, $I''_M = I'_M$ and $I''_m = I'_m$. ■

Theorem 2.34 ($SplitLast$ is mapping-equivalent). (split-me)

By definition 2.7, $SplitLast$ is mapping equivalent iff:

- Ⓐ $SplitLast^s_{\rightarrow}$ is well defined
- Ⓑ $SplitLast^i_{\leftarrow}$ is well defined
- Ⓒ $SplitLast^i_{\leftarrow}$ is operationally bijective

Proof. Ⓐ is proven in lemma 2.31

Ⓑ is proven in lemma 2.32

Ⓒ is proven in lemma 2.33 ■

2.8 FoldLast2

$FoldLast2$ is the inverse of $SplitLast$. It decreases the dimensionality by folding the last two dimensions together. This means that the length of the new dimension is equal to the product of the lengths of the last two original dimensions.

Definition 2.35 ($FoldLast2$). (def-fold)

- We define the major and minor dimensions to be: $M = n - 2$ the major dimension and $m = n - 1$ the minor dimension.

$$FoldLast2^s_{\rightarrow} : *S_d^n \rightarrow *S_d^{n-1'}$$

$$FoldLast2^s_{\rightarrow}(U) = U'$$

For the definition of U' we separate dimensions as in definition 2.11.

$$u' = \begin{cases} u & | d < M \\ U_M * U_m & | d = M \end{cases}$$

$$FoldLast2^i_{\leftarrow} : S_d^{n-1'} \rightarrow S_d^n$$

$$FoldLast2^i_{\leftarrow}(I) = I$$

For the definition of I we separate dimensions as in definition 2.11.

$$i = \begin{cases} i' & | d < M \\ \left\lfloor \frac{i'}{U_m} \right\rfloor & | d = M \\ i' \% U_m & | d = m \end{cases}$$

Lemma 2.36 ($FoldLast2_{\rightarrow}^s$ is well defined). (*fold-wds*)

Given any $*S_d^n \in \mathbb{S}_d^n$. We now have to prove that
 $SplitLast_{\rightarrow}(*S_d^n) \in *S_d^{n-1'}$.

Proof. Using lemma 2.8 for S_d^n , we now only have to prove that $U' \in \mathbb{N}$. We will separate dimensions as in definition 2.11.

We distinct two cases:

- $d < M$

In this case, $u' = u$, and we know that $u \in \mathbb{N}$, as $*S_d^n \in \mathbb{S}_d^n$.

- $d = M$

In this case, $u' = U_M * U_m$. As both $U_M, U_m \in \mathbb{N}$, we know $u' \in \mathbb{N}$. ■

Lemma 2.37 ($FoldLast2_{\leftarrow}^i$ is well defined). (*fold-wdi*)

Given any $I' \in *S_d^{n-1'}$. We now have to prove that $FoldLast2_{\leftarrow}^i(I') \in S_d^n$.

Proof. Using lemma 2.9 for codomain S_d^n , we can now prove for any given I' and $I = FoldLast2_{\leftarrow}^i(I')$: $I' \in S_d^{n-1'} \rightarrow I \in S_d^n$. We separate dimensions as in definition 2.11\$. For $d < M$, we have an identity mapping and we can use lemma 2.14\$. For the rest of the proof, we only have to prove the cases (A) $d = M$ and (B) $d = m$.

(A) $d = M$

$i' < u' \rightarrow i < u$

\iff

$i' < U_M * U_m \rightarrow \left\lfloor \frac{i'}{U_m} \right\rfloor < U_M$

Expand definitions of i and u'

\iff

$\frac{i'}{U_M} < U_m \rightarrow \left\lfloor \frac{i'}{U_m} \right\rfloor < U_M$

Divided left side of the implication by U_M

This holds, as $\lfloor X \rfloor \leq X$ for any X .

(B) $d = m$

$I'_M < U'_M \rightarrow i < u$

\iff

$I'_M < U'_M \rightarrow i'_M \% U_m < U_m$

Expand definition of i

This holds, as $X \% Y < Y$ holds for any $X, Y \in \mathbb{N}$. ■

Lemma 2.38 ($FoldLast2_{\leftarrow}^i$ is operational bijective). (*fold-opbi*)

Take the operational parts of $S_d^n : \mathcal{O}(S_d^n)$ and $S_d^{n-1'} : \mathcal{O}_{FoldLast2_{\leftarrow}^i}(S_d^n)$. We now have to prove that $FoldLast2_{\leftarrow}^i : \mathcal{O}_{FoldLast2_{\leftarrow}^i}(S_d^n) \rightarrow \mathcal{O}(S_d^n)$ is bijective.

Proof. We separate dimensions as in definition 2.11. If $d < M$, we have an identity mapping and we will use lemma 2.14. For the rest of this proof, we can assume that $d = M$ or $d = m$.

Using lemma ??, we now have to prove (A) surjectivity and (B) injectivity.

