

Ranger: Parallel Analysis of Alloy Models by Range Partitioning

Nicolás Rosner
Department of Computer Science,
FCEyN, UBA,
Argentina,
nrosner@dc.uba.ar

Junaid H. Siddiqui
Department of Computer Science,
LUMS School of Science and Engineering,
Pakistan,
junaid.siddiqui@lums.edu.pk

Nazareno Aguirre
Department of Computer Science,
FCEFQyN, UNRC,
Argentina,
naguirre@dc.exa.unrc.edu.ar

Sarfraz Khurshid
Department of Electrical and Computer Engineering,
The University of Texas at Austin,
USA,
khurshid@ece.utexas.edu

Marcelo F. Frias
Department of Software Engineering,
Instituto Tecnológico de Buenos Aires,
Argentina,
mfrias@itba.edu.ar

Abstract—We present a novel approach for parallel analysis of models written in Alloy, a declarative extension of first-order logic based on relations. The Alloy language is supported by the fully automatic Alloy Analyzer, which translates models into propositional formulas and uses off-the-shelf SAT technology to solve them. Our key insight is that the underlying constraint satisfaction problem can be split into subproblems of lesser complexity by using *ranges* of candidate solutions, which partition the space of all candidate solutions. Conceptually, we define a total ordering among the candidate solutions, split this space of candidates into ranges, and let independent SAT searches take place within these ranges’ endpoints. Our tool, **Ranger**, embodies our insight. Experimental evaluation shows that **Ranger** provides substantial speedups (in several cases, superlinear ones) for a variety of hard-to-solve Alloy models, and that adding more hardware reduces analysis costs almost linearly¹.

Index Terms—Static analysis, Alloy, Parallel analysis, SAT-solving.

I. INTRODUCTION

Declarative formal models of software are valuable at a number of different stages of the software development process. They are particularly useful during requirements elicitation, as a means to express requirements in a language that is precise and expressive enough to document the needs of the stakeholders. Moreover, declarative models allow us to document rationales behind design decisions, and even to analyze properties of software designs prior to implementation. During the implementation phase, declarative models allow programmers to document expected properties of their classes and methods, e.g., using class invariants, method contracts, and loop invariants, which can also be exploited for different kinds of analyses.

¹Authors of this article have submitted a second paper on efficient computation of tight bounds for bounded verification of code. While that article uses Alloy as an intermediate language, its goal is to improve the sequential analysis of procedural code. This article, instead, deals with parallel analysis of Alloy declarative models, an unrelated problem.

A number of modeling languages today allow writing formal, declarative models [18], [34], [25], [22], [2]. Our specific focus is Alloy [18], a declarative extension of first-order based on relations. Alloy’s small yet expressive notation, together with its fully automated SAT-based Alloy Analyzer tool [3], make the language particularly appealing for modeling and analysis. Indeed, Alloy has already been used effectively in requirements [19], [39], design [24], [20], testing [21], [1], and as an intermediate language in static program analysis [10], [8], [13], [14], [27]. Section II provides the necessary details regarding the Alloy language and the Alloy Analyzer.

Alloy’s analysis technique, known as *scope-bounded checking*, analyzes a model correctness with respect to a bounded universe of discourse, by searching for violations of assertions that the user may expect to hold in the model. The assertions on a model are evaluated on model instances whose domains are bounded in size. The bound on the size of model instances is termed the *scope*, and is given by the user. Clearly, assertions that pass the analysis are not necessarily valid in general – they are valid for the given scope. Thus, to enhance their confidence in the correctness of their models, Alloy users must run the analysis for larger scopes. However, the cost of the underlying SAT-based analysis is exponential in those bounds, and the analysis bounds that can be used in practice for a number of Alloy models tend to be quite small. Our work in this paper is driven by our desire to effectively leverage the increasing availability of commodity hardware for effective parallelization schemes, and to enable Alloy analysis to scale better.

We present **Ranger**, a novel parallel analysis technique for Alloy models, based on what we call *range partitioning*. Essentially, the technique relies on the definition of a linear ordering on the state space of an Alloy model. A partition of the state space can then be defined by splitting the linear ordering into non-overlapping intervals, called *ranges*. Each

restriction of the original problem to a particular range thus becomes an independent subproblem, which can be analyzed by a separate processor on a cluster of computers. Further details are presented in Section III. We also discuss some more technical, implementation-related issues in Section IV.

We perform an experimental evaluation of our approach for range partitioning using a benchmark consisting of 10 hard-to-analyze properties from 7 different Alloy models (Section V). The benchmark includes valid and invalid cases (i.e., unsatisfiable and satisfiable problems) from a variety of problem domains. We show that, for 64 workers, the average speedup obtained by Ranger is XXX, including XXX cases where superlinear speedups are achieved. In all cases, Ranger was able to analyze the assertions for scopes that exceed the capabilities of the Alloy Analyzer; in some of the cases, it was able to do so for scopes that stand no chance whatsoever of being tractable by the Analyzer.

In Section VI we discuss existing techniques aimed at improving the scalability of Alloy, including parallel analysis tools and techniques related or applicable to Alloy. Finally, our conclusions and some proposals for further work are presented in Section VII.

In summary, this paper makes the following contributions:

- **Range partitioning.** We introduce the idea of distributing an Alloy problem into several subproblems of lesser complexity by defining *ranges* of candidate solutions;
- **Parallel analysis for Alloy.** We present Ranger, our technique for effective parallelization of Alloy problems using dynamic work stealing;
- **Experimental evaluation.** We embody Ranger into a prototype implementation and present experimental results that show the effectiveness of Ranger in analyzing a variety of Alloy models.

II. ALLOY AND THE ALLOY ANALYZER

Alloy is a declarative modeling language whose syntax incorporates features that are ubiquitous in object orientation. This amenable syntax is complemented with a relational semantics whose comprehension requires elementary concepts from discrete mathematics. More formally put, Alloy’s relational logic is an extension of first-order logic with reflexive-transitive closure. In the remaining part of this section we introduce Alloy’s syntax and semantics through an example. Figure 1 shows an Alloy model for a heap allocated binary tree data structure.

Data domains are defined using *signatures* (denoted by the keyword `sig`), which are represented as sets. Signature `Node`, for instance, declares a set of node objects. Akin to classes in object oriented languages, signatures may extend other signatures, in which case domains defined by the extending signatures are subsets of the domains defined by the extended ones. Signatures may be *abstract*. Domains defined by abstract signatures only contain elements that belong to extending signatures. Like classes, signatures may contain fields, which are captured by relations. For example, field `root` denotes a (total and functional) binary relation contained

```

module binTrees

one sig null {}

abstract sig Object {}

sig BinTree extends Object {
  root : Node + null }

sig Node extends Object {
  left, right : Node + null }

pred Acyclic[t : BinTree] {
  all n : t.root.*(left + right) |
  n !in n.^(left + right) &&
  (n.(left) & n.(right)) in null &&
  (n != null => (lone n.^(left + right))) }

pred NumNodesEqualsNumEdgesPlusOne[t: BinTree] {
  // in a non-empty tree, numNodes = numEdges + 1
  t.root != null =>
  #(t.root.*(left+right)-null) =
  #(left.Node)+#(right.Node)+1 }

pred NoUnreachableNodes[t : BinTree] {
  t.root.*(left+right) = (Node + null) }

fact { all t : BinTree | NoUnreachableNodes[t] }

check { all t : BinTree |
  Acyclic[t] <=> NumNodesEqualsNumEdgesPlusOne[t]
} for 0 but 1 BinTree, exactly 5 Node

```

Fig. 1. A sample Alloy model for binary trees.

in $\text{BinTree} \times (\text{Node} \cup \text{null})$. It is worth emphasizing that Alloy fields may also denote relations of arity greater than 2. Predicates allow us to name properties, while functions name terms. They may be combined to write axioms, which are called *facts* in Alloy.

