

CatSAT: A Practical, Embedded, SAT Language for Runtime PCG

Ian Douglas Horswill

Northwestern University, Evanston IL, USA
ian@northwestern.edu

Abstract

Answer-set programming (ASP), a family of SAT-based logic programming systems, is attractive for procedural content generation. Unfortunately, current solvers present significant barriers to runtime use in games. In this paper, I discuss some of the issues involved, and present *CatSAT*, a solver designed to better fit the run-time resource constraints of modern games. Although intended only for small problems, it allows designers to compactly specify simple PCG problems such as NPC generation, solve them in a few tens of microseconds, and to adapt solutions dynamically based on the changing needs of gameplay. We hope that by making adoption as convenient as possible, we can increase the uptake of declarative techniques among developers.

Introduction

Many procedural content generation (PCG) programs amount to making a set of random choices subject to domain constraints. Constraint programming (Rossi, Van Beek, & Walsh, 2006) is an attractive approach for such systems because it allows designers to specify the choices and constraints without having to develop a bespoke search algorithm for solving them (G. Smith, Whitehead, & Mateas, 2011).

Boolean Satisfiability (SAT) has been extensively studied as a constraint programming framework, since it is highly expressive and supports surprisingly fast solvers (Biere, Heule, Maaren, & Walsh, 2009). Answer-set programming (ASP) is a particularly convenient way to formulate SAT problems for PCG (A. M. Smith, 2017; A. M. Smith, Andersen, & Mateas, 2012; A. M. Smith, Nelson, & Mateas, 2010; A. M. Smith & Mateas, 2011). It allows programmers to specify finite-domain constraint satisfaction problems as a set of Prolog-like first-order rules. The ASP system expands them into an equisatisfiable SAT problem, a process known as *grounding*, and solves the resulting problem, generally using some variant of Contradiction-Driven Clause Learning (Marques-Silva & Sakallah, 1999), a backtracking-based systematic search algorithm.

As a simple example, suppose we want to generate a party of 3 non-player characters. Characters have:

- Three possible races (human, electroid, insectoid)
- Four possible classes (fighter, magic user, cleric, thief).
- Humans additionally have one of 3 possible nationalities
- Clerics have one of 4 possible religions.

In addition, there are constraints on the possible solutions:

- Party members should have different classes.
- Electroids can't be clerics.
- One of the religions is outlawed in one of the nations
- Another religion is mandatory in one of the other nations.

This can be written as a 17-line ASP program. Clingo (Eiter, Faber, Fink, & Woltran, 2008), the most commonly used ASP solver, can generate a party in 6ms on a modern laptop.

This is very appealing. It makes it easy to phrase PCG problems and solve them efficiently. Designers are free to incrementally add options and constraints as they see fit, without having to redesign the generator algorithm each time they make a change.

Moreover, it's easy to tailor generation on the fly by adding and removing constraints based on immediate gameplay needs. Provided the constraints aren't inconsistent, the system will simply "solve around" whatever those needs are.

Barriers to in-game execution

Unfortunately, using ASP for in-game PCG faces several challenges.

Designer transparency

One of the key factors in the success of behavior trees was the existence of designer-facing tools that allowed non-programmers to understand and manipulate them (Isla, 2005). Equivalent tools for constraint-based PCG, such as (G.

$m = \text{CatSAT.Solve}()$
 $c =$ all inequalities marked as true in m
Use rejection sampling to solve the inequalities c

The rejection sampler here is very simple (<230 lines of C#), but is sufficient to add the generation of stats to our NPC generator, including different class-dependent constraints, such as requiring fighters to have higher strength than intelligence, or magic users to have higher intelligence than strength. Results are shown in Table 2 under “NPC generator with stats”

Probability patching

One issue with any random generator is that some kinds of configurations have more solutions than others, and so are chosen more frequently. As discussed above, 98% of possible parties in our example involve a human and/or a cleric. If this doesn’t bother the designer or player, then it’s not a problem. If it is problematic, various techniques can be used to adjust the sample distribution.

The simplest is to decide a race and class in advance using a random number generator, then force their values in the `Problem` object. This gives the designer direct control over the distribution of those specific variables, and allows the SAT solver to run faster. This is another example of hybrid solving.

The probability of particular combinations can be increased by giving those combinations extra, hidden attributes to choose values for. This increases the number of notional solutions for those combinations, and thereby their frequency of occurrence. The additional attributes can then be ignored.

Conversely, combinations that are judged to occur too frequently can be controlled using rejection sampling. To reduce the frequency of insectoid fighters by 50%, check the generated character to see if they’re an insectoid fighter. If so, regenerate it with a probability of 50%.

Future work

There are many obvious additions that would make the system more useful. The most obvious of these would be to integrate SMT support in the solver. It would also be useful for the system to ship with a standardized implementation of a floating-point solver, perhaps based on *Craft* (Horswill, 2015). Another useful and straightforward extension would be to modify the solver to support MAXSAT (optimization). Another possible improvement would be to add optional incremental generation of loop formulae, allowing the use of non-tight programs. However, it’s unclear how well this would work with stochastic local search.

There are many performance improvements that are possible, as the current system is not especially well optimized. The use of watched literals (Moskewicz, Madigan, Zhao, Zhang, & Malik, 2001) instead of counters, for example, may improve performance.

Finally, although it may be easier to write a Sudoku generator in *CatSAT* than in raw C#, it still requires considerable comfort with both C# programming and logical axiomatization. A designer-facing tool, a la *Tracery* (Compton, Filstrup, & Mateas, 2014), that would help designers build generators will be important in the future.

Conclusion

SAT-based systems provide a flexible and highly expressive framework for finite-domain PCG problems, but have traditionally been difficult to implement in a running game. *CatSAT* shows that modest-sized problems can be efficiently expressed, solved, and integrated with a performance profile that’s acceptable for a wide range of games. Generators can be easily created with a few lines of code, and solved in tens of microseconds. No search algorithms need be written or debugged. Generators can be incrementally tuned, both in terms of their constraints, and their generation frequencies. Moreover, their behavior can be dynamically steered at run-time to suit gameplay needs.

In addition to making it easier to add PCG to a game, the steerability of the system allows it to be put to unusual uses. It could, for example, be used for dynamic difficulty adjustment (DDA) to adjust the level of challenge of enemies or puzzles. It can be used for bespoke boss and item generation to fit the narrative needs of the game, based on the player’s current stats and inventory. And it can be used as an aid in configuration interfaces (e.g. character or ship designers) to allow the player to specify some attributes, while allowing the system to generate reasonable default values for the remaining attributes. If the player doesn’t like a chosen value, they can change it or just ask the system to choose a different random configuration.

Most excitingly, easy constraint satisfaction could allow the creation of fundamentally new game mechanics, such as the narrative puzzles of Benmergui’s *Storyteller*. The best way to discover such mechanics is to get deployable versions of these new technologies into the hands of practicing game designers.

Acknowledgements

I would like to thank Ethan Robison, Adam Smith, Rob Zubek, Robby Findler, Spencer Florence, and the reviewers for their helpful references, advice, and comments.