(A) Given $I \in \mathcal{O}(S_d^n)$, $I'_M = U_m * I_M + I_m$ and $I'_d = I_d$ for any $d < M$. We now have to prove that (A.1) $FoldLast2_{\leftarrow}^i(I') = I$ and (A.2) $I' \in S_d^{n-1'}$.

(A.1) We distinct three cases: $d < M$, $d = M$ and $d = m$.

- $d < M$

In this case we have an identity mapping, and we can use lemma 2.14.

- $d = M$

$FoldLast2_{\leftarrow M}^i(I')$

$= \left\lfloor \frac{I'_M}{U_m} \right\rfloor$

Expand definition of $FoldLast2_{\leftarrow}^i$

$= \left\lfloor \frac{U_m * I_M + I_m}{U_m} \right\rfloor$

Expand definition of I'_M

$= \left\lfloor \frac{U_m * I_M}{U_m} \right\rfloor$

As $U_m * I_M$ is a multiple of U_m and $I_m < U_m$, we can remove I_m from the equation

$= I_M$

- $d = m$

$FoldLast2_{\leftarrow m}^i(I')$

$= I'_M \% U_m$

Expand definition of $FoldLast2_{\leftarrow}^i$

$= (U_m * I_M + I_m) \% U_m$

Expand definition of I'_M

$= I_m \% U_m$

As $U_m * I_M \% U_m = 0$, we can remove it

$= I_m$

As $I_m < U_m$ by definition of $I \in S_d^n$

(A.2) For dimensions $d < M$ this is trivial, as this is the identity mapping. We now only have to discuss dimension M .

$$4 \quad I_M < U_M \wedge I_m < U_m \rightarrow I'_M < U'_M \\ \iff$$

$$I_M < U_M \wedge I_m < U_m \rightarrow U_m * I_M + I_m < U_M * U_m$$

Expand definitions of I'_M and U'_M

$$8 \quad \iff$$

$$I_M < U_M \wedge I_m < U_m \rightarrow \\ U_m * I_M \leq U_M * U_m - U_m \wedge I_m < U_m$$

Split the smaller then into two smaller smaller then's

$$12 \quad \iff$$

$$I_M < U_M \wedge I_m < U_m \rightarrow I_M \leq U_M - 1 \wedge I_m < U_m$$

Divide by U_m

This is trivially true.

16 (B) Given an $I \in S_d^n$ and $I', I'' \in S_d^{n-1}$. We now have to prove that if $FoldLast2_{\leftarrow M}^i(I) = I_M$, $FoldLast2_{\leftarrow M}^i(I'') = I_M$, $FoldLast2_{\leftarrow m}^i(I') = I_m$ and $FoldLast2_{\leftarrow m}^i(I'') = I_m$, then $I'_M = I''_M$. We prove this by contradiction. Assume there are I', I'' with $I'_M \neq I''_M$ for which this holds:

$$I'_M \neq I''_M \wedge \\ FoldLast2_{\leftarrow M}^i(I') = I_M \wedge FoldLast2_{\leftarrow M}^i(I'') = I_M \wedge \\ 24 \quad FoldLast2_{\leftarrow m}^i(I') = I_m \wedge FoldLast2_{\leftarrow m}^i(I'') = I_m$$

$$\iff$$

$$I'_M \neq I''_M \wedge \\ \left[\frac{I'_M}{U_m} \right] = I_M \wedge \left[\frac{I''_M}{U_m} \right] = I_M \wedge \\ 28 \quad I'_M \% U_m = I_m \wedge I''_M \% U_m = I_m$$

Expand definition of $FoldLast2_{\leftarrow}^i$

For $\left[\frac{I'_M}{U_m} \right] = I_M \wedge \left[\frac{I''_M}{U_m} \right] = I_M$ to hold, I'_M and I''_M can differ at most $U_m - 1$.

32 For $I'_M \% U_m = I_m \wedge I''_M \% U_m = I_m$ to hold, I'_M and I''_M must differ an exact multiple of U_m .

Combined, this ensures that $I'_M = I''_M$, which contradicts with $I'_M \neq I''_M$. \blacksquare

36

Theorem 2.39 (FoldLast2 is mapping-equivalent). (fold-me)

By definition 2.7, FoldLast2 is mapping equivalent iff:

- 40 • (A) $FoldLast2_{\rightarrow}^s$ is well defined
- (B) $FoldLast2_{\leftarrow}^i$ is well defined
- (C) $FoldLast2_{\leftarrow}^i$ is operationally bijective

Proof. (A) is proven in lemma 2.31

(B) is proven in lemma 2.32

(C) is proven in lemma 2.33 \blacksquare

2.9 Permute

Permute changes the order of dimensions, preparing it for a mapping operating on specific dimensions, such as *FoldLast2* or *PadLast*. The mapping does not change the dimensions themselves.

Permute takes an extra argument $P \in \mathbb{Z}^n$. We assume P to be a valid permutation of $0..[n-1]$.