Alloy expressions are built using set-theoretical and relational operators and constants. Constants `univ`, `iden` and `none` denote the set containing all elements, the identity binary relation on such set, and the empty set, respectively. Operations `+`, `-` and `&` denote set union, difference and intersection. Relational operators include composition (called *navigation* in Alloy), transpose, and (reflexive-)transitive closure. They are defined as follows²:

$$\begin{aligned}
 R.S &= \{ \langle a_1, \dots, a_{n-1}, b_2, \dots, b_m \rangle : a_n = b_1 \wedge \\
 &\quad \langle a_1, \dots, a_n \rangle \in R \wedge \langle b_1, \dots, b_m \rangle \in S \} \text{ (navigation)} \\
 \sim R &= \{ \langle b, a \rangle : \langle a, b \rangle \in R \} \text{ (transpose)} \\
 \wedge R &= \bigcup_{i>0} R^i \text{ (transitive closure)} \\
 *R &= \bigcup_{i\geq 0} R^i \text{ (reflexive-transitive closure)}
 \end{aligned}$$

Notice that transpose and closures are defined in Alloy only for binary relations.

²We define $R^0 = \text{iden}$ and $R^i = \underbrace{R \dots R}_{i \text{ times}}$.

Function # computes the cardinality of a relation. For example, term $\#(t.\text{root}.*(\text{left}+\text{right})-\text{null})$ in predicate $\text{NumNodesEqualsNumEdgesPlusOne}$ denotes the number of nodes reachable from the root of tree t by traversing t along fields left and right .

Alloy formulas are built from the atomic predicate in (inclusion), using standard connectives from first-order logic. Java notation is used for propositional connectives. Quantifiers are denoted by all (universal) and some (existential).

Notice that our sample model also includes an assertion – a property that is expected to hold in valid instances of the model (i.e., those satisfying the structural constraints imposed by signature definitions, and the model facts). Checking the validity of assertions is a means to verify the correctness of the model. As explained in Section I, assertions are analyzed with the aid of the Alloy Analyzer and within user-prescribed bounds. Check commands are issued in the model, and set the bounds for data domains. In this example, the Analyzer will analyze all configurations with at most one tree and exactly 5 nodes.

III. RANGE PARTITIONING

In this section we present *range partitioning*, our new technique for parallel analysis of Alloy models. As we will show later on, this technique provides a substantial improvement to the scalability of the SAT-based analysis, compared to the sequential Alloy Analyzer.

The technique is essentially composed of the following stages:

- Given an Alloy model, all candidate configurations that would be explored by the Analyzer are linearly ordered. This step establishes a linear ordering C_1, C_2, \dots, C_n (notice that the number of configurations, although usually very large, is finite due to the imposed bounds).
- We select arbitrary configurations $C_{j_1}, C_{j_2}, \dots, C_{j_i}$, and use them to split the above ordering into ranges using the arbitrary configurations as partition points (notice that we do not demand these configurations to satisfy the model axioms). We then obtain ranges $[C_1, C_2, \dots, C_{j_1}], [C_{j_1}, \dots, C_{j_2}], \dots, [C_{j_i}, \dots, C_n]$.
- The source Alloy model is constrained in such a way that the resulting models correspond to the different ranges; these models are distributed to different processors and analyzed in parallel.

The above described process requires addressing the following technical challenges:

- Define a linear ordering on the set of candidate configurations.
- Provide an algorithm for selecting appropriate configurations as partition points (ideally, we want ranges to contain roughly the same number of configurations).
- Solve the distribution problem in an efficient way.

We will deal with each one of these challenges in Sections III-A–III-C.

A. A Linear Ordering on Configurations

We begin this section by describing how configurations are internally handled by the Alloy Analyzer. The Alloy Analyzer translates models to Kodkod’s [36] language. From the scopes in the check command, Kodkod generates a uniform naming for atoms in each domain. For the example in Fig. 1, Kodkod produces the naming in Table I.

| Sig | scope | naming |
|--------|-------|-----------------------|
| Object | 1 | Object\$0 |
| null | 1 | null\$0 |
| Node | 5 | Node\$0, ..., Node\$4 |

TABLE I
NAMING TABLE FOR THE MODEL FROM FIG. 1.

From the naming and other information (see Section IV-A), Ranger builds a vector specification (vecSpec), i.e., a mapping

$$\text{vecSpec} : \text{AtomNames} \times \text{RelNames} \rightarrow \mathcal{P}(\text{AtomNames})$$

that, given an atom name n and a relation name R , retrieves the atom names that may be considered as the result of computing $n.R$. A vecSpec is then used to capture the state space of configurations. For the sake of simplifying the presentation, we will restrict ourselves to total and functional signature fields. Notice that, in the example from Fig. 1, all fields satisfy this constraint. Figure 2 provides a graphical representation of the vecSpec associated to the model in Fig. 1, as an actual vector.




Fig. 2. Graphical representation of a vecSpec .

In the figure, $\text{Object}\$0$ may relate to $\text{null}\$0, \text{Node}\$0, \dots, \text{Node}\$4$ via relation root . No element relates (via any relation) to $\text{Object}\$0$. While a vecSpec describes the state space, a *configuration* is a particular state. Configurations can be described by choosing, for each entry in the vecSpec , one of the possible values. In Fig. 3 we show a configuration vector as well as the binary tree described by the configuration.

We take the ordering in which Kodkod lists atom names as the strict linear ordering on atom names. We will denote this ordering on atom names by $<_K$. This ordering can be extended to a lexicographical ordering between configurations (denoted $<_{LC}$) as follows (notice that all vector configurations have the same size, which we denote by s):

$$C <_{LC} C' \iff (\exists i, 0 \leq i < s)(C[0, i-1] = C'[0, i-1] \wedge C[i] <_K C'[i]).$$




Fig. 3. Sample configuration for the model in Fig. 1.

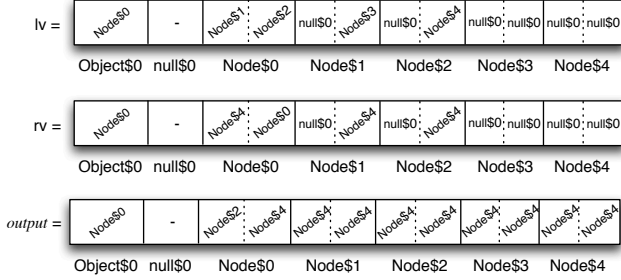


Fig. 4. Application of partitioning (standard case).

THEOREM 1. *Relation $<_{LC}$ is a strict total order on the set of configurations.*

Proof. We must prove that $<_{LC}$ is irreflexive, transitive and total. Irreflexivity and totality follow from the irreflexivity and totality of $<_K$, respectively. Let us focus, then, on transitivity. Let $C_1 <_{LC} C_2$ and $C_2 <_{LC} C_3$. Let i_0, i_1 such that $C_1[0, i_0 - 1] = C_2[0, i_0 - 1] \wedge C_1[i_0] <_K C_2[i_0]$ and $C_2[0, i_1 - 1] = C_3[0, i_1 - 1] \wedge C_2[i_1] <_K C_3[i_1]$. Let $I = \min(i_0, i_1)$. Notice that $C_1[0, I - 1] = C_3[0, I - 1]$. If $I = i_0 < i_1$, $C_1[I] = C_1[i_0] <_K C_2[i_0] = C_3[i_0]$. If $I = i_1 < i_0$, $C_1[I] = C_1[i_1] = C_2[i_1] <_K C_3[i_1]$. If $I = i_0 = i_1$, then $C_1[I] = C_1[i_0] <_K C_2[i_0] = C_2[i_1] <_K C_3[i_1]$. Thus, since $<_K$ is transitive, $C_1[I] <_K C_3[I]$, which implies $C_1 <_{LC} C_3$. \square

Theorem 1 allows us to adopt $<_{LC}$ as the strict linear ordering on configurations.

