

Approaching Symbolic Parallelization by Synthesis of Recurrence Decompositions

Grigory Fedyukovich and Rastislav Bodík

Computer Science and Engineering, University of Washington, Seattle, Washington, USA

We present GRASSP, a novel approach to perform automated parallelization relying on recent advances in formal verification and synthesis. GRASSP augments an existing sequential program with an additional functionality to decompose data dependencies in loop iterations, to compute partial results, and to compose them together. We show that for some classes of the sequential prefix sum problems, such parallelization can be performed efficiently.

1 Introduction

Parallelization of software is important for improving its effectiveness and productivity. Industrial applications (studied, e.g., in [17]) operate with large inputs and therefore could run days. But even a small program that iterates over a single array and incrementally computes a single numerical output could be challenging for parallelization. It would require partitioning of the input data into a sequence of segments, processing each segment separately, and merging the partial outputs for the segments.

Data dependencies prevent parallel processing of the segments. Therefore, the parallel loop should break the dependencies by devising an alternative function. This additional processing could invoke some analysis either *before* the parallel execution: in order to move segment boundaries, or *after* the parallel execution: in order to repair the outputs broken by violated dependencies. In this paper, we address the challenge of synthesizing such additional code automatically.

Preserving equivalence is a crucial requirement to automatic parallelization, and it aims at confirming that pairwise equivalence of inputs implies pairwise equivalence of outputs [14, 18, 5, 10, 15, 8]. Rather than merely relying on the existing solutions to *verify* equivalence between programs, we aim to *synthesize* a parallel program that is equivalent to the given sequential program. Given arbitrary segments of the input data, we build on an assumption that segment boundaries can always be adjusted, and the actual loop computation can be performed for the updated segments.

Program synthesis is an approach to generate a program implementation from the given specification. State-of-the-art synthesis tools employ the Counter-Example-Guided Inductive Synthesis (CEGIS) [20] paradigm, i.e., assume a space of candidate implementations and check whether there exists a candidate among them that matches the given specification. In this work, we exploit the capabilities of SMT-based bounded model checking and template-based synthesis in ROSETTE [21, 22].

Our novel approach GRASSP automatically generates a search space of candidate decompositions for the given program, each of which consists of function *prefix* (that would identify during runtime how segment boundaries should be moved), and function *merge* (that would identify how the partial outputs should be combined to get the final output). Then, for a given number of segments and for a given maximal length of each segment, GRASSP chooses a candidate decomposition and formally verifies if it produces results equivalent to the given sequential program. Our preliminary experiments with GRASSP confirm that for reasonably small programs the decomposition can be found in seconds.

2 Sequential Recurrence Decomposition

We start by introducing a functional notation for array-handling programs, and proceed by formulating the parallelization criteria for them. We focus on functions taking a single finite-sized array as input, and recurrently computing a single output.

An *input* and an *output* are the designated variables respectively of type In and Out . An n -sized array is a finite sequence of inputs:

$$A : In^n$$

In this paper, we consider functions of the following type:

$$f : D \times In \rightarrow D$$

where D is a domain of any type. An element d of type D is called *state*. Intuitively, f takes a state and updates it with respect to a given input. The state calculated by function f can further be taken by function h to compute an *output*:

$$h : D \rightarrow Out \qquad h(f(d, input)) = output$$

In the scope of the paper, we are interested in iterative application of function f to the elements of an n -sized array by means of the higher-order function *fold*:

$$fold : (D \times In \rightarrow D) \times D \times In^n \rightarrow D$$

For example, consider array $A = \langle input_1, \dots, input_n \rangle$. In the first iteration, the initial state d_0 is taken by f and updated with respect to the first element $input_1$. Then, the updated state is taken by f again and updated with respect to the second element $input_2$. A state obtained after n iterative applications of f to d_0 is called *final*. It is represented as the following first-order *recurrent relation*:

$$fold(f, d_0, A) = f(input_n, f(input_{n-1}, \dots, f(input_2, f(input_1, d_0))))$$

Notably, *output* has to be computed only for the final state:

$$h(fold(f, d_0, A)) = output$$

Throughout the paper, we rely on an operator that concatenates m arrays:

$$append : In^{n_1} \times \dots \times In^{n_m} \rightarrow In^{n_1 + \dots + n_m}$$

It is important to ensure the following functional property of *append*:

$$fold(f, d_0, append(A_1, \dots, A_m)) = fold(f, fold(f, \dots, fold(f, d_0, A_1), \dots, A_{m-1}), A_m) \quad (1)$$

The left-hand-side of equation (1) denotes the initial state d_0 iteratively updated by f with respect to all elements of $append(A_1, \dots, A_m)$. The right-hand-side of the equation consists of m groups of consequent applications of *fold* to each of the arrays $\{A_i\}$. That is, the final state obtained for an i -th application of *fold* is further used for the $(i+1)$ -th application of *fold*. We refer to this property to as *sequential recurrence decomposition*, since it guarantees equivalence between a single application of *fold* to $append(A_1, \dots, A_m)$ and m recurrent applications of *fold*.