Definition 2.40 (*Permute*). (def-permute)

$$\bullet \quad Permute_{\rightarrow}^s : \mathbb{Z}^n \rightarrow *S_g^n \rightarrow *S_g^n$$

$$Permute_{\rightarrow}^s(P)(L, U, T, W) = (L', U', T', W')$$

For definitions of L', U', T', W' , we separate dimensions as in definition 2.11.

$$\begin{aligned} - l' &= L_{p'} \\ - u' &= U_{p'} \\ - t' &= T_{p'} \\ - w' &= W_{p'} \end{aligned}$$

$$\bullet \quad Permute_{\leftarrow}^i : S_d^{n'} \rightarrow S_n^n$$

$$Permute_{\leftarrow}^i(P)(I') = I$$

For definition of I , we separate dimensions as in definition 2.11.

$$\begin{aligned} - \text{let } p' &\in \mathbb{Z} \text{ with } P_{p'} = d \\ i &= I'_{p'} \end{aligned}$$

Lemma 2.41 ($Permute_{\rightarrow}^s$ is well defined). (permute-wds)

Given any $*S_g^n \in *S_g^n$. We now have to prove that $Permute_{\leftarrow}^i(*S_g^n) \in *S_g^n$.

Proof. Using lemma 2.8 for S_g^n , we now have to prove that (A) $L', U', T', W' \in \mathbb{Z}^n$ and (B) $L \leq U$ and $W \leq T$. We separate dimensions as in definition 2.11.

(A) We know that $l = L_p$. By definition, $L_p \in \mathbb{Z}$. So, we can conclude that $l \in \mathbb{Z}$ also holds. Similarly, $u, t, w \in \mathbb{Z}$ also hold.

(B) We know that $l = L_p$ and $u = U_p$. By definition, $L_p \leq U_p$. So, we can conclude that $l \leq u$ also holds. Similarly, $w \leq t$ also holds. \blacksquare

Lemma 2.42 ($Permute_{\leftarrow}^i$ is well defined). (permute-wdi)

Given any $I' \in S_g^{n'}$. We now have to prove that $Permute_{\leftarrow}^i(I') \in S_g^n$.

Proof. Using lemma 2.9, we can now prove for any given I' and $I = \text{PruneGrid}_{\leftarrow}^i(I')$. We split separate dimensions for I .

$$\begin{aligned} 4 \quad L' &\leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ i &= \perp \vee (l \leq i < \downarrow u \wedge (i - l) \xrightarrow{\%} t < \downarrow w) \end{aligned}$$

\iff

$$\begin{aligned} 8 \quad L' &\leq I' < \downarrow U' \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ I'_{p'} &= \perp \vee (L'_{p'} \leq I'_{p'} < \downarrow U'_{p'} \wedge (I'_{p'} - L'_{p'}) \xrightarrow{\%} T'_{p'} < \downarrow W'_{p'}) \\ &\text{By definition of } \text{Permute}_{\leftarrow}^i, \text{ with } P_{p'} = d. \text{ Because } P \\ &\text{is a permutation, we know such a } p' \text{ exists.} \end{aligned}$$

This holds, as the right hand side of the implication is equal to the left hand side in a single dimension. \blacksquare

Lemma 2.43 ($\text{Permute}_{\leftarrow}^i$ is operational bijective). (*permute-opbi*)

Take the operational parts of $S_g^n : \mathcal{O}(S_g^n)$ and $S_g^n : \mathcal{O}_{\text{Permute}_{\leftarrow}^i}(S_g^n)$. We now have to prove that $\text{Permute}_{\leftarrow}^i : \mathcal{O}_{\text{Permute}_{\leftarrow}^i}(S_g^n) \rightarrow \mathcal{O}(S_g^n)$ is bijective.

Proof. Using lemma 2.10, we now have to prove \textcircled{A} surjectivity and \textcircled{B} injectivity.

\textcircled{A} Given an $I \in \mathcal{O}(S_g^n)$. For each $d' \in [0..n-1]$, we take I' with $\textcircled{1}$ (tfi-opbi-IP) $I'_{d'} = I_{P_{d'}}$. We now have to prove that $\textcircled{A.1}$ $\text{Permute}_{\leftarrow}^i(I') = I$ and $\textcircled{A.2}$ $I' \in S_g^{n'}$.

$\textcircled{A.1}$ By definition, for each $d \in [0..n-1]$ we have a $p' \in \mathbb{Z}$ with $\textcircled{2}$ (tfi-opbi-Pp) $P_{p'} = d$. We have to prove that $I_d = I'_{p'}$.

$$\begin{aligned} 28 \quad I_d &= I'_{p'} \\ &= I_{P_{p'}} \textcircled{1} \\ &= I_d \textcircled{2} \end{aligned}$$

$\textcircled{A.2}$ We separate dimensions using definition 2.11 for I' . By definition, we now need to prove that:

$$i' = \perp \vee (l' \leq i' < u' \wedge (i' - l') \% t' < w')$$

\iff

$$\begin{aligned} 36 \quad L_{p'} &\leq I_{P_d} < U_{p'} \wedge (I_{p'} - L_{p'}) \% T_{p'} < W_{p'} \\ &\text{Truncated left hand side of the 'or,' expanded definition} \\ &\text{of } \text{Permute}_{\rightarrow}^s \text{ for } l', u', t', w', \text{ expanded definition of } i' \\ &\textcircled{1} \end{aligned}$$

\iff

$$40 \quad L_{p'} \leq I_{p'} < U_{p'} \wedge (I_{p'} - L_{p'}) \% T_{p'} < W_{p'}$$

This holds, as it holds for any dimension $p' \in [0..n-1]$, and $p' \in [0..1]$ as P is a permutation of $[0..n-1]$.