B. Selection of the Partition Points

Partition points are configurations that serve as boundaries between ranges. In this section we show how, given a range $R = [C_1, C_2]$ (with $C_1 <_{LC} C_2$) and a positive number n , we can select configurations $X_1, \dots, X_{n-1} \in R$ such that ranges $[C_1, X_1], [nextConf(X_1), X_2], \dots, [nextConf(X_{n-1}), C_2]$ are all contained in $[C_1, C_2]$ and are somewhat balanced with respect to the number of configurations they contain. We first show, in Alg. 1, pseudocode for partition implementation of a range into two subranges. Lines 1–15 describe the most frequent case, where the first position in which the vectors disagree contains elements that are far apart enough that a middle element can readily be found. The scenario, as well as the result returned by Alg. 1 are illustrated in Fig. 4.

Notice that once a partition point has been found, the source range $[C_1, C_2]$ is split into the ranges $[C_1, C]$ and

```

1 Config binRangePartition(lv, rv : Config, vs : vecSpec)
2   i = 0; // Skip common prefix
3   while lv[i] == rv[i] do
4     i = i + 1;
5   end
6   fdl, fdr = lv[i], rv[i]; // First values that differ
7   options = [x ∈ vs[i] : fdl <_K x <_K fdr]; // Sorted list
8   if options != ∅ then
9     midOption = options[(len(options) - 1)/2];
10    output[0, i - 1] = lv[0, i - 1];
11    output[i] = midOption;
12    for i < j < len(vs) do
13      | output[j] = max(vs[j]);
14    end
15    return output;
16  else
17    // Values at pos i differ by 1
18    // e.g., lv=[3,8,5,...], rv=[3,8,6,...]
19    if i == len(vs) then
20      | return lv;
21    end
22    i = i + 1;
23    while lv[i] == max(vs[i]) && rv[i] == min(vs[i]) do
24      i = i + 1;
25      if i == len(vs) then
26        // e.g. lv=[3,8,5,9,9,9]
27        // and rv=[3,8,6,0,0,0]
28        return lv;
29      end
30    end
31    if lv[i] != max(vs[i]) then
32      // e.g. lv=[3,8,5,9,9,7,...]
33      // and rv=[3,8,6,0,0,3,...]
34      output[0, i - 1] = lv[0, i - 1];
35      for i ≤ j < len(vs) do
36        | output[j] = max(vs[j]);
37      end
38      return output;
39    else
40      // e.g. lv=[3,8,5,9,9,9,...]
41      // and rv=[3,8,6,0,0,3,...]
42      output[0, i - 1] = rv[0, i - 1];
43      for i ≤ j < len(vs) do
44        | output[j] = min(vs[j]);
45      end
46      return output;
47    end
48  end
49 end

```

Algorithm 1: The binary partitioning algorithm.

$[nextConf(C), C_2]$. Pseudocode for method $nextConf$, which retrieves the next configuration according to ordering $<_{LC}$, is presented in Alg. 2. Intuitively, the behavior of Alg. 2 is quite similar to adding 1 in elementary arithmetic – yet, instead of using digits 0 thru 9 (or those of any other fixed, uniform base), the algorithm uses, for each cell of a vector, the sorted list of options available for that cell according to the $vecSpec$. The reader should keep in mind that the options for each cell can be an arbitrary subset of the set of all atoms.

We now focus on generation of range partitions in the general case. Algorithm 3 presents pseudocode for the actual partition algorithm. The algorithm starts from an input range, and iterates over a list of already generated ranges as long as the requested number of ranges has not been reached and there are still some ranges left to be split (ranges of the form $[C, C]$ cannot be split any further).

```

1 Config nextConf(c : Config, vs : vecSpec)
  // c must not be the last config.
2   i = len(vs) - 1;
3   output = c;
4   while true do
5     options = [x ∈ vs[i]]; // Sorted list
6     pos = position of options[i] in options;
7     if pos < len(options) - 1 then
8       output[i] = options[pos + 1];
9       return output;
10    else
11      output[i] = options[0];
12      i = i - 1;
13    end
14  end
15 end

```

Algorithm 2: Pseudocode for method *nextConf*.

```

1 List[Range] rangePartitioning(R : Range, n : int, vs : vecSpec)
2   L = addToEmptyList(R); // L = [R]
3   i = 0; // index for traversing the list
4   while len(L) < n && hasSplittableRanges(L) do
5     if splittable(L[i]) then
6       Config C1 = leftEndpoint(L[i]);
7       Config C2 = rightEndpoint(L[i]);
8       Config C3 = binRangePartition(C1, C2, vs);
9       removeRangeAtPos(L, i);
10      addRangeAtPos(L, i, [C1, C3]);
11      addRangeAtPos(L, i + 1, [nextConf(C3, vs), C2]);
12      if i < len(L) - 2 then
13        i = i + 2;
14      else
15        i = 0;
16      end
17    else
18      if i < len(L) - 1 then
19        i = i + 1;
20      else
21        i = 0;
22      end
23    end
24  end
25 end

```

Algorithm 3: The range partitioning algorithm.

C. Parallel Range Analysis

In Section III-B we proposed a method to partition a range into sub-ranges; we now explain how to use the aforementioned in order to parallelize the analysis of an Alloy model.

Our scenario for parallel analysis makes use of a cluster of computers. Worker processes execute commands sent by a master process, which runs on a dedicated processor. Actions that workers can perform include: solving a task sequentially, aborting the ongoing analysis, splitting the current task into a given number of ranges (and locally enqueueing the resulting subtasks), fetching a new task from the local queue, and moving tasks between local and remote queues (requesting and obtaining tasks from other workers). Further details on *RangeR*'s implementation will be provided in Section IV.

In the remaining parts of this section we present two alternatives for parallelization, both based on range partitioning. Before doing so, we discuss the impact of Alloy's symmetry-breaking predicates on the ranges generated by partitioning.

1) *Range Partitioning and Symmetry Breaking:* Atom names are irrelevant when building an Alloy configuration: given a configuration that satisfies (or not) the facts of an Alloy model, any other configuration obtained by mere permutation of atom names (while, of course, preserving typing constraints) will behave in the same way. Symmetry-breaking axioms are introduced by Kodkod [36] during the translation of the Alloy model to a propositional formula, and greatly improve the performance of the underlying SAT-solver by avoiding the exploration of many such superfluous isomorphisms. Therefore, symmetry-breaking has a direct impact on what non-superfluous configurations will look like. For example, for the model in Fig. 1, field *root* can only relate to atom names *null\$0* or *Node\$4*. Hence, any ranges in which *root* points to nodes *Node\$0*, *Node\$1*, *Node\$2* or *Node\$3* contain configurations that cannot satisfy the propositional formula generated by Kodkod. For our sample Alloy model from Fig. 1 (but using a scope of 10 *Node*, since 5 *Node* is too easy), Table II shows how partitioning the full range into increasingly large numbers of ranges yields very small numbers of nontrivial subproblems. The remaining ranges can all be proved unsatisfiable in under a millisecond each. As we will see in Sections III-C2 and III-C3, the fact that a significant number of the tasks resulting from a partition may become trivial can lead to hardware being severely underused. In order to measure the actual utilization of the assigned hardware during parallel analysis, we define the metric

$$\text{Hardware Use Efficiency} = \frac{\text{Total non-idle seconds}}{\text{Number of workers} \times t},$$

where *t* is the wallclock runtime in seconds taken by the whole parallel analysis (as perceived by the end user), and the numerator is the sum, over all workers, of the number of seconds during which some task was actually being analyzed.

| Num. Ranges | 512 | 1024 | 2048 | 4096 | 8192 | 16384 |
|--------------|-------|-------|-------|-------|-------|-------|
| Nontrivial | 10 | 20 | 35 | 56 | 66 | 90 |
| Nontrivial % | 0.019 | 0.019 | 0.017 | 0.013 | 0.008 | 0.005 |

TABLE II
RANGE PARTITIONING: NUMBER AND PERCENTAGE OF NONTRIVIAL SUBPROBLEMS (OF THOSE GENERATED FOR BINTREES WITH 10 *Node*).