Trivially, the equivalence of final states entails equivalence of the outputs computed out of these states.

array-max

```

1 ; given functions:
2   (define (f element current-max) (max element current-max))
3   (define (array-max A) (foldl f -inf.0 A))
4
5 ; function needed for parallel decomposition:
6   (define (merge-max out) (array-max out))

```

Figure 1: Calculating a maximal element of the array and a suitable function *merge* in ROSETTE.

3 Parallel Recurrence Decomposition

Application of *fold* to each of the arrays $\{A_i\}$ in parallel requires the recurrent relation to be decomposed. We assume that each application of *fold* to A_i takes the same initial state d_0 , and refer to the corresponding final states $\{d_i\}$ as to *partial* states (and to the corresponding outputs as to *partial* outputs).

$$\forall i \cdot d_i \triangleq \text{fold}(f, d_0, A_i) \qquad \forall i \cdot \text{output}_i \triangleq h(d_i)$$

The question is how to find a suitable function *merge* that takes all partial results as input and produces an output that is equivalent to the output of the corresponding sequential computation:

$$\text{merge} : \text{Out}^m \rightarrow \text{Out}$$

In the rest of the section, we aim at establishing a property of *parallel* (as opposed to sequential) *recurrence decomposition* that binds together all ingredients of the parallel processing of m arrays. Interestingly, there are several possible ways of defining this property, depending on existence of function *merge* for each particular f . We consider three such cases.

3.1 Direct Decomposition

The first case assumes existence of function *merge* that is *directly* applicable to all partial outputs obtained after applications of *fold* to each A_i :

$$h(\text{fold}(f, d_0, \text{append}(A_1, \dots, A_m))) = \text{merge}(\text{output}_1, \dots, \text{output}_m) \quad (2)$$

Example. Consider function `array-max` that calculates a maximal element of a given array in ROSETTE¹ (shown in Fig. 1). Function `f` is essentially a wrapper for the built-in function `max`, i.e., it outputs a maximal element of pair $\langle \text{element}, \text{current-max} \rangle$. Function `array-max` delegates its functionality to `foldl` (which stands for left-fold) that iteratively applies `f` to the initial state $-\infty$ and to the elements of array A . For some partitioning $A = (\text{append } A_1 \dots A_m)$, individual applications of `array-max` to each segment produce array `out` consisting of m elements, where `outi` is maximal element for A_i . It is easy to see that the maximal element of A is the maximal element of `out`. So for `array-max` there exists function *merge*, and it is equal to `array-max`. \square

We return to this scheme in Sect. 5 and show how function *merge* can be synthesized.

¹We provide *fold*-like implementation to comply with the chosen formalization. An alternative implementation using `apply` is also possible (similar to `max` in Fig. 3).

3.2 Decomposition with Constant Prefixes

When no function *merge* meeting the condition (2) exists, then the computations $fold(f, d, A_i)$ and $fold(f, d, A_{i+1})$ depend on each other and cannot be *correctly* performed from the speculative initial state $d = d_0$. If we are in the lucky situation that the impact of the speculative initial state is localized to a prefix of A_i then we can repair the incorrect execution by recomputing the affected prefix. In this subsection, we show how such a repair can be performed on a constant-size prefix, if one exists. First, computations $fold(f, d, A_i)$ are performed in parallel from the initial state $d = d_0$, computing the partial result d_i . Subsequently, the computations are repeated for a constant-size prefix of A_{i+1} , starting from the state d_i . The result of processing the prefix is then supplied instead of d_i to a suitable function *merge*.

We say that A' is a *prefix* of A if $A = append(A', A'')$ for some A'' . We allow the prefix to be an empty array. We denote by $prefix(A)$ the prefix of A , and by $prefix_{length}$ its predetermined length. We call such prefix the *constant prefix*.

Assuming the existence of constant prefixes for each A_{i+1} , the corresponding partial outputs are computed and repaired as follows:

$$\forall i \cdot d_i^{completed} \triangleq fold(f, d_i, prefix(A_{i+1})) \quad \forall i \cdot output_i^{completed} \triangleq h(d_i^{completed})$$

Assuming the existence of a suitable function *merge*, the repaired partial outputs are combined:

$$h(fold(f, d_0, append(A_1, \dots, A_m))) = merge(output_1^{completed}, \dots, output_{m-1}^{completed}, output_m) \quad (3)$$

Note that the partial output produced for the last array A_m is always completed since no subsequent array needs to be repaired.

