

Finding Kernels in Non-Linear Data-Driven CHC Solving

Michael Eden

Trustable Programming
Georgia Institute of Technology

Abstract

Program verification has seen a lot of progress, but its still unable to automatically find proofs for industry programs. This paper builds on data-driven approaches from previous work [11] to provide a more robust automatic prover for programs with non-linear loop invariants. It does so by attempting to find the correct kernel for the relation that makes the invariant linear. This is an easy addition to existing systems and can be used with any data-driven approach, allowing it to be easily implemented on top of them. By finding a suitable kernel, many difficult non-linear invariants are easily found.

1 Background

Questions on computability and decidability have perplexed theoretical computer scientists for many years. Early research showed that some classes of problems are undecidable, i.e. no algorithm could determine an answer in a finite amount of time [9]. However there are incredibly useful problems that fall into this class, such as determining if an algorithm will ever end or determining if it contains errors. Since malicious actors exploit software by finding errors in its code, the process of automatically finding those errors becomes invaluable to people, businesses, and the government. There are some who believe that, since the problem is impossible, its not worth pursuing, but only a few years after being ruled out a solution was discovered and published [9]. These two discoveries weren't conflicting, rather the solution simply didnt work in all cases.

Even though it was incomplete, that first solution began a mission to find an ever more (but not quite) complete method. The focus of safety verification has been to attempt to determine if some programs could reach any error states and not to develop a sys-

tem that works for all programs [4]. When rephrased like this the search seems more optimistic. After all, humans can do this task in a finite amount of time for many programs. Many systems have been proposed, but currently encoding programs into a system of Horn Clauses (a special form of logical equation) stands out [8]. If such a system has no solution, then there is no way for a program to reach an error state. This method has had significant success but is incomplete in an important way. One of its vital tasks is being able to summarize a loop body so it can avoid essentially simulating a run of the input program. This is an impossible task, one basically has to find a function that maps every iteration of a loop to a program state efficiently. A common problem is that this function becomes too specific, i.e. it becomes a table of previous iterations and their states. Such a table is not amenable to generalization or abstraction; it is too specific to past iterations to help understand future ones. Specifically, the problem described here is the issue of finding weaker interpolants when finding loop invariants.

This research focuses on finding more general (weaker) loop invariants in the realm of boolean and integer arithmetic. It reviews existing techniques for finding linear invariants, especially data-driven methods, and explores how they can be extended to fit non-linear loops. To accomplish this, a mapping from the state space to a higher-dimensional and non-linear space is found (a kernel). Since the focus is augmenting existing linear solvers, the method is limited by the SMT solver used (Z3 here). Hopefully these weaker non-linear interpolants will generalize better than their linear counterparts.

2 Related Work

In 1936 Alan Turing proved that no algorithm could determine in a finite amount of time if a program will

always terminate or continue to run forever [9]. Unfortunately this also means that many related and useful questions are also undecidable about general programs. For instance one cannot always know whether a program will have errors in it until its actually run, an obvious property of interest in industry. Still thirteen years later Turing also published a now widely used method of proving program termination [9]. In the face of a provably impossible problem, Turings method simply never halts when it cannot decide on an answer.

This led to the comparatively easier problem of answering “Will Halt”, “Will Not Halt” and “Unknown” for every program. Since then the fields of program termination, software verification, program equivalence, etc. have all been attempting to slim down the cases in which their algorithms result in “Unknown” [4] [10] [3]. A driving and practical motivation for program verification is that software embedded into the United States’ electric grid, medical instruments, nuclear weapons storage, and other vital components of infrastructure could contain errors. Proving that these systems work correctly is both beneficial at face value and prevent attackers from compromising them.