\textcircled{B} Given $I \in \mathcal{O}(S_g^n)$ and $I', I'' \in S_g^{n'}$. We now have to prove that if $\text{Permute}_{\leftarrow}^i(I') = I$ and $\text{Permute}_{\leftarrow}^i(I'') = I$, then $I' = I''$. We prove this by contradiction. Assume there are two $I' \neq I''$ with $\text{Permute}_{\leftarrow}^i(I') = I$ and $\text{Permute}_{\leftarrow}^i(I'') = I$. We separate dimensions for I' and I'' using definition 2.11

$$i' \neq i'' \wedge \text{Permute}_{\leftarrow}^i(i') = i \wedge \text{Permute}_{\leftarrow}^i(i'') = i$$

\iff

$$i' \neq i'' \wedge P_{p'} = d \wedge I'_{p'} = i \wedge I''_{p'} = i$$

\iff

$$I' \neq I'' \wedge P_{p'} = d \wedge I' = I \wedge I'' = I$$

As P is a permutation, the above holds for all $p' \in [0..n-1]$.

This is a contradiction by transitivity of $=$ \blacksquare

Theorem 2.44 (Permute is mapping-equivalent). (*permute-me*)

By definition 2.7, Permute is mapping equivalent iff:

- \textcircled{A} $\text{Permute}_{\rightarrow}^s$ is well defined
- \textcircled{B} $\text{Permute}_{\leftarrow}^i$ is well defined
- \textcircled{C} $\text{Permute}_{\leftarrow}^i$ is operationally bijective

Proof. \textcircled{A} is proven in lemma 2.41

\textcircled{B} is proven in lemma 2.42

\textcircled{C} is proven in lemma 2.43 \blacksquare

2.10 Padlast

PadLast extends the length of the last dimension to a multiple of a given integer. This is to prepare it for the GPU warp size, or to be split using *SplitLast*.

PadLast takes an extra argument $p \in \mathbb{Z}$ to indicate the size of the padding.

Definition 2.45 (*PadLast*). (def-pad)

- $\text{PadLast}_{\rightarrow}^s : \mathbb{Z} \rightarrow *S_g^n \rightarrow *S_g^n$

$$\text{PadLast}_{\rightarrow}^s(p)(L, U, T, W) = (L', U', T', W'):$$

$$- L' = L$$

$$- T' = T$$

$$- W' = W$$

- For the definition of U' , we separate dimensions as in definition 2.11.

$$u' = \begin{cases} u & | \quad d \neq n-1 \\ \left[\frac{u}{p} \right] * p & | \quad d = n-1 \end{cases}$$

- $PadLast_{\leftarrow}^i : S_g^{n'} \rightarrow S_g^n$
 $PadLast_{\leftarrow}^i(p)(I') = I$

–

$$I = \left\{ \begin{array}{l|l} I' & I' < \downarrow U \\ \perp & I' \not< \downarrow U \end{array} \right.$$

Lemma 2.46 ($PadLast_{\rightarrow}^s$ is well defined). (*pad-wds*)

- Given any $*S_g^n \in *S_g^n$. We now have to prove that $PadLast_{\rightarrow}^s(*S_g^n) \in *S_g^n$.

Proof. Using lemma 2.8 for S_g^n , we now have to prove that $\textcircled{A} L', U', T', W' \in \mathbb{Z}^n$ and $\textcircled{B} L \leq U$ and $W \leq T$. We separate dimensions as in definition 2.11.

\textcircled{A} By definition of $PadLast_{\rightarrow}^s$: $L' = L, T' = T, W' = W$. Because $L, T, W \in \mathbb{Z}$, $L', T', W' \in \mathbb{Z}$ as well. Similarly, $U' = \left\lfloor \frac{U}{p} \right\rfloor * 32$, so $U' \in \mathbb{Z}$ as well.

- \textcircled{B} We know that $L' = L, U' \geq U, T' = T, W' = W$. We also know that $L \leq U$ and $W \leq T$. Together, this proves $L' \leq U'$ and $W' \leq T'$. ■

Lemma 2.47 ($PadLast_{\leftarrow}^i$ is well defined). (*pad-wdi*)

- Given any $I' \in S_g^{n'}$. We now have to prove that $PadLast_{\leftarrow}^i(I') \in S_g^n$.