An important observation is that computing each $d_i^{completed}$ requires processing $prefix(A_{i+1})$ twice, with the second processing serialized after d_i has been computed. The inefficiency is mitigated by the observation that the prefixes can be processed in parallel, so the critical path of the computation grows only by the processing of the constant prefix.

This scheme is applicable when there exists a constant prefix of each A_{i+1} that limits the scope of the necessary recomputation. Crucially, the prefix must not cover the whole A_{i+1} or else the repair of $fold(f, d_i, prefix(A_{i+1}))$ would modify d_{i+1} , which would in turn necessitate the repair of $fold(f, d_{i+1}, A_{i+2})$, serializing the repairs.

Example. Consider function `is-sorted` (shown in Fig. 2) that returns 1 if for each pair of consequent elements of array `A`, the former is smaller than the latter, and returns 0 otherwise. Function `f` takes `element` and `pair state` whose first element gets compared to `element`, and whose second element is either passed to `output` (if `element` is greater) or dropped to zero (otherwise).

The merge function is `min`: it takes array `out` of zeroes and ones, and returns 1 if all the arrays from $\{A_i\}$ are sorted (i.e., all elements of `out` are ones), and 0 if at least one array is not sorted (i.e., at least one element of `out` is zero). Each individual result is necessary but not sufficient for concluding that $(append\ A_1 \dots A_m)$ is sorted. It remains to establish that all elements of each A_i are smaller than any element of A_{i+1} , and we do it via constant prefixes.

The prefix of each A_{i+1} is used for repairing `outi` and it is defined as the array containing only the first element² of A_{i+1} . Note that the prefixes should be processed twice, e.g., for checking that $(append\ A_i\ (take\ A_{i+1}\ 1))$ is sorted and for checking that A_{i+1} itself is sorted. \square

We return to this scheme in Sect. 5 and show how function *merge* as well as constant $prefix_{length}$ can be synthesized.

²Following the ROSETTE notation, $(take\ A_{i+1}\ 1)$ stands for the sub-array of A_{i+1} containing just its first element.

is-sorted

```

1 ; given functions:
2   (define (f element state)
3     (cond
4       [(>= element (first state)) '(element (second state))]
5       [else '(0 0)]
6     )
7   )
8   (define (is-sorted A) (foldl f '(-inf.0 1) 1))
9
10 ; functions needed for parallel decomposition:
11   (define (merge-sorted out) (apply min out))
12   (define (prefix-length-sorted) 1)

```

Figure 2: Checking if array is sorted, a suitable function *merge* and a suitable value of *prefix.length*.

3.3 Decomposition with Conditional Prefixes

When there is no function *merge* meeting condition (3) for all possible constant prefixes, then there could exist more complicated functions to express the length of prefixes. Indeed, the number of elements in $prefix(A_i)$ might depend on the elements themselves. The problem of finding such *conditional* prefixes can be reduced to iterative evaluation of predicate $prefix_{cond}$ for each element of the given array:

$$prefix_{cond} : In \rightarrow bool$$

That is, for each i , the length of $prefix(A_i)$ is the position number k of some element in A_i , such that $prefix_{cond}$ evaluates to *true* for the k -th element, and $prefix_{cond}$ evaluates to *false* for all j -th elements in A_i where $j < k$. Thus, contrary to constant prefixes that could be identified by some static analysis of f , the calculation of conditional prefixes requires running f for the particular given arrays. For the tasks enjoying existence of a suitable predicate $prefix_{cond}$ and existence of a suitable function *merge*, the parallel recurrence decomposition property has the same form as (3). The difference is the way of computing prefixes, whose length is not constant any longer.

Example. Consider function *seen-2-after-1* that checks whether “2” appeared in the given array some time after “1” (shown in Fig. 3). Function f updates two flags indicating whether “1” or “2” has been already seen in the array.

The merge function is *max*: it takes the array out of zeroes and ones, and returns 1 if at least one array has “2” appeared some time after “1”. The prefix of each A_{i+1} contains all elements from the beginning of A_{i+1} until the first appearance of “2”. Indeed, consider a case when $m = 2$, A_1 contains “1”, but does not contain “2”, and A_2 contains “2”, but does not contain “1”. In this case, it is important to keep searching “2” in A_2 (i.e., traverse all elements until “2” is found. In contrast, processing A_1 and A_2 solely only lead to incorrect outputs. \square

We return to this scheme in Sect. 5 and show how function *merge* and predicate $prefix_{cond}$ can be synthesized.

seen-2-after-1

```

1 ; given functions:
2   (define (f element state)
3     (define seen-one (first state))
4     (define seen-two (second state))
5     (cond
6       [(= element 1) '(1 seen-two)]
7       [(= element 2) (if (= seen-one 1) '(seen-one 1) '(0 0))]
8       [else '(seen-one seen-two)]
9     )
10  )
11  (define (seen-2-after-1 A) (foldl f '(0 0) A))
12
13 ; functions needed for parallel decomposition:
14   (define (merge-search out) (apply max out))
15   (define (prefix-cond-search element) (= element 2))

```

Figure 3: Searching if “2” appeared in the array some time after “1” and the possible implementation of *merge* and *prefix*.