Verification has thankfully seen a lot of progress since it was able to transform programs into a concise logical representation known as Horn Clauses and then solve the resulting systems efficiently [8]. As reported by Albarghouthi, there is now a fission in the approach used to solve these systems [1]. One style attempts to find properties at every point in the program while the other only tries to find properties when they are needed to prove some goal [1]. It might seem that the former is clearly worse, however it can solve some programs which the latter cannot and vice-versa. The disagreement here, at its core, is about finding the right property to prove. A lot of research effort has been poured into this problem, and so far the consensus is leaning towards the latter “lazy” approach [1] [8] [7].

A quick consequence of the “lazy” approach is the ability to completely automate the proving process. A key in the fully automated verification process is the ability to find some invariant that is consistent across every iteration of a loop. Since programs without loops (and without recursion) are relatively simple, it is this automatic finding of loop invariants that has taken hold of a lot of research today [8] [10]. Issues with finding these invariants are consistent with the overall theme and difficulty of finding the right property to prove. In McMillan’s research, he demonstrates an automatic proof that a loop starting at zero and incrementing a counter by two can

never be forty-one [8].

```

function COUNT( $N$ )
   $i \leftarrow 0$ 
  while  $i < N$  do
     $i \leftarrow i + 2$ 
  end while
  assert  $i \neq 41$ 
end function

```

Figure 1: Program described in [8]

However the tool in [8] does not prove “the loop starts at an even number and adds an even number, so it cannot be odd” as many people might hope. Instead it shows “once the loop gets to 40, the only next possible value is 42” [8]. In theory, McMillan’s method can find any invariant expressed in linear arithmetic (and the modulo operator is linear for a constant modulus, since $q = p \text{ mod } c \rightarrow p = k \cdot c + q$). In practice, Z3 does not generalize the multiple unfoldings of this loop to arrive at that invariant. It tends to generate facts that are too specific, large, and complex about the current program state and does not provide a more general and useful fact [2] [5]. According to Albarghouthi, the principle of Occams Razor should somehow be embedded into the invariant finding process in the hopes a simpler fact will be a more useful one. This is akin to “regularization” in many statistical methods.

Albarghouthi has made significant progress in finding incredibly concise and general interpolants (a precursor to loop invariants) in Quantifier Free Linear Rational Arithmetic (QFLRA). This encompasses a significant but not total portion of programs in industry [2]. Finding these loop invariants are vital since they are the biggest roadblock to automatic verification. Automatic is important here since most use cases involve large complex software that would take a long time to annotate with invariants.

Another gap in the current research is the availability of a dataset of real-world programs. Many of the seminal papers here are from Microsoft Research, which tests their methods on Windows Drivers [3]. This research benefits from the existing set of testable programs but, being a data-driven approach, will always benefit more from more data. Hopefully there will be a large de-facto dataset in the future, since this is a major hurdle when quantifying new methods and improving old ones.

3 Implementation

The implementation of this method is based off the idea that [11] does a great job with linear data and that it can be extended to fit non-linear data by simply applying a different kernel. The kernel can be “guessed” by analyzing the source code of the incoming program. In the method described in [11] there was no direct analysis of the program source, so the process would be more difficult. It’s also possible to learn the correct space purely through the data, but this requires a large dataset. Solving logical systems, therefore, has a strange relationship with statistical methods since they contain very little data and demand a perfect fit (completely “over-fitted”). As described earlier a “simpler” solution is preferable so that it can generalize, but this generalization must not cause any loss in accuracy.

3.1 Linear CHCs

The function described in Figure 2 is a simple quadratic relationship between x and y . One might think that this non-linearity would pose a problem, but that’s not the case.

```

function QUADRATIC
   $x \leftarrow 1, y \leftarrow 1$ 
  loop
     $y \leftarrow y + x$ 
     $x \leftarrow x + 2$ 
    assert  $y \geq x$ 
  end loop
end function