2) *Flat Range Partitioning:* One way of parallelizing the analysis based on range partitioning consists in determining a large enough value of *n* (called the *fan-out* of the analysis), and having a worker partition the full range into *n* ranges using algorithm *rangePartitioning*. The *n* newly generated tasks are then solved in parallel by the available workers. Each worker receives a task and solves it sequentially until a SAT/UNSAT verdict is obtained. Unfortunately, this approach seldom performs as well as expected. In most cases the solving time variance (hardest vs. easiest subproblems obtained after initial partitioning) is very high. Thus, if the initial range is only partitioned once, eventually only a few active workers will remain, while (almost all) other workers will be idle until the end of the run. Table III reports the HUE for analysis of the model in Fig. 1 when flat range partitioning is used.

| Num. Ranges | 512 | 1024 | 2048 | 4096 | 8192 | 16384 |
|-------------|-----|------|------|------|------|-------|
| HUE | | 0.04 | 0.04 | 0.11 | 0.11 | 0.12 |

TABLE III

EFFICIENCY RATES OBTAINED FOR FLAT RANGE PARTITIONING OF BINTREE EXAMPLE, WITH 10 NODE, USING 64 WORKERS.

3) *Recursive Range Partitioning*: The main drawback of flat range partitioning is its static nature. Trying to balance the number of candidate configurations in each range is indeed a reasonable starting point, but we cannot predict, in the general case, where the harder subranges might lie. A more dynamic approach to determining the location of nontrivial subranges is therefore desirable. In recursive range partitioning, the oldest active subproblem (i.e., the oldest among those that are still being SAT-solved) can be re-partitioned by its assigned worker. This yields sub-subproblems, and so on, recursively. Recursive partitioning of a range may occur under two circumstances:

- the UNSAT frequency, i.e. the number of UNSAT verdicts per second, falls below a user-defined threshold, or
- there are idle workers.

The first condition aims at achieving progress during analysis by avoiding to waste time analyzing tasks that are still too hard to be solved sequentially by a worker. The second condition targets the HUE metric and strives at making good use of resources by avoiding idle workers.

Unlike flat partitioning, where the fan-out must be large and fixed beforehand, in recursive partitioning we use a small fan-out; its value is set to the number of workers. For instance, in the experiments reported in Section V, the fan-out is 64, since that is the total number of worker cores in the cluster. In this way, whenever many of the 64 tasks turn out to be trivial (or shortly after idle workers start to abound) the recursive partitioning process will react by “zooming in”, i.e. focusing the computational effort, on the remaining nontrivial tasks.

Using recursive range partitioning, the HUE value for our example (for the same 10-Node scope) becomes 0.84, which is about 8 times higher than with flat range partitioning. Even if we only counted time invested in *successful* solving attempts as non-idle time (i.e., if the partial solving attempts before re-splitting were to be considered completely wasted effort), the HUE value for this run would be 0.59 – still a significant improvement over the low efficiency of flat partitioning.

The following theorem discusses the correctness of recursive range partitioning. A detailed proof is omitted due to space constraints.

THEOREM 2. *Recursive range partitioning is sound and complete, i.e., an Alloy model has a satisfying configuration if and only if one can be found using recursive range partitioning.*

Proof sketch: Recursive range partitioning splits ranges whenever tasks are aborted. Proving the theorem then requires showing that each time a range is partitioned according to Alg. 3, no configurations are lost. Algorithm 3 iteratively splits a list of ranges using the binary split implemented in Alg. 1. Then, given a range $[C_1, C_2]$ visited by Alg. 3, it

suffices to show, according to program lines 10 and 11, that $[C_1, C_2] = [C_1, C_3] \cup [nextConf(C_3, vs), C_2]$ (where vs is the global `vecSpec` and C_3 is the configuration returned by Alg. 1). We now need to prove that $C_3 \in [C_1, C_2]$ and that $nextConf(C_3, vs)$ indeed returns the next configuration. The first proof is completed by considering the alternatives provided by the guards in the algorithm. For instance, if the set `options` is nonempty, the proof is immediate (see lines 9–15). Regarding the second property, we must prove that Alg. 2 terminates, and that when it terminates it produces the next configuration. Termination is guaranteed because index i iterates from the back of the array until it finds a position in the configuration in which the stored atom name is not the maximum possible. Such a position must exist because the input configuration is required not be the largest possible one. Proving that the configuration produced by Alg. 2 is the next one according to ordering $<_{LC}$ reduces to showing that if a configuration C exists such that $C_3 <_{LC} C$ and $C \leq_{LC} nextConf(C_3, vs)$, then $C = nextConf(C_3, vs)$.

IV. IMPLEMENTATION DETAILS

A. Initial model translation and `vecSpec` construction phase

Given a user-provided Alloy model, as a first step, `Ranger` interfaces with the Alloy Analyzer to obtain a list of suitable fields for range partitioning (currently, all functional binary relations are used; see Section V-E) and to have the model translated to CNF. During the translation, it also interacts with `Kodkod` in order to obtain the necessary information to build the `vecSpec`: a copy of the atom universe, details on which atoms appear in each relevant relation’s domain and range, and on which propositional variable is being used to represent presence or absence of each tuple in each relevant relation.

The model is only translated once. The resulting CNF file is broadcasted by the master to all workers, along with a description of the `vecSpec`, just before the distributed analysis phase starts. All further range-related restrictions are to be injected by the workers, directly at the clausal level, every time they load a new task into their local sequential solver. This eliminates the cost of re-running the Alloy translation toolchain, and allows for subproblems to be very lightweight objects: each pending task is represented by a pair of vectors (which require less than a few hundred bytes each, even for the largest scopes and models in our benchmark).

B. Clauses added to enforce SAT-solving within range

In Alg. 4 we show pseudocode illustrating what each worker does when loading a new subproblem. Three groups of clauses are injected. The first group limits the search to candidate configurations that are no smaller (as per $<_{LC}$) than the left endpoint of the range, whereas the second group requires that they be no larger than the right endpoint. For the third group, both vectors are scanned from left to right until the end of their common prefix (if any) is found. A unit clause is added for every cell that could only have one possible value (i.e., within the common prefix). At the first cell where left and right values differ, an all-positive clause is generated.

The third group does not really add any new constraints: its clauses could also be derived (with some effort) from the first two groups and the rest of the translated model. However, empirical evidence suggests that the presence of this positive formulation often promotes faster propagation.

Let l be the vector length, and c_i (with $0 \leq i < l$) the number of choices for the i -th cell according to a `vecSpec`. Then $\sum_i (c_i - 1)$ is a worst-case upper bound for the number of clauses added by either of the first two groups; as for the third group, it cannot add more than l clauses. In practice, we have not seen any case where the total number of added clauses reached 1% of the number of clauses in the CNF.

```

1  addRangeClauses(ss: SATSolver, R : Range, vs: vecSpec)
   // Forbid anything smaller than left endpoint
2  Config lvec = leftEndpoint(R);
3  antecedent = []; // Empty list
4  for ith, atom in enumerate(lvec) do
5     options = vs.atoms[ith]; // Sorted list
6     position = options.indexOf(atom);
7     forbidden = options[0, position - 1];
8     for f_ith, f_atom in enumerate(forbidden) do
9        f_pvar = vs.pvars[ith][f_ith];
10       ss.addClause(antecedent ++ [-f_pvar]);
11     end
12     atom_pvar = vs.pvars[ith][position];
13     antecedent.append(-atom_pvar);
14 end
   // Forbid anything greater than right endpoint
15 Config rvec = rightEndpoint(R);
16 antecedent = []; // Empty list
17 for ith, atom in enumerate(lvec) do
18    options = vs.atoms[ith]; // Sorted list
19    position = options.indexOf(atom);
20    forbidden = options[position + 1, len(options) - 1];
21    for f_ith, f_atom in enumerate(forbidden) do
22       f_pvar = vs.pvars[ith][f_ith + position + 1];
23       ss.addClause(antecedent ++ [-f_pvar]);
24    end
25    atom_pvar = vs.pvars[ith][position];
26    antecedent.append(-atom_pvar);
27 end
   // Add a unit clause per common prefix cell
   // and a clause for the first differing cell
28 atompairs = zip(lvec, rvec);
29 for ith, (latom, ratom) in enumerate(atompairs) do
30    options = vs.atoms[ith];
31    lpos = options.indexOf(latom);
32    rpos = options.indexOf(ratom);
33    if latom == ratom then
34       // still within common prefix
35       ss.addClause([vs.pvars[ith][lpos]]);
36    else
37       // first cell where values differ
38       ss.addClause([vs.pvars[ith][lpos], rpos]);
39       break;
40    end
41 end

```

Algorithm 4: Adding clauses to enforce ranged analysis.