4 Bringing It All Together

Assuming the existence of function *merge* and (if applicable) constant $prefix_{length}$ or predicate $prefix_{cond}$ for some function f , we show how f can be parallelized in Fig. 4. The diagram considers three processors and represents different segments of the given data (original and updated with the prefixes) with rectangular boxes, and functions to iterate over data with ovals.

We estimate the time spent at different stages of the parallel algorithm as a means of number of *fold*-iterations. Due to (1), the total time T_s for three sequential applications of *fold* (denoted respectively s_1 , s_2 and s_3) performed for three segments is:

$$T_s = s_1 + s_2 + s_3$$

For the second and the third arrays, the calculation of the corresponding prefixes requires respectively times p_2 and p_3 . Since in the worst case the prefix calculation is an iterative process, p_i for some i might be equal to s_i .³ Note that every process should wait until prefixes for all arrays are computed, and the total time T_p is:

$$T_p = \max(p_2, p_3)$$

Then, for each array (already coupled with the subsequent prefix), the original function f should be iteratively applied. Here, every process should also wait for all other processes, and the total time T_f is:

$$T_f = \max(s_1 + p_2, s_2 + p_3, s_3)$$

³There could also be the case when a prefix spans over the entire segment, so its computation continues for the subsequent segments. We do not consider this scenario in this paper.