Proof. Using lemma 2.9, we can now prove for any given I' and $I = PadLast_{\leftarrow}^i(I')$. We distinct two cases:

- $I' < \downarrow U$

$$L' \leq I' < \downarrow U \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ I = \perp \vee (L \leq I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W)$$

\iff

$$L \leq I' < \downarrow U \wedge (I' - L) \xrightarrow{\%} T < \downarrow W \rightarrow \\ L \leq I' < \downarrow U \wedge (I' - L) \xrightarrow{\%} T < \downarrow W$$

- Expand definition of $PadLast_{\rightarrow}^s$ for L', T', W' .
- Replace $I' < \downarrow U$ with distinction assumption $I' < \downarrow U$.
- Prune left hand side of the ‘or’
- Expand definition of $PadLast_{\leftarrow}^i$ for I using $I' < \downarrow U$

- $I' \not< \downarrow U$

$$L' \leq I' < \downarrow U \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ I = \perp \vee (L \leq I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W)$$

\iff

$$L' \leq I' < \downarrow U \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ \perp = \perp$$

Expand definition of $PadLast_{\leftarrow}^i$ for I using $I' \not< \downarrow U$, prune right hand side of the ‘or’

4

Lemma 2.48 ($PadLast_{\leftarrow}^i$ is operational bijective). (*pad-opbi*)

Take the operational parts of $S_g^n : \mathcal{O}(S_g^n)$ and $S_g^{n'} : \mathcal{O}_{PadLast_{\leftarrow}^i}(S_g^n)$. We now have to prove that $PadLast_{\leftarrow}^i : \mathcal{O}_{PadLast_{\leftarrow}^i}(S_g^n) \rightarrow \mathcal{O}(S_g^n)$ is bijective.

Proof. Using lemma 2.10, we now have to prove \textcircled{A} surjectivity and \textcircled{B} injectivity.

\textcircled{A} Given an $I \in \mathcal{O}(S_g^n)$ and $I' \in S_g^n$ with $I' = I$. We now have to prove that $\textcircled{A.1} PadLast_{\leftarrow}^i(I') = I$ and $\textcircled{A.2} I' \in S_g^{n'}$.

$\textcircled{A.1}$ By definition 2.45 of $PadLast_{\leftarrow}^i$, $PadLast_{\leftarrow}^i(I') = I \iff I' < \downarrow U$. We know that $I' = I$, and by definition of $I \in \mathcal{O}(S_g^n)$, $I < \downarrow U$. So, we can conclude that $PadLast_{\leftarrow}^i(I') = I$ holds.

$\textcircled{A.2}$

$$L \leq I < \downarrow U \wedge (I - L) \xrightarrow{\%} T < \downarrow W \rightarrow \\ L' \leq I' < \downarrow U \wedge (I' - L') \xrightarrow{\%} T' < \downarrow W'$$

\iff

$$L' \leq I < \downarrow U \wedge (I - L') \xrightarrow{\%} T' < \downarrow W' \rightarrow \\ L' \leq I < \downarrow U \wedge (I - L') \xrightarrow{\%} T' < \downarrow W'$$

By expanding the definition of $PadLast_{\rightarrow}^s$ for L, T, W and the definition of I'

This holds, as U can only be bigger than U' .

\textcircled{B} Given $I \in \mathcal{O}(S_g^n)$ and $I', I'' \in S_g^{n'}$. We now have to prove that if $PadLast_{\leftarrow}^i(I') = I$ and $PadLast_{\leftarrow}^i(I'') = I$, then $I' = I''$. We prove this by contradiction. Assume there are two $I' \neq I''$ with $PadLast_{\leftarrow}^i(I') = I$ and $PadLast_{\leftarrow}^i(I'') = I$.

$$I' \neq I'' \wedge PadLast_{\leftarrow}^i(I') = I \wedge PadLast_{\leftarrow}^i(I'') = I$$

\iff

$I' \neq I'' \wedge I' = I \wedge I'' = I$
 $PadLast_{\leftarrow}^i(I')$ cannot be \perp , as $\perp \notin \mathcal{O}(S_g^n)$. Similarly for I'' .

This is a contradiction by transitivity of $= \neq$ ■

Theorem 2.49 ($PadLast$ is mapping-equivalent). (*pad-me*)

By definition 2.7, $PadLast$ is mapping equivalent iff:

- $\textcircled{A} PadLast_{\rightarrow}^s$ is well defined

- \textcircled{B} $\text{PadLast}_{\leftarrow}^i$ is well defined
- \textcircled{C} $\text{PadLast}_{\leftarrow}^i$ is operationally bijective

Proof. \textcircled{A} is proven in lemma 2.46

- 4 \textcircled{B} is proven in lemma 2.47
 \textcircled{C} is proven in lemma 2.48

■

Chapter 3

Implementation

The implementation of these mappings was done in the SAC compiler [1], which is implemented in C. The compiler takes a SAC program as input, generates C code, and uses a C or cuda C compiler to create an executable. We will distinct the implementation of the SAC compiler in three different parts: The existing compiler, the implementation of the mappings as defined in chapter ??, and the mechanisms that determine what mappings will be executed, in what order, and with what parameters. We will discuss each part in it's own section below.