V. EXPERIMENTAL RESULTS

In this section we first describe the hardware and software setup (V-A). We then evaluate `Ranger` on a benchmark of models that includes valid (V-B) and invalid (V-C) assertions. In V-D we evaluate how the speedup achieved by `Ranger`

evolves as the amount of hardware used for the analysis varies. Finally, in V-E we discuss some possible threats to the validity of the presented experimental results.

A. Setup and conventions

All experiments were run in a cluster of 16 identical PCs, each featuring an Intel Core i7-2600 4-core, 8-thread processor with a 3.40 GHz clock speed and 8 GB DDR3 RAM, running Linux 3.2.0. All `Ranger` evaluation runs were executed on 8x8 (eight physical nodes, each running eight workers) except where otherwise indicated (e.g., the “adding more hardware” experiments, which were run on 4x8, 8x8, 12x8, and 16x8).

All times are given in wallclock seconds. All timeouts (“TO”) mean failure to complete within 36,000 seconds (10 hours), except where otherwise indicated. “OofM” means failure to complete due to exhausting 8 GB of main memory.

To compensate for some variance due to subtask scheduling, each `Ranger` experiment was run three times; the reported timing is the average thereof.

B. UNSAT cases (valid properties)

`BINARYTREES` is the model that we introduced as a running example in Section II. Property `TWODEFSEQUIVALENT` asserts the equivalence of two different characterizations of the binary tree structure. As shown on Table IV, its difficulty curve is particularly steep: although the property can be proven for scope 9 in under a minute, scope 10 requires over 7 hours. `Ranger` can prove the latter in under 4 minutes – a 119x speedup. It can also prove the property for scope 11 in under an hour, whereas the Alloy Analyzer fails to prove it within the 10-hour timeout. Note that in this case the speedup is conservatively reported as being “>10.83”, since we do not know how much longer than 10 hours the Analyzer would need. The actual speedup is likely to be much higher.

`LINKEDLISTS` is a model involving singly linked lists. In this case the goal is to verify that 3 different definitions are equivalent. The model includes two separate properties to that effect: `PAIRWISE`, which asserts that $D_1 \Leftrightarrow D_2 \wedge D_2 \Leftrightarrow D_3$, and `CIRCULAR` (i.e., $D_1 \Rightarrow D_2 \wedge D_2 \Rightarrow D_3 \wedge D_3 \Rightarrow D_1$). Tables V and VI show similar behavior for both properties, with `Ranger` obtaining about 14x speedup on the largest scope that the Alloy Analyzer can handle (within 10 hours), and then being able to prove the properties for 2 additional scopes.

`CHORD` is a model of the Chord [35] distributed hash table lookup protocol. It is one of the case studies bundled with the Alloy distribution. The model contains one property, called `FINDSUCCESSORWORKS`, that is particularly hard to prove. Table VII shows a speedup of at least 103x (again, merely a floor value) for the first scope that is not tractable sequentially. It also shows that distributed analysis pushes the tractability barrier another 2 scopes for this property.

`STABLEMUTEXRING`, another Alloy-bundled example, is a model of Dijkstra’s K-state mutual exclusion algorithm for a ring [9]. There are two hard-to-prove properties in this model. Both use the notion of a “bad tick” – an instant in time where two or more distinct processes try to run their critical

sections simultaneously. NOBADSAFETYTRACE asserts that it is impossible to find a trace with a loop containing a bad tick (such that the algorithm would never stabilize). CLOSURE asserts that there can be no bad ticks if the first tick is “good”. As seen on Tables VIII and IX, Ranger pushes the ten-hour tractability limit 10 scopes (from 12 to 22) for the former, and 4 scopes (from 13 to 17) for the latter. At the last AA-tractable scopes (12 and 13, respectively), the speedups exceed 200x for NOBADSAFETYTRACE and 40x for CLOSURE.

FIREWIRE describes the behavior of the leader election protocol used in the IEEE 1394 [17] standard for connecting consumer electronic devices. This is another case study included with the Alloy distribution. The hardest property in the model is ATMOSTONEELECTED, which asserts that two or more devices cannot be elected as leader in the same state. As shown on Table X, distributed analysis yields nearly 20x speedup for scope 5. For scope 6, where the Alloy Analyzer fails to yield a result within 10 hours, Ranger proves the property in under 4 hours.

C. SAT cases (instance generation, invalid properties)

Many Alloy-borne SAT cases are easy; typically, when the translation of an Alloy property results in a satisfiable problem instance, finding a satisfying valuation is a quick and simple matter. However, hard SAT problems do come up in practice, and can be very challenging indeed. Therefore, we also evaluate Ranger on some difficult SAT instances.

BINOMIALHEAP is the translation to Alloy of a Java binomial heap class implementation, taken from [38]. One of its methods, `extractMin()`, contains a bug that can only be detected for some sufficiently large input structures. Property EXTRACTMINCORRECT asserts the correctness of said method. Its translation to CNF yields UNSAT problems up to scope 12, but nontrivial SAT problems for scopes 13 and above. Although the speedups obtained were modest (around 3x, on average), it was important for us to confirm that Ranger did not miss the counterexample whenever one existed.

AVLTREES is a model of AVL trees which was written for automated test input generation purposes. The goal, for scope n , is to find some configuration that represents a valid AVL tree of size n . An easy task for small n , this becomes much harder as n grows. Table XII shows that it took the Analyzer over 1 hour to produce an AVL tree of size 19, while Ranger achieved the same in 138 seconds – a 27x speedup. At scope 22, sequential analysis was no longer possible because the solver quickly exhausted 8 GB RAM, whereas distributed analysis succeeded in producing AVLs of sizes 22 and 23.

The FIREWIRE model also contains an instance generation command called NOREPEATS. This is an auxiliary property: the model’s author suggests executing it repeatedly while increasing the number of states, until a counterexample is no longer found, as a way to determine how many states are sufficient for a certain scope (number of devices and links). For large scopes, finding these intermediate SAT instances becomes a hard problem in its own right. We ran these analyses sequentially for up to 16 states per scope, and for each scope,

re-ran the most demanding analysis using Ranger. As shown on Table XIII, the distributed approach yielded over 40x speedup for the last sequentially-tractable scope (22), and was able to raise the tractability limit from 22 to 34.

D. Adding more hardware

For each of the aforementioned series, we took the hardest scope that was tractable by Ranger on 8x8 and re-ran it using half as much hardware, 50% more hardware, and twice as much hardware (i.e., on 4x8, 12x8, and the full 16x8 capacity of the cluster). The results are reported on Table XIV. In all UNSAT cases, the actual runtimes were close to the linear extrapolation (200%, 100%, 66%, 50%) of the 8x8 timing.