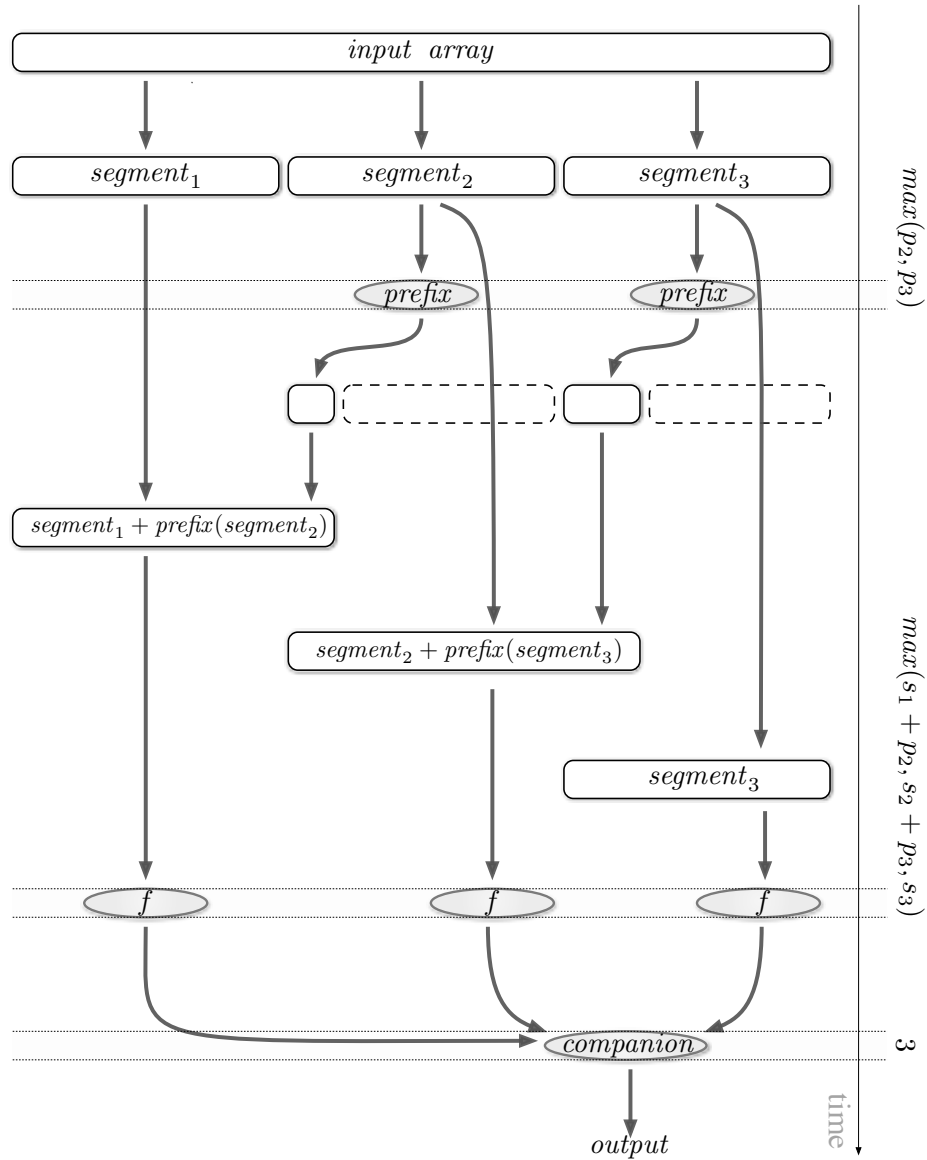


Figure 4: Executing function f in parallel.

The last step for executing *merge* (called *companion* in the diagram) takes time 3, since it just aggregates three integers:

$$T_c = 3$$

Finally, the speedup earned by the parallel version of the function compared to the sequential one can be calculated by the following formula:

$$X \triangleq \frac{T_s}{T_p + T_f + T_c}$$

5 Parallelization Synthesis Problem

Given function f and initial state d_0 , as declared in Sect. 2, we wish to find such function implementations for $prefix$ and $merge$, as declared in Sect. 3, so for any possible sequence of input arrays, $prefix$ and $merge$ correctly decompose recurrence relations foisted by f :

$$\begin{aligned} \exists prefix, merge \cdot \forall A_1, \dots, A_m \cdot \\ h(\text{fold}(f, d_0, \text{append}(A_1, \dots, A_m))) = \\ \text{merge}(h(\text{fold}(f, d_0, \text{append}(A_0, \text{prefix}(A_1)))), \dots, \\ h(\text{fold}(f, d_0, \text{append}(A_{m-1}, \text{prefix}(A_m))), h(\text{fold}(f, d_0, A_m))) \end{aligned}$$

5.1 Our Solutions

We present GRASSP, an algorithm to deliver solutions to the synthesis problem, and outline its pseudocode in Alg. 1. GRASSP aims at parallelizing iterative applications of f to d_0 for any possible input arrays A_1, \dots, A_m . It treats each A_i nondeterministically, by allowing them to contain only symbolic elements. Thus, while parallelizing $\text{fold}(f, d_0, \text{append}(A_1, \dots, A_m))$, the algorithm considers all possible resolution of nondeterminism in each A_i and if a solution is found, it is guaranteed to be general enough to satisfy all arrays containing concrete (i.e., non-symbolic) elements. To ensure the finiteness of the search space of the solutions, the lengths of A_i are bounded.

GRASSP exploits our observations made in Sect. 3, and gradually attempts synthesizing functions $merge$ and $prefix$ for f and d_0 under the following hypotheses:

- 1) there exists function $merge$ that makes (2) hold, so there is no need for prefixes (method SYNTNOPREFIX),
- 2) there exist function $merge$ and function $prefix$ with a fixed length $prefix_{length}$ that make (3) hold (method SYNTCONSTANTPREFIX),
- 3) there exist function $merge$ and function $prefix$ with predicate $prefix_{cond}$ that make (3) hold (method SYNTCONDITIONALPREFIX).

Each of the three methods verifies whether the corresponding hypothesis is true. In particular, it traverses the search space of candidate implementations of $merge$ (and, for the hypotheses 2 and 3, $prefix$) function and checks whether one of those candidates is a witness for the hypothesis. The increasing complexity of the methods allows saving time while parallelization, and delivering simple solutions first. To further improve efficiency, GRASSP bounds the search space of $merge$ to contain relatively simple operators (e.g., “+”, “min”, “max”), and the search space of $prefix_{cond}$ to contain conjunctions of simple terms in linear arithmetic with equality. Notably, the efficiency comes with a price: the more restrictions are applied to the search space, the more risks are taken by the algorithm to produce “unknown” output.

Given solutions to the synthesis problem, it is straightforward to compile the synthesized functions into a self-contained parallelized function that is equivalent to the sequential one, but behaves in a fashion described in Sect. 4.