```

Figure 2: Simple quadratic program

The iteration is simple, positive (brown red) and negative (blue) points are loaded into a classifier to be separated (Figure 3). Positive points are simply variables values that the program will encounter (found by unfolding the loop) and that agree with the goal (if they do not, one can show the goal is wrong). While negative points are counterexamples that Z3 generates that ensure the invariant found isn’t too weak or trivial. A classifier is built and its function is assigned to the loop invariant in Z3. If the classifier finds an invariant strong enough to prove the goal, Z3 stops and the process completes. If the invariant is too weak, Z3 finds a counterexample which then gets added to the pool of negative points and the process restarts. To bootstrap the first negative point, `true` or the goal function is used as the initial invariant. So far this tool has

only emulated [11] but it will soon be extended to do more. Source code is freely available online [6].

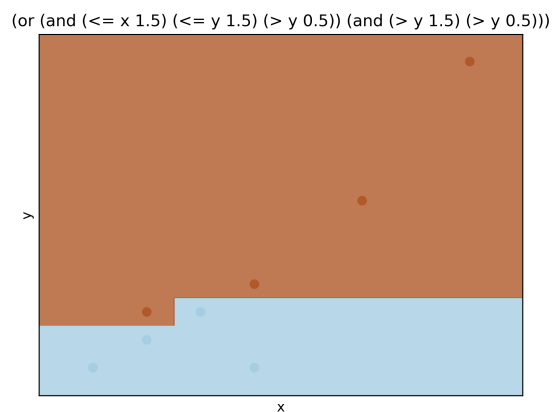


Figure 3: Decision boundary created by Figure 2

Clearly Figure 3 demonstrates that linear decision boundary is quickly found for the quadratic program (Figure 2) using the method described in [11], the classifier found three linear boundaries tied together by a decision tree.

Figure 3 shows the output of the tool developed for this project on the program in Figure 2. Its title is the decision boundary encoded into Z3’s s-expressions, a necessary step in the process, it is the equivalent of:

$$(x \leq 1.5 \wedge 0.5 < y \leq 1.5) \vee (y > 1.5 > 0.5)$$

Naturally, this depends on Z3’s ability to produce counterexamples, which means Z3’s limited ability to handle transcendental functions (e^x , $\sin x$, $x!$, etc.) also applies here, limiting the scope to polynomials.

3.2 Non-Linear CHCs

Non-Linear CHCs have direct application in finding the complexity bound of a program (since many useful programs have polynomial complexity). Writing a simple recurrence containing a multiplication (Figure 4), however, grinds the previous method to a halt.

Solving the recurrence $g(n) = g(n) + n^2$ and $g'(n) = g'(n) + n^2 - 2n$ (where x is $g(n)$, y is $g'(n)$ and z is n) shows us that both x and y are functions of degree three. However it might not always be easy to solve recurrences directly, instead it is easier to approximate the degree and attempt to classify the

```

function CUBIC
   $x \leftarrow 1, y \leftarrow 1, z \leftarrow 1$ 
  loop
     $x \leftarrow x + z^2$ 
     $y \leftarrow y + z^2 - 2z$ 
     $z \leftarrow z + 1$ 
    assert  $G(x, y, z)$ 
  end loop
end function

```

Figure 4: Slightly more complex program

points in a higher space. This is where picking the correct space becomes crucial, too large and sparse and it will take too long to traverse, too small and it won't be expressive enough.

The degree isn't solely decided by the program either, the goal $G(x, y, z)$ is also relevant. For instance if Figure 4 contained

$$G(x, y, z) = x \geq y$$

then a simple linear boundary would suffice. However if the goal was more complex like

$$G(x, y, z) = y + z^2 \geq x$$

then linear boundaries won't do, and the method in [11] iterates forever generating lines that approximate a parabola. It seems that x and y are of close enough degree to each other and z is far away from both. By analyzing the source we can tell the maximum degree these variables might become (but Z3 does not check this bound):