SAT cases are less predictable since, rather than exhausting the search space, success depends on quickly finding the first needle in the haystack. While AVL instance generation scaled even better than expected, the other two cases did not do as well, and SAT runs on 4x8 performed poorly in general.

| Scope | AA | Ranger | Speedup |
|-------|-----------|----------|---------|
| 8 | 6.00 | 5.49 | 1.09 |
| 9 | 43.69 | 16.58 | 2.63 |
| 10 | 25,552.44 | 215.22 | 118.72 |
| 11 | TO | 3,324.20 | > 10.83 |
| 12 | TO | TO | |

TABLE IV
BINARYTREES: TWODEFSEQUIVALENT

| Scope | AA | Ranger | Speedup |
|-------|-----------|-----------|---------|
| 13 | 24.25 | 14.56 | 1.67 |
| 14 | 86.15 | 37.57 | 2.29 |
| 15 | 346.48 | 91.09 | 3.80 |
| 16 | 1,862.96 | 197.98 | 9.41 |
| 17 | 11,580.27 | 819.81 | 14.13 |
| 18 | TO | 4,107.56 | > 8.76 |
| 19 | TO | 22,845.55 | ≥ 1.58 |
| 20 | TO | TO | |

TABLE V
LINKEDLISTS: THREEDEFSEQUIVALENT (PAIRWISE)

| Scope | AA | Ranger | Speedup |
|-------|-----------|-----------|---------|
| 13 | 22.17 | 14.73 | 1.50 |
| 14 | 74.94 | 35.61 | 2.10 |
| 15 | 360.01 | 85.65 | 4.20 |
| 16 | 1,602.04 | 189.79 | 8.44 |
| 17 | 11,484.27 | 859.62 | 13.36 |
| 18 | TO | 4,299.79 | > 8.37 |
| 19 | TO | 24,077.44 | ≥ 1.50 |
| 20 | TO | TO | |

TABLE VI
LINKEDLISTS: THREEDEFSEQUIVALENT (CIRCULAR)

E. Threats to Validity

When building a vecSpec from an Alloy model, the current implementation of Ranger considers those relations that are binary and functional. Note that this does not restrict input models to those that only use such relations. Ranger can analyze Alloy models with relations of arbitrary type and

| Scope | AA | Ranger | Speedup |
|-------|----------|----------|----------|
| 6 | 94.90 | 23.95 | 3.96 |
| 7 | 1,447.98 | 67.86 | 21.34 |
| 8 | TO | 349.76 | > 102.93 |
| 9 | TO | 3,569.07 | >> 10.09 |
| 10 | TO | TO | |

TABLE VII
CHORD: FINDSUCCESSORWORKS

| Scope | AA | Ranger | Speedup |
|-------|-----------|-----------|-----------------|
| 10 | 322.39 | 35.79 | 9.01 |
| 11 | 1,326.09 | 51.10 | 25.95 |
| 12 | 24,239.91 | 118.04 | 205.35 |
| 13 | TO | 330.74 | > 108.85 |
| 14 | TO | 850.46 | >> 42.33 |
| 15 | TO | 1,672.21 | >>> 21.53 |
| 16 | TO | 3,802.20 | >>>> 9.47 |
| 17 | TO | 5,263.09 | >>>>> 6.84 |
| 18 | TO | 7,400.67 | >>>>>> 4.86 |
| 19 | TO | 10,859.77 | >>>>>>> 3.31 |
| 20 | TO | 16,404.06 | >>>>>>>> 2.19 |
| 21 | TO | 23,982.52 | >>>>>>>>> 1.50 |
| 22 | TO | 29,705.61 | >>>>>>>>>> 1.21 |
| 23 | TO | TO | |

TABLE VIII
STABLEMUTEXRING: NOBADSAFETYTRACE

| Scope | AA | Ranger | Speedup |
|-------|----------|-----------|-----------|
| 10 | 343.50 | 33.03 | 10.40 |
| 11 | 924.96 | 57.08 | 16.20 |
| 12 | 2,835.47 | 111.72 | 25.38 |
| 13 | 9,459.15 | 231.50 | 40.86 |
| 14 | TO | 707.88 | > 50.86 |
| 15 | TO | 2,427.98 | >> 14.83 |
| 16 | TO | 9,771.50 | >>> 3.68 |
| 17 | TO | 30,607.91 | >>>> 1.18 |
| 18 | TO | TO | |

TABLE IX
STABLEMUTEXRING: CLOSURE

| Scope | AA | Ranger | Speedup |
|-------|----------|-----------|---------|
| 3 | 3.98 | 3.16 | 1.26 |
| 4 | 141.04 | 23.67 | 5.96 |
| 5 | 6,269.58 | 319.87 | 19.60 |
| 6 | TO | 14,297.74 | > 2.52 |
| 7 | TO | TO | |

TABLE X
FIREWIRE: ATMOSTONEELECTED

| Scope | AA | Ranger | Speedup |
|-------|----------|--------|---------|
| 8 | 102.10 | 40.56 | 2.52 |
| 9 | 185.05 | 106.64 | 1.74 |
| 10 | 243.12 | 132.81 | 1.83 |
| 11 | 563.47 | 196.93 | 2.86 |
| 12 | 700.69 | 239.04 | 2.93 |
| 13 | 80.41 | 31.20 | 2.58 |
| 14 | 122.87 | 60.84 | 2.02 |
| 15 | 251.49 | 80.29 | 3.13 |
| 16 | 349.60 | 187.09 | 1.87 |
| 17 | 847.28 | 270.93 | 3.13 |
| 18 | 483.16 | 116.55 | 4.15 |
| 19 | 542.40 | 381.70 | 1.42 |
| 20 | 1,022.42 | 149.97 | 6.82 |

TABLE XI
BINOMIALHEAP: EXTRACTMINCORRECT

| Scope | AA | Ranger | Speedup |
|-------|-----------|-----------|----------|
| 15 | 47.09 | 13.87 | 3.40 |
| 16 | 121.76 | 51.12 | 2.38 |
| 17 | 195.37 | 81.05 | 2.41 |
| 18 | 1,703.20 | 125.95 | 13.52 |
| 19 | 3,715.74 | 137.69 | 26.99 |
| 20 | 3,839.97 | 241.02 | 15.93 |
| 21 | 17,588.57 | 1,422.13 | 12.37 |
| 22 | OofM | 4,993.58 | ∞ |
| 23 | OofM | 13,654.49 | ∞ |
| 24 | OofM | TO | |

TABLE XII
AVLTREES: GENERATEINSTANCE

| Scope | AA | Ranger | Speedup |
|-------|--------|--------|----------|
| 16 | 148.97 | 17.01 | 8.76 |
| 18 | 334.40 | 28.48 | 11.74 |
| 20 | 497.93 | 41.12 | 12.11 |
| 22 | 765.93 | 18.21 | 42.06 |
| 24 | OofM | 46.20 | ∞ |
| 26 | OofM | 63.52 | ∞ |
| 28 | OofM | 100.62 | ∞ |
| 30 | OofM | 89.99 | ∞ |
| 32 | OofM | 95.79 | ∞ |
| 34 | OofM | 171.11 | ∞ |
| 36 | OofM | OofM | |

TABLE XIII
FIREWIRE: NOREPEATS

arity; it simply ignores nonfunctional and/or ternary relations for the purposes of building the vecSpec, and therefore, for those of range partitioning. So far, this does not seem to be an impediment. Most of the models in the benchmark use some nonfunctional and/or ternary relations; some even use many of them, and comparatively few functional binary ones (for instance, CHORD uses 5 of the former and just 3 of the latter, while FIREWIRE uses 8 of the former and 4 of the latter). However, this does imply that models using no functional binary relations at all would not be splittable by the current version of Ranger, as the vecSpec would be empty.