5.2 Current Limitations and Future Improvements

Since prefixes are the key to identify the amount of overhead of parallel computation against sequential computation, GRASSP can be required to minimize the length of discovered prefixes. In case of constant

Algorithm 1: GRASSP (f, d_0)

Input: Function f , initial state d_0 **Output:** $merge, [prefix]$

- 1 $A_1, \dots, A_m \leftarrow \text{NONDET}()$
 - 2 Try **return** SYNTNOPREFIX(f, d_0, A_1, \dots, A_m) ▷ see Sect. 3.1
 - 3 Try **return** SYNTCONSTANTPREFIX(f, d_0, A_1, \dots, A_m) ▷ see Sect. 3.2
 - 4 Try **return** SYNTCONDITIONALPREFIX(f, d_0, A_1, \dots, A_m) ▷ see Sect. 3.3
 - 5 **return** “unknown”
-

prefixes, the algorithm can enumerate positive numbers starting from 0 and check whether the current number constitutes a sufficient prefix length for all arrays.

Interestingly, in case of conditional prefixes, it is not obvious how to construct the minimal prefix. Since depending on evaluation of $prefix_{cond}$, the length of each prefix can vary from array to array. Furthermore, $prefix_{cond}$ can be evaluated to *false* for all iterations of the i -th application of *fold*, implying that the entire A_{i-1} and A_i should be concatenated, and construction of prefixes should be continued with A_{i+1} . In general, the existence of $prefix_{cond}$ does not even guarantee that a prefix can eventually be found. And in the worst case, the entire input array should be processed sequentially.

As we mentioned in Sect. 3.2, for each i , the elements of $prefix(A_i)$ are processed twice: for completing the $(i-1)$ -th, and for starting the i -th applications of *fold*. For the big picture, we should also mention the routine for calculating $prefix(A_i)$ itself that requires yet another iterative run over the elements of A_i . To avoid such a redundancy, the prefix calculation engine can be augmented by construction of symbolic summaries that would capture the effect of applying f to the given prefix and an arbitrary initial state. A similar idea was presented in [17], but GRASSP could enable its application in a new context. That is, as opposed to summarizing the effect of application of *fold* to the entire arrays, we would like to summarize the effect of application of *fold* just to the prefixes, and to leave the remaining elements untouched. Then, the computed summaries could be applied directly 1) to repair the output of the $(i-1)$ -th application of *fold*, and 2) to start the i -th application of *fold*, thus avoiding duplicate iterations. We believe, such summarization-based approach to automatic parallelization is also suitable for the cases when a prefix spans over the entire arrays(s).

6 Evaluation

We implemented a prototype of GRASSP in ROSETTE. Given a specification in a form of array-handling sequential function f , ROSETTE maintains a space of candidate decompositions of f and verifies equivalence between f and each candidate separately. A candidate decomposition that fulfils the specification is returned by ROSETTE as output.

We evaluated GRASSP on a set of ROSETTE implementations of some prefix sum problems. We used a bound of 15 for the length of the input array. Interestingly, this bound was sufficient, and gave correct decompositions also for bigger arrays. In general, of course, the soundness of recurrence decomposition for big arrays is not guaranteed.

Table 1 summarizes our preliminary results of GRASSP for several benchmarks. It measures the complexity of each benchmark by giving the number of variables in the state of each f (“# Vars”). The

Benchmark	# Vars	Hypothesis	<i>merge</i>	$prefix_{length} / prefix_{cond}$	Synt time (s.)
array-count	1	SYNTNOPREFIX	+	—	3
array-max	1	SYNTNOPREFIX	max	—	3
is-sorted	1	SYNTCONSTPREFIX	min	1	3
alternation-of-1-2	1	SYNTCONSTPREFIX	min	1	2
number-of-123	2	SYNTCONSTPREFIX	+	2	3
seen-2-after-1	2	SYNTCONDPREFIX	max	(= element 2)	3
alternation-of-11-22	3	SYNTCONDPREFIX	min	(= element EOF)	4

Table 1: Preliminary evaluation of GRASSP.

results of recurrence decomposition are reported by the hypothesis which was true for the function, the operator for a synthesized function *merge*, and a synthesized function that identifies the length of the prefix (i.e., $prefix_{length}$ for the constant prefixes, and $prefix_{cond}$ for the conditional prefixes). Finally, we provide the total time spent by ROSETTE for realizability check of the correspondent hypothesis (i.e., to realize the necessity of either constant or conditional prefixes), and for synthesis of the witnessing functions.

Functions `array-max`, `is-sorted` and `seen-2-after-1` were already discussed in Examples 1, 2, and 3 respectively. Function `array-count` calculates the size of the array. Functions `alternation-of-1-2` and `alternation-of-11-22` check if the entire array is an alternation of “1” and “2” and “11” and “22” respectively. Function `number-of-123` searches for appearances of word “123”. Interestingly, the synthesis time for all decompositions was insignificant. In future work, we plan to enhance the set of experiments by more challenging benchmarks.