3.1 The SAC compiler

First of all, there is the existing compiler. The part we're mainly interested in is the part with Jing's implementation for executing SAC with-loops on the GPU. This implementation is fairly limited and has many cases where a with-loop cannot be executed. If such a case occurs, it either results in a compiler or runtime error, while the program would have been valid in the normal sequential compiler backend.

AST Traversals

When the compiler compiles a SAC program for the CUDA backend, it will do so using many different AST traversals. We will discuss the traversals relevant to the execution of with-loops on the GPU using a small code example in figure ?? 3.1. Note that some of the code examples use a bit of pseudo code to illustrate internal AST representations. Also note that this example will probably not work like this when compiled, as some optimizations may take out the loop altogether.

- Annotate cudarizable loops: This compiler phase determines whether a with-loop is eligible for execution on a GPU. Examples of with-loops not eligible are very small loops, nested with-loops, or with loops with non-cudarizable function calls in their bodies. This compiler phase sets a boolean

```
1 int[.,.] plusone(int[.,.] a) {
2   b = with {
3     (0 <= iv < shape(a)) : a[iv] + 1;
4   } : genarray (shape(a), 0);
5   return b;
6 }
```

(trav-ex)

code example 3.1: Example SAC program

flag in with-loop AST nodes. Note that the dimensionality is always known. When a cudarizable function has an unknown dimensionality, e.g. it can be used with any dimensionality, a separate instance of this function is created for each used dimensionality.

- Insert primitives to prepare for GPU execution: These are multiple phases responsible for transferring variables, constants and arrays to the GPU before execution, and the generated array back to the CPU after execution. Figure 3.2 displays the state of the example program after this phase.
- Create CUDA kernels: Create a CUDA kernel function for each partition in the with loop. Note that multiple partitions can exist within a with-loop. Each kernel function will contain the body of the partition it was created for. Information as the lowerbound, upperbound, step and width will be passed in as an argument, as well as GPU pointers to the arrays. The current iv can be derived from the variables threadIdx and blockIdx, which both contain the properties x, y and z. Figure 3.3 displays the state of the example program after this phase.
- Generate C code: The actual generated C code does not contain any pseudocode anymore, but should be compilable by the C-compiler.

```

1  int[.,.] plusone(int[.,.] a) {
2      a_d = cudaMemCpyHost2Dev(a);
3      b_d = cudaMemCreate(...);
4
5      b_d = with {
6          (0 <= iv < shape(a_d)) : a_d[iv]+1;
7      } : genarray (shape(a_d), 0);
8
9      b = cudaMemCpyDev2Host(a_d);
10     cudaMemFree(a_d);
11     cudaMemFree(b_d);
12
13     return b;
14 }
(trav-mem)

```

code example 3.2: Example SAC program after memory transfers have been inserted

Grid and block

In figure 3.3, there are a few lines of code left undefined. The most important of these are the lines annotated with the comments *Compute CUDA grid and block sizes* and *Recompute iv from threadIdx and blockIdx*. The original implementation is pretty straightforward. Let us define the lengths of the dimensions as $d_0\dots d_n$. Now we make a case distinction for the first five dimensionalities, and map them as described in table 3.4. For 6 or more dimensions, the current implementation gives a compiler error. Note that the block size has a maximum 1024 threads (may vary per GPU), so d_0*d_1 can be at most 1024. If it is not, a runtime error is thrown.

Inside the kernel, the original iv is reconstructed from the threadIdx and blockIdx variables. For dimensionalities 3 to 5, we can directly take the values of blockIdx.z, blockIdx.y, blockIdx.x, threadIdx.y and threadIdx.x. However, for dimensionalities 1 and 2, iv_0 and iv_1 have to be recomputed, and because the dimension length has been padded to be divisible by 32, we also have to check it against the original dimension length. For iv_0 , we do this as shown in code example 3.5. iv_1 is computed similarly. Note that in the generated C code, these values do not exist as an array. Instead, all values for each dimensions have their own variables.

Lowerbound, upperbound, step and width

In table 3.4, we used dimension lengths $d_0\dots d_n$. However, in our C program we have $lb_0\dots lb_n$, $ub_0\dots ub_n$, $step_0\dots step_n$ and $width_0\dots width_n$. This means we

```

1  int[.,.] plusone(int[.,.] a) {
2      a_d = cudaMemCpyHost2Dev(a);
3      b_d = cudaMemCreate(...);
4
5      // Compute CUDA grid and block sizes
6      grid = ...
7      block = ...
8
9      plusone_kernel<grid, block>(
10         lb_0, lb_1, ub_0, ub_1, a_d, b_d);
11
12     b = cudaMemCpyDev2Host(a_d);
13     cudaMemFree(a_d);
14     cudaMemFree(b_d);
15
16     return b;
17 }
18
19 void plusone_kernel_0 (
20     lb_0, lb_1, ub_0, ub_1, a_d, b_d) {
21     // Recompute iv from
22     // threadIdx and blockIdx
23     iv = ...
24
25     b_d[iv] = a_d[iv] + 1;
26 }
(trav-kernel)