A binary non functional relation $R \in A \times B$ can be seen as a ternary relation $T_R \in A \times B \times \{true, false\}$, where $(a, b) \in R \iff (a, b, true) \in T_R$. Therefore, the problem reduces to handling ternary relations. To that effect, we can see a ternary relation $T \in A \times B \times C$, as a total function $F_T : A \rightarrow (B \times C)$, where $F_T(a) = \{(b, c) \in B \times C \mid (a, b, c) \in T\}$.

VI. RELATED WORK

As we said in Section I, scaling the analysis to larger scopes is necessary to improve the confidence levels obtainable by users when writing and analyzing Alloy models. An important step in this direction is the inclusion of symmetry-breaking predicates during the translation to a propositional formula, which significantly enhances the Analyzer's analysis capabilities [29], [36]. Also, since Alloy's translation target are propositional formulas, it leverages the frequent improvements in SAT-solving technology. Surprisingly, developments on parallel and/or distributed analysis of Alloy models are scarce.

One first option to consider is using a parallel SAT-solver. Multi-core SAT-solver research has gained a lot of momen-

| Model/Property | Scope | 4x8 | 8x8 | 12x8 | 16x8 |
|------------------------------|-------|---------|--------|--------|--------|
| LinkedLists: Equiv. Pairwise | 19 | TO | 22,846 | 14,446 | 10,197 |
| | | | 100% | 63% | 45% |
| LinkedLists: Equiv. Circular | 19 | TO | 24,077 | 14,368 | 10,059 |
| | | | 100% | 60% | 42% |
| BinTrees: Equivalence | 11 | 6,584 | 3,324 | 2,273 | 1,688 |
| | | 198% | 100% | 68% | 51% |
| Chord: FindSuccWorks | 9 | 7,041 | 3,569 | 2,270 | 1,807 |
| | | 197% | 100% | 64% | 51% |
| SMRing: Closure | 17 | TO | 30,608 | 18,806 | 12,848 |
| | | | 100% | 61% | 42% |
| SMRing: BadSafetyTrace | 22 | TO | 29,706 | 20,526 | 17,354 |
| | | | 100% | 69% | 58% |
| FireWire: AtMostOneElected | 6 | 27,680 | 14,298 | 8,843 | 6,707 |
| | | 194% | 100% | 62% | 47% |
| BHeap: ExtractMinCorrect | 19 | 1,989 | 382 | 322 | 272 |
| | | 521% | 100% | 84% | 71% |
| AVL: instance generation | 19 | TO | 13,654 | 5,516 | 4,958 |
| | | | 100% | 40% | 36% |
| Firewire: NoRepeats | 34 | 19,949 | 171 | 177 | 195 |
| | | 11,659% | 100% | 104% | 114% |

TABLE XIV
ADDING MORE HARDWARE

tum. ManySAT [16] and plingeling [4] are award-winning multithreaded SAT solvers. As shown on Table XV, Ranger frequently outperforms both of them even when running on a single machine (1x8), possibly due to the synergy between range partitioning and Alloy’s symmetry breaking. But another important advantage of Ranger is its distributed nature, which makes it possible to add more machines and combine their computational power. Multithreaded solvers heavily depend on shared memory and are thus confined to a single computer.

Unfortunately, while multithreaded SAT-solvers are starting to take off, working distributed ones are hard to come by. PMSat [15], an MPI-based, cluster-oriented version of the MiniSat solver, is available for experimentation but reports generally small speed-ups. GrADSAT [7] reported experiments showing an average 3.27x and a maximum 19.9x speed-up using various numbers of workers ranging between 1 and 34. C-sat [26] is a SAT-solver for clusters. It reports linear speed-ups, but the tool is not available for experimentation. Also, relying on a parallel SAT-solver prevents making use of Alloy-level information that may contribute to better analyses.

In [28], the notion of transcompiling is introduced as an aid to improve parallel analysis of Alloy models. Since Alloy analyses occur within given bounds, transcompiling proposes to explore small scopes first in order to extrapolate the best way to distribute the analysis of larger scopes. Ranger may contribute to the development of transcompiling, given that it introduces a new technique for distributing the analysis.

Although little research has been done on parallelizing its analysis, Alloy has been used as an intermediate language by different tools that parallelize code analysis. In [33], parallel analysis of code is performed by splitting the program control flow graph and using JForge [8] (which relies on Kodkod) to analyze each slice. Notice that the parallelization, as is also the case in [32], occurs at the code level, and not at the level of the Alloy intermediate representation. In [27], parallel analysis of Java methods is performed by translating complete methods

| Model/Property | Scope | Ranger 1x8 | plingeling 1x8 | ManySAT 1x8 |
|------------------------------|-------|---------------|-------------------|----------------|
| LinkedLists: Equiv. Pairwise | 15 | 321.70 | 657.90 | 545.53 |
| | 16 | 1,169.76 | TO | TO |
| | 17 | TO | TO | TO |
| BinTrees: Equivalence | 9 | 23.20 | 97.80 | 15.57 |
| | 10 | 1,467.65 | TO | TO |
| | 11 | TO | TO | TO |
| Chord: FindSuccWorks | 6 | 33.87 | 173.80 | 51.34 |
| | 7 | 269.88 | 1,514.10 | 606.86 |
| | 8 | TO | TO | TO |
| SMRing: Closure | 11 | 221.37 | 837.50 | 361.60 |
| | 12 | 639.62 | 1,657.60 | 1,039.99 |
| | 13 | TO | TO | TO |
| SMRing: BadSafetyTrace | 11 | 180.05 | TO | 426.72 |
| | 12 | 695.54 | TO | TO |
| | 13 | TO | TO | TO |
| FireWire: AtMostOneElected | 4 | 51.85 | 14.10 | 35.02 |
| | 5 | TO | 71.80 | 1,529.32 |
| BHeap: ExtractMinCorrect | 14 | 65.15 | 435.50 | 12.89 |
| | 15 | 445.18 | 506.70 | 139.19 |
| | 16 | TO | 501.40 | 65.22 |
| AVL: instance generation | 15 | 20.07 | 8.10 | 15.95 |
| | 16 | 175.42 | 12.90 | 31.65 |
| | 17 | TO | 18.00 | 78.93 |
| Firewire: NoRepeats | 22 | 19.76 | 178.10 | 201.56 |
| | 24 | 741.24 | 293.70 | 199.97 |
| | 26 | TO | 502.50 | 450.80 |

TABLE XV
RANGER ON 1X8 VS. MULTITHREADED SAT-SOLVERS. (TO=30 MIN)

to Alloy. The partitions required to parallelize the analysis are obtained from the Alloy intermediate representation. Unfortunately, the efficiency of the technique depends on the presence of class invariants or the lack of aliasing, concepts usually absent in more general Alloy models such as the ones considered in this article.

The vector-based representation of Alloy configurations is adopted from Korat’s [5] *candidate vectors*. Korat is a tool for test input generation, for which a parallel version (PKorat [30]) exists. Yet neither the problems addressed by PKorat and Ranger (testing and Alloy analysis, respectively), nor the partition techniques, are related.

Ranging techniques for symbolic execution [31] and explicit state model checking [12] of imperative programs were introduced recently in the context of the KLEE symbolic execution tool for C [6] and the JPF model checker for Java [37] respectively. Ranging to analyze declarative programs in Alloy is very different from ranging to analyze imperative programs in C or Java. Specifically, ranges in symbolic execution and model checking are based on program execution paths, specifically sequences of control-flow branches. Such paths do not exist in declarative models. Our technique for ranging for Alloy defines a novel form of ranges – at the black-box input space level, not white-box control-flow level.

VII. CONCLUSIONS AND FURTHER WORK

This paper introduced a novel technique for scaling Alloy’s SAT-based analysis using ranging. The key idea is to distribute the Alloy problem into several subproblems of lesser complexity by defining ranges that partition the space of candidate solutions such that each subproblem only explores

its corresponding range. Experiments using a variety of hard-to-solve Alloy formulas show the efficacy of ranging in scaling Alloy.