7 Related work

The problem of parallelizing recurrent programs dates back to seventies [13]. Today, from purely mathematical solutions it grew up to the mature large-scale industrial applications within MapReduce [7], Dryad [12], Spark [26] and Hadoop [24] distributed programming platforms. To automatically generate MapReduce programs, [16] proposes to translate sequential code into parallel one based on the set of rewrite rules. Alternatively, the MapReduce program synthesis can be driven by input/output examples [19]. Instead of synthesizing the MapReduce code, [17] proposes to synthesize symbolic summaries over concrete segments of the input data, the combination of which in a predefined order produces the final output. GRASSP can be also viewed as a technique to synthesize a MapReduce-like code, and in contrast to the aforementioned works, it exploits template-based synthesis and bounded model checking.

Other approaches targeting synthesis of concurrent programs include [23, 3, 25]. In particular, [25] deals with partial SPMD programs which are automatically converted into sequential partial programs, for which the requested code fragments are synthesized. The approach of [23] synthesizes synchronization primitives to be used in concurrent programs. The approach of [3] performs automatic vectorization, i.e., replaces loop-free blocks of code from the loop bodies with sequences of synthesized SIMD instructions. All these works use formal methods, but they are applicable to the different classes of problems than the one of GRASSP.

Programming By Examples (PBE) is nowadays one of the most promising directions in synthesis [11, 2, 9, 1, 19]. Since the applications of PBE do not require an explicit specification (contrary

to GRASSP that treats the sequential program as specification), PBE is orthogonal to our synthesis approach. However, one can still find a connection between bounded synthesis and PBE. Indeed, while fixing a bound for the number of elements in the given array, we force GRASSP to implicitly consider all possible examples consisting of the bounded arrays and the corresponding outputs of the given sequential program.

Formal methods are getting more involved into the data-parallel and array-handling computing: [4] studies the commutativity problem of MapReduce, [6] proposes an approach to prove counting properties in unbounded array-handling programs. Finally, equivalence checking is one of the most intriguing branches of formal methods (e.g., [14, 18, 5, 10, 15, 8]). Currently, GRASSP uses equivalence checking only in the bounded context, to verify that a candidate decomposition is equivalent to the given sequential program. In future, it would be interesting to see which unbounded equivalence checking methods can also be applicable to GRASSP.

8 Conclusion

In this paper, we addressed the challenge of bringing the existent sequential array-handling code into a parallel fashion. We presented an approach called GRASSP that applies for the cases when the data is partitioned into a sequence of segments, each segment processed separately, and the partial outputs for the segments are merged together. GRASSP augments the given sequential program with an additional functionality to decompose data dependencies in loop iterations. GRASSP gradually considers several parallelization scenarios and attempts to find easier solutions first. We presented a prototype of the algorithm GRASSP to automatically synthesize decompositions of small programs in Racket, and we envision its further improvements in future.

Acknowledgments. This work is supported in part by the SNSF Fellowship P2T1P2_161971, NSF Grants CCF-1139138, CCF-1337415, and NSF ACI-1535191, a Grant from U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences Energy Frontier Research Centers program under Award Number FOA-0000619, and grants from DARPA FA8750-14-C-0011 and DARPA FA8750-16-2-0032, as well as gifts from Google, Intel, Mozilla, Nokia, and Qualcomm.

References

- [1] Rajeev Alur, Pavol Cerný & Arjun Radhakrishna (2015): *Synthesis Through Unification*. In: *CAV, LNCS 9207*, Springer, pp. 163–179, doi:10.1007/978-3-642-31612-8.
- [2] Daniel W. Barowy, Sumit Gulwani, Ted Hart & Benjamin G. Zorn (2015): *FlashRelate: extracting relational data from semi-structured spreadsheets using examples*. In: *PLDI*, ACM, pp. 218–228, doi:10.1145/2737924.2737952.
- [3] Gilles Barthe, Juan Manuel Crespo, Sumit Gulwani, César Kunz & Mark Marron (2013): *From relational verification to SIMD loop synthesis*. In: *PPoPP*, ACM, pp. 123–134, doi:10.1145/2442516.2442529.
- [4] Yu-Fang Chen, Lei Song & Zhilin Wu (2016): *The Commutativity Problem of the MapReduce Framework: A Transducer-Based Approach*. In: *CAV, 9780*, Springer, pp. 91–111, Part II, doi:10.1007/978-3-319-41540-6_6.
- [5] Edmund M. Clarke, Daniel Kroening & Karen Yorav (2003): *Behavioral consistency of C and Verilog programs using bounded model checking*. In: *DAC*, ACM, pp. 368–371, doi:10.1145/775832.775928.