```

code example 3.3: Example SAC program after kernels have been generated

have to make a few slight changes still:

$\forall i < dims :$

- We take the upperbound as the dimension lengths

```

1  // Computing grid and block sizes
2  d_i = ub_i

```

- We subtract the lowerbound from the dimension length. Note that we have to add it again inside the kernels.

```

1  // Computing grid and block sizes
2  d_i = d_i - lb_i
3
4  // Recomputing the index vector
5  iv_i = iv_i + lb_i

```

- The step and width do not influence the lengths of the dimensions. Instead, we will check inside of the kernel whether iv is inside the grid:

```

1  // Recomputing the index vector
2  if (iv_i % step_i > width_i)
3      return;

```

d	grid.z	grid.y	grid.x	blk.y	blk.x
1	1	1	$d_0/32 + 1$	1	32
2	1	$d_1/32 + 1$	$d_0/32 + 1$	32	32
3	1	1	d_2	d_1	d_0
4	1	d_3	d_2	d_1	d_0
5	d_4	d_3	d_2	d_1	d_0

(jing-mapping)

table 3.4: Case distinction for mapping n -dimensional index spaces onto the GPU, using Jing’s heuristics

```

1  iv_0 = blockIdx.x * 32 + threadIdx.x;
2  if (iv_0 >= d_0) return;
(trav-kernel-iv)

```

code example 3.5: Recomputation of iv_0 in the kernel, for dimensionalities 1 and 2. iv_1 is computed similarly.

If we combine all this together, we get code example 3.6 as output of the compiler:

Optimizations

As SAC is a functional language, the SAC compiler can do many different optimizations that are unique to functional languages. For example, the SAC compiler will, in this case, probably decide that the memory of variable a can be reused. This means that no b or b_dev have to be allocated. Furthermore, the SAC compiler may just omit the with-loop altogether, and replace it with an accessor that adds one to the original array value. The code fragments we have been looking at merely serve as an example, so there are no guarantees that the SAC compiler will, in this case, behave exactly as illustrated here.

3.2 Mapping execution

The mappings introduced in ?? change how the grid and block spaces are generated, and how the iv is computed from the grid and block coordinates. This also means that the mapping execution implementation only changes the generated code responsible for those two tasks. In code example 3.3, this would be lines 5-7 and 21-23.

Differences from theoretical model

Because C works a bit different then theoretical maths, the C implementation differs a bit from the theoretical model we implemented in ?. Besides these implementation details, there are also a few conceptual differences in the implementation.

Grid index spaces

The representation of the grid index spaces in the compiler are very similar to the definition of a grid index space as defined in section 3.2. There are, however, a few small differences.

- The arrays of I, L, U, T and W variables only exist in the compiler. In the generated code, the values are represented either in individual variables or, if possible, as constants.
- The filter function f does not exist in the runtime. The check is immediately executed alongside of the mappings, and if it fails we either return immediately in branching mode, or we set an “outside of grid” variable in branchless mode. More on branching/branchless modes in section TODO.

3.3 Mapping generation using strategies

```

1  int[.,.] plusone(int[.,.] a) {
2      // Assume all lb_i, ub_i, step_i
3      // and width_i variables exist
4
5      a_d = cudaMemCpyHost2Dev(a);
6      b_d = cudaMemCreate(...);
7
8      // Compute CUDA grid and block sizes
9      grid.x = (ub_0 - lb_0) / 32;
10     grid.y = (ub_1 - lb_1) / 32;
11     block.x = 32;
12     block.y = 32;
13
14     plusone_kernel<grid, block>(
15         lb_0, lb_1, ub_0, ub_1, a_d, b_d);
16
17     b = cudaMemCpyDev2Host(a_d);
18     cudaMemFree(a_d);
19     cudaMemFree(b_d);
20
21     return b;
22 }
23
24 void plusone_kernel_0 (
25     lb_0, lb_1, ub_0, ub_1, a_d, b_d) {
26     // Recompute iv from
27     // threadIdx and blockIdx
28     iv_0 = blockIdx.x * 32 + threadIdx.x;
29     iv_0 = iv_0 + lb_0
30     if (iv_0 >= ub_0) return;
31     // No step or width, so check omitted
32
33     iv_1 = blockIdx.y * 32 + threadIdx.y;
34     iv_1 = iv_1 + lb_1
35     if (iv_1 >= ub_1) return;
36     // No step or width, so check omitted
37
38     iv = {iv_1, iv_0};
39
40     b_d[iv] = a_d[iv] + 1;
41 }
(trav-complete)

```

code example 3.6: Example SAC program after it has been compiled

Chapter 4

Listings

List of figures

List of tables

4	3.4 Case distinction for mapping n-dimensional index spaces onto the GPU, using Jing's heuristics	26
---	---	----