Our work opens a new direction in scaling analysis of declarative programs. With the increasing availability of multi-core and multi-processor systems, such parallel techniques have a vital role to play in substantially enhancing our ability to develop more reliable software. We plan to next explore the application of ranging to other declarative domains, such as SMT solving as well as deep static checking where the program and its specification are represented together using a formula, which captures a violation of the specification by the program for goal-directed counterexample generation.

REFERENCES

- [1] Abad P., Aguirre N., Bengolea V., Ciolek D., Frias M.F., Galeotti J., Maibaum T., Moscato M., Rosner N., Vissani L., *Tight Bounds + Incremental SAT = Better Test Generation under Rich Contracts*, in Proceedings of Sixth IEEE International Conference on Software Testing, Verification and Validation (ICST) 2013.
- [2] Abrial J. R., *The B-Book: Assigning Programs to Meanings*. Cambridge, UK, Cambridge University Press, 1996.
- [3] Alloy Analyzer, available at <http://alloy.mit.edu/alloy/download.html>.
- [4] Biere A., *Lingeling, Plingeling, PicoSAT and PrecoSAT at SAT Race 2010*, in Solver description, Special Track 1 (Parallel CNF), SAT-Race 2010, available at http://baldur.iti.uka.de/sat-race-2010/descriptions/solver_1+2+3+6.pdf.
- [5] Boyapati C., Khurshid S., Marinov D., *Korat: automated testing based on Java predicates*. ISSTA 2002: 123-133.
- [6] C. Cadar, D. Dunbar, and D. R. Engler. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *Proc. 8th Symposium on Operating Systems Design and Implementation (OSDI)*, pages 209–224, 2008.
- [7] Chrabakh W., and Wolski R., *GrADSAT: A Parallel SAT Solver for the Grid*, in UCSB Computer Science Technical Report Number 2003-05.
- [8] Dennis, G., Chang, F., Jackson, D., *Modular Verification of Code with SAT*. in ISSTA'06, pp. 109–120, 2006.
- [9] Dijkstra E. W., *A belated proof of self-stabilization*, Distributed Computing, Vol. 1, Issue 1, pp.5–6, 1986.
- [10] Dolby J., Vaziri M., Tip F., *Finding Bugs Efficiently with a SAT Solver*, in ESEC/FSE'07, pp. 195–204, ACM Press, 2007.
- [11] Een N., and Sörensson N., *An Extensible SAT-Solver*. In SAT 2003.
- [12] D. Funes, J. H. Siddiqui, and S. Khurshid. Ranged model checking. In *Proc. Java PathFinder Workshop (JPF)*, 2012.
- [13] Galeotti J. P., Rosner N., López Pombo C., Frias M. F., *Analysis of invariants for efficient bounded verification*. In proceedings of ISSTA 2010, pp. 25-36, 2010.
- [14] Galeotti J. P., Rosner N., López Pombo C., Frias M. F., *TACO: Efficient SAT-Based Bounded Verification Using Symmetry Breaking and Tight Bounds*. IEEE Transactions on Software Engineering, to appear.
- [15] Gil L., Flores P., and Silveira L. M., *PMSat: a parallel version of MiniSAT*, Journal on Satisfiability, Boolean Modeling and Computation 6 (2008) 71-98.
- [16] Hamadi Y., Jabbour S., and Sais L., *ManySAT: a Parallel SAT Solver*, International Journal on Satisfiability, Boolean Modeling and Computation (JSAT), Volume 6, Special Issue on Parallel SAT, IOS Press, 2009.
- [17] *IEEE Standard for a High-Performance Serial Bus*, available at <http://ieeexplore.ieee.org/servlet/opac?punumber=4659231>
- [18] Jackson, D., *Software Abstractions*. MIT Press, 2006.
- [19] Kang E., Jackson D., *Formal Modeling and Analysis of a Flash Filesystem in Alloy*. in Proceedings of ABZ 2008, LNCS 5238, Springer, 294–308.
- [20] Kim J. S., and Garlan D., *Analyzing Architectural Styles*, Journal of Systems and Software, Vol. 83, Issue 7, Elsevier, 1216-1235.
- [21] Khurshid S., and Marinov D., *TestEra: Specification-Based Testing of Java Programs Using SAT*, Automated Software Engineering 11(4): 403–434 (2004)
- [22] Leavens G.T., Baker A.L., and Ruby C. *JML: a notation for detailed design*. In Behavioral Specifications of Businesses and Systems, Chapter 12, pp. 175-188, Amsterdam, Kluwer, 1999.
- [23] Leino K. R. M., Mülcer P., *Using the Spec# Language, Methodology, and Tools to Write Bug-Free Programs*, Manuscript KRML 189, 17 September 2009, Available at <http://specsharp.codeplex.com/wikipage?title=Tutorial>
- [24] Maoz S., Ringert J.O., and Rumpe B., *CD2Alloy: Class Diagrams Analysis Using Alloy Revisited*. In Proceedings of MODELS 2011, LNCS 6981, Springer, 592-607.
- [25] Object Management Group. *OCL Specification V. 2.3.1*. January 1st, 2012. Available at <http://www.omg.org/spec/OCL/2.3.1/PDF/>.
- [26] Ohmura K., and Ueda K., *c-sat: A Parallel SAT Solver for Clusters*, in SAT 2009, LNCS 5585, 2009.
- [27] Rosner N., Galeotti J. P., Bermúdez S., Marucci Blas G., Perez de Rosso S., Pizzagalli L., Zemín L., and Frias M. F., *Parallel Bounded Analysis in Code with Rich Invariants by Refinement of Field Bounds* to appear in Proceedings of ISSTA 2013.

- [28] Rosner N., López Pombo C. G., Aguirre N., Jaoua A., Mili A., and Frias M. F., *Parallel Bounded Verification of Alloy Models by TranScoping*, in Proceedings of VSTTE 2013, to appear.
- [29] Shlyakhter I., *Generating effective symmetry-breaking predicates for search problems*. In Proceedings of LICS 2001 Workshop on Theory and Applications of Satisfiability Testing, June 2001, Boston, MA. Henry Kautz and Bart Selman (eds.), Electronic Notes in Discrete Mathematics, Vol. 9, 2001.
- [30] Siddiqui J. H., and Khurshid S., *PKorat: Parallel generation of structurally complex test inputs*. 2nd International Conference on Software Testing, Verification, and Validation (ICST 2009). Denver, CO. Apr 2009.
- [31] J. H. Siddiqui and S. Khurshid. Scaling symbolic execution using ranged analysis. In *Proc. 27th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, 2012.
- [32] Shao D., Gopinath D., Khurshid S., Perry D. E., *Optimizing Incremental Scope-Bounded Checking with Data-Flow Analysis*. ISSRE 2010: 408–417.
- [33] Shao D., Khurshid S., Perry D. E., *An Incremental Approach to Scope-Bounded Checking Using a Lightweight Formal Method*. FM 2009: 757–772.
- [34] Spivey J. M., *The Z Notation: A Reference Manual*, 2nd ed. Upper Saddle River, NJ, Prentice Hall, 1992.
- [35] Stoica I., and Morris R., and Liben-Nowell D., Karge D., and Kaashoek M. F., and Balakrishnan H., *Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications*, IEEE Transactions on Networking, vol. 11, 2003.
- [36] Torlak E., Jackson, D., *Kodkod: A Relational Model Finder*. in TACAS '07, LNCS 4425, pp. 632–647.
- [37] W. Visser, K. Havelund, G. Brat, and S. Park. Model checking programs. In *Proc. 15th Conference on Automated Software Engineering (ASE)*, Grenoble, France, 2000.
- [38] Visser W., Păsăreanu C. S., Pelánek R., *Test Input Generation for Java Containers using State Matching*, in ISSTA 2006, pp. 37–48, 2006.
- [39] Zave, P., *Compositional binding in network domains*. In Proceedings of FM 2006. LNCS 4085, Springer, 332-347.