- [6] Przemyslaw Daca, Thomas A. Henzinger & Andrey Kupriyanov (2016): *Array Folds Logic*. In: *CAV*, 9780, Springer, pp. 230–248, Part II, doi:10.1007/978-3-319-41540-6.
- [7] Jeffrey Dean & Sanjay Ghemawat (2004): *MapReduce: Simplified Data Processing on Large Clusters*. In: *OSDI*, USENIX Association, pp. 137–150, doi:10.1145/1327452.1327492.
- [8] Grigory Fedyukovich, Arie Gurfinkel & Natasha Sharygina (2016): *Property Directed Equivalence via Abstract Simulation*. In: *CAV*, 9780, Part II, Springer, pp. 433–453, doi:10.1007/978-3-319-41540-6.
- [9] John K. Feser, Swarat Chaudhuri & Isil Dillig (2015): *Synthesizing data structure transformations from input-output examples*. In: *PLDI*, ACM, pp. 229–239, doi:10.1145/2737924.2737977.
- [10] Benny Godlin & Ofer Strichman (2009): *Regression verification*. In: *DAC*, ACM, pp. 466–471, doi:10.1145/1629911.1630034.
- [11] Sumit Gulwani, William R. Harris & Rishabh Singh (2012): *Spreadsheet data manipulation using examples*. *Commun. ACM* 55(8), pp. 97–105, doi:10.1145/2240236.2240260.
- [12] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell & Dennis Fetterly (2007): *Dryad: distributed data-parallel programs from sequential building blocks*. In: *EuroSys*, ACM, pp. 59–72, doi:10.1145/1272996.1273005.
- [13] Peter M. Kogge (1974): *Parallel Solution of Recurrence Problems*. *IBM Journal of Research and Development* 18(2), pp. 138–148, doi:10.1147/rd.182.0138.
- [14] A. Kuehlmann & F. Krohm (1997): *Equivalence checking using cuts and heaps*. In: *DAC*, IEEE, pp. 263–268, doi:10.1145/266021.266090.
- [15] Shuvendu K. Lahiri, Chris Hawblitzel, Ming Kawaguchi & Henrique Rebêlo (2012): *SYMDIFF: A Language-Agnostic Semantic Diff Tool for Imperative Programs*. In: *CAV*, LNCS 7358, pp. 712–717, doi:10.1007/978-3-642-31424-7_54.
- [16] Cosmin Radoi, Stephen J. Fink, Rodric M. Rabbah & Manu Sridharan (2014): *Translating imperative code to MapReduce*. In: *OOPSLA*, ACM, pp. 909–927, doi:10.1145/2660193.2660228.
- [17] Veselin Raychev, Madanlal Musuvathi & Todd Mytkowicz (2015): *Parallelizing user-defined aggregations using symbolic execution*. In: *SOSP*, ACM, pp. 153–167, doi:10.1145/2815400.2815418.
- [18] Rahul Sharma, Eric Schkufza, Berkeley R. Churchill & Alex Aiken (2013): *Data-driven equivalence checking*. In: *OOPSLA*, ACM, pp. 391–406, doi:10.1145/2509136.2509509.
- [19] Calvin Smith & Aws Albarghouthi (2016): *MapReduce program synthesis*. In: *PLDI*, ACM, pp. 326–340, doi:10.1145/2908080.2908102.
- [20] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodík, Sanjit A. Seshia & Vijay A. Saraswat (2006): *Combinatorial sketching for finite programs*. In: *ASPLOS*, ACM, pp. 404–415, doi:10.1145/1168857.1168907.
- [21] Emina Torlak & Rastislav Bodík (2013): *Growing solver-aided languages with Rosette*. In: *Onward!*, ACM, pp. 135–152, doi:10.1145/2509578.2509586.
- [22] Emina Torlak & Rastislav Bodík (2014): *A lightweight symbolic virtual machine for solver-aided host languages*. In: *PLDI*, ACM, p. 54, doi:10.1145/2594291.2594340.
- [23] Martin T. Vechev & Eran Yahav (2008): *Deriving linearizable fine-grained concurrent objects*. In: *PLDI*, ACM, pp. 125–135, doi:10.1145/1375581.1375598.
- [24] Tom White (2009): *Hadoop - The Definitive Guide: MapReduce for the Cloud*. O’Reilly. Available at <http://www.oreilly.de/catalog/9780596521974/index.html>.
- [25] Zhilei Xu, Shoab Kamil & Armando Solar-Lezama (2014): *MSL: A Synthesis Enabled Language for Distributed Implementations*. In: *SC*, IEEE, pp. 311–322, doi:10.1109/SC.2014.31.
- [26] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauly, Michael J. Franklin, Scott Shenker & Ion Stoica (2012): *Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing*. In: *NSDI*, USENIX Association, pp. 15–28.