List of code examples

	3.1 Example SAC program	24
8	3.2 Example SAC program after memory transfers have been inserted	25
	3.3 Example SAC program after kernels have been generated	25
12	3.5 Recomputation of <code>iv_0</code> in the kernel, for dimensionalities 1 and 2. <code>iv_1</code> is computed similarly.	26
16	3.6 Example SAC program after it has been compiled	27

List of definitions

	2.1	Definition (Preliminary assumptions)	7
	2.2	Definition (Posets using product order)	7
4	2.3	Definition (Index space transformations)	8
	2.4	Definition (Index spaces)	8
	2.6	Definition (Operative and excess subsets)	9
	2.7	Definition (Mapping equivalence)	9
8	2.11	Definition (Separating dimensions)	10
	2.12	Definition (Dimension dependency)	11
	2.15	Definition (<i>ShiftLB</i>)	11
	2.20	Definition (<i>CompressGrid</i>)	12
12	2.25	Definition (<i>PruneGrid</i>)	16
	2.30	Definition (<i>SplitLast</i>)	17
	2.35	Definition (<i>FoldLast2</i>)	18
	2.40	Definition (<i>Permute</i>)	20
16	2.45	Definition (<i>PadLast</i>)	21

List of lemmas

2.5	Lemma (<i>gen</i> of normalized and dense index spaces)	8	
2.8	Lemma (Well-definedness of tf_{\rightarrow}^s)	9	
2.9	Lemma (Well-definedness of tf_{\leftarrow}^i)	10	4
2.10	Lemma (Operational bijectivity of tf_{\leftarrow}^i)	10	
2.13	Lemma (Proofs with separated dimensions)	11	
2.14	Lemma (Mapping equivalence for identity dimensions)	11	
2.16	Lemma ($ShiftLB_{\rightarrow}^s$ is well-defined)	11	8
2.17	Lemma ($ShiftLB_{\leftarrow}^i$ is well-defined)	12	
2.18	Lemma ($ShiftLB_{\leftarrow}^i$ is operational bijective)	12	
2.21	Lemma ($CompressGrid_{\rightarrow}^s$ is well defined)	13	
2.22	Lemma ($CompressGrid_{\leftarrow}^i$ is well defined)	13	12
2.23	Lemma ($CompressGrid_{\leftarrow}^i$ is operational bijective)	14	
2.26	Lemma ($PruneGrid_{\rightarrow}^s$ is well defined)	16	
2.27	Lemma ($PruneGrid_{\leftarrow}^i$ is well defined)	16	
2.28	Lemma ($PruneGrid_{\leftarrow}^i$ is operational bijective)	16	16
2.31	Lemma ($SplitLast_{\rightarrow}^s$ is well defined)	17	
2.32	Lemma ($SplitLast_{\leftarrow}^i$ is well defined)	17	
2.33	Lemma ($SplitLast_{\leftarrow}^i$ is operational bijective)	17	
2.36	Lemma ($FoldLast2_{\rightarrow}^s$ is well defined)	19	20
2.37	Lemma ($FoldLast2_{\leftarrow}^i$ is well defined)	19	
2.38	Lemma ($FoldLast2_{\leftarrow}^i$ is operational bijective)	19	
2.41	Lemma ($Permute_{\rightarrow}^s$ is well defined)	20	
2.42	Lemma ($Permute_{\leftarrow}^i$ is well defined)	20	24
2.43	Lemma ($Permute_{\leftarrow}^i$ is operational bijective)	21	
2.46	Lemma ($PadLast_{\rightarrow}^s$ is well defined)	22	
2.47	Lemma ($PadLast_{\leftarrow}^i$ is well defined)	22	
2.48	Lemma ($PadLast_{\leftarrow}^i$ is operational bijective)	22	28

List of theorems

	2.19 Theorem (ShiftLB is mapping-equivalent)	12
	2.24 Theorem (CompressGrid is mapping-equivalent)	15
4	2.29 Theorem (PruneGrid is mapping-equivalent)	16
	2.34 Theorem (SplitLast is mapping-equivalent)	18
	2.39 Theorem (FoldLast2 is mapping-equivalent)	20
	2.44 Theorem (Permute is mapping-equivalent)	21
8	2.49 Theorem (PadLast is mapping-equivalent)	22

References

- [1] URL: <https://www.sac-home.org/>.
- 4 [2] Jing Guo, Jeyarajan Thiyagalingam, and Sven-Bodo Scholz. “Breaking the GPU programming barrier with the auto-parallelising SAC compiler”. In: *Proceedings of the sixth workshop on Declarative aspects of multicore programming*. 2011, pp. 15–24.
- 8 [3] Hee Sik Kim J. Neggers. *Basic posets*, pp. 64–78. ISBN: 9789810235895. URL: https://books.google.nl/books?id=-ip3-wejeR8C&pg=PA64&redir_esc=y#v=onepage&q&f=false.
- 12