$$z : 1, \quad y : 1 + 2(z) = 3, \quad x : 1 + 2(z) = 3,$$

Further we can look at all the possible ways each term is used (1), or simply take each variable to its maximum degree (2), or look the goals terms (3). This gives us a collection of possible kernels:

$$T = \{x, y, z, z^2\} \tag{1}$$

$$P = \{x^1, x^2, x^3, y^1, y^2, y^3, z^1\} \tag{2}$$

$$Q = \{x, y, z^2\} \tag{3}$$

For the quadratic G itself is enough to act as the invariant, meaning the space spanned by Q is the optimal space to search through. The general solution to the recurrence, however, would not be covered by Q :

$$x = y + z^2 - z \quad \wedge \quad z > 0$$

Since this is the ground truth for the loop invariant, we'd like all classifiers to look through no space larger than this.

4 Experiments

A tool was made to test different heuristics when evaluating these non-linear programs [6]. It analyzes the source and pulls out terms such as T , P , and Q , adding them to the dataset and trying to classify them. A large dataset of positive and negative points was also used from the indefinitely-iterating purely linear method. When these points were classified with added dimensions they were clearly linearly separable. However simply using a polynomial SVM with a decision tree took hours to complete on a test machine.

The most successful heuristic ended up being Q , or terms only in the goal, since the goal has to be somehow expressed within the invariant. Using a simple linear SVM was much faster than the polynomial SVM (a few milliseconds for 1000 data points) so the best strategy so far is to attempt to fit many different combinations of degrees in parallel. Since the number of data points is usually small, given a max degree of Δ and number of variables n that would generate:

$$\text{number of classifiers} = 2^{n\Delta}$$

The example above would train $2^9 = 512$ linear SVMs, which is a lot, but since `scikit-learn`'s SVM is so fast it can generally be done rather quickly. Further if the list of term combinations is ordered with more likely sets first (such as Q), it will likely find a separable line quicker.

5 Conclusion

Overall finding non-linear invariants is harder, but not so much harder than linear invariants. On the other hand even simple exponential functions will cause Z3 to give "unknown". This method is a simple addition to many data-driven methods that can make a lot of programs that users expect to be solvable actually solvable. What this and many other methods need is a larger and more standard data set to learn from. Future work to aid this automatic verification would definitely include manually collecting large data-sets from important open source projects. Another important addition to this work would be to add it to a layer in SeaHorn so that it could be used in production with any LLVM language.

References

- [1] ALBARGHOUTHI, A., GURFINKEL, A., AND CHECHIK, M. Under-approximations to over-approximations and back. *TACAS* (2012).
- [2] ALBARGHOUTHI, A., AND MCMILLAN, K. L. Beautiful interpolants. *CAV* (2013).
- [3] COOK, B., GOTSMAN, A., PODELSKI, A., RYBALCHENKO, A., AND VARDI, M. Y. Proving that programs eventually do something good. *POPL* (2007).
- [4] COOK, B., PODELSKI, A., AND RYBALCHENKO, A. Proving program termination. *ACM* 54, 5 (2011), 88–98.
- [5] D’SILVA, V., KROENING, D., PURANDARE, M., AND WEISSENBACHER, G. Interpolant strength. *VMCAI* (2010).
- [6] EDEN, M. Non-linear chcs. <https://github.com/illegalprime/nonlinear-chcs>, 2018.
- [7] MCMILLAN, K. L. Lazy abstraction with interpolants. *CAV* (2006).
- [8] MCMILLAN, K. L., AND RYBALCHENKO, A. Solving constrained horn clauses using interpolation. *Technical Report MSR-TR-2013-6* (2013).
- [9] TURING, A. Checking a large routine. *Conference on High Speed automatic Calculating Machines* (1949).
- [10] ZHOU, Q., HEATH, D., AND HARRIS, W. Completely automated equivalence proofs. *ACM* (2017).
- [11] ZHU, H., MAGILL, S., AND JAGANNATHAN, S. Data-driven chc solver. *PLDI* (2018).