The Theory Blogs Column

BY

LUCA TREVISAN

Bocconi University
Via Sarfatti 25, 20136 Milano, Italy
L.Trevisan@UniBocconi.it
https://lucatrevisan.github.io

A big difference between academic blogs and other forms of academic communication is the immediacy: the ability to post about results that have just been announced, about events that just took place, or about thoughts that just happened, and to start an interactive discussions. In a complementary way, this column can be an opportunity to look back.

For the first few editions of this column, I will ask writers of some of the most popular theoretical computer science blogs to write guest columns in which they look back at how they started and at some of their favorite posts and interactions.

To get started, Lance Fortnow, the OG of theoretical computer science blogging, will tell us what inspired him to get started, how he set the tone for his blog, what's up with the green background, and he will give us an updated version of one of his favorite posts.

Lance's blog, which is now co-written with Bill Gasarch, is at https://blog.computationalcomplexity.org/

Computational Complexity

Lance Fortnow
Illinois Institute of Technology

Thanks to Luca Trevisan for starting this column about theoretical computer science blogs and for asking me to write one of the first columns. Luca said he was interested in knowing what inspired me to start my blog, how I thought about topics (that is, the balance between technical and non-technical content and at what level to address the technical content), why I chose the green background and so on. He also suggested I update a favorite post.

This column thus expands and updates two June 2009 blog posts, The Story of the Blog [3] and A Kolmogorov Complexity Proof of the Lovász Local Lemma [2], the latter about the most mind-blowing STOC talk I ever attended.

1 The Story of the Blog

In 2002 I read a Newsweek article [7] on the then new phenomenon on blogging and I decided to give it a try. I wanted to blog on a specific topic and so I started blogging on the topic I know best. On August 22, 2002 I started "My Computational Complexity Web Log", the first major blog devoted to theoretical computer science. I used the Blogger platform before it was bought out by Google. Back then it had few choices for themes and I chose one with a green background, a color that has become our trademark. I can tell across a large conference room who is looking at the blog.

The blog gained in readership after a sardonic post [4] entitled "Finding the order of the multiplicative group (mod n)... zuh?" in Jason Kottke's popular blog.

> Weblogs are usually pretty easy for readers to get into. Lance Fortnow's Computational Complexity Web Log is probably the most difficult-to-read weblog I've ever come across. But that's OK because if you're into computational complexity, it's just the thing.

All of a sudden I had thousands of readers interested in computational complexity. In that first year I wrote mostly technical posts. I did a set of posts entitled Foundations of Complexity giving an introduction to the field and a series Complexity Class of the Week where I would review results and questions about a specific complexity class. Later on I wrote a monthly series of my Favorite Theorems in the field.

The technical posts take considerable effort to write and proofread, and seemed less interesting to the majority of my readers who didn't come from the theoretical computer science community. I started writing more opinion and academic-oriented posts. The blog became a meeting place for the theoretical computer science community where we would have some discussions, sometimes quite heated, over the issues of concern to our community. Most of the young people I would meet at conferences knew me more for the blog than for my research. For a while a Google search on my name led to the blog before it led to me.

Changes happen. I shortened the name of the blog. In March of 2007 after a post on turtles [1], I felt I was just going through the motions and decided to retire from blogging. Bill Gasarch took over the blog and kept it going. But I had too much I wanted to say and rejoined the blog in January 2008. Since then Bill and I have co-written this blog and have together experimented with podcasting, vidcasting and typecasting, where we simply transcribe a discussion between us. We often hosted guest bloggers, some of whom, like Scott Aaronson, would go on to write popular blogs on their own. There was a time this blog had new posts every working day but as I took on more administrative roles, we typically do about one to two posts a week. I have supplemented the blog with a Twitter account for short comments and announcements.

Now we have many excellent blogs in theoretical computer science ranging from very technical to very amusing. We strive to be the blog of record, the weblog people turn to to learn about the issues and happenings in the community. We try to cover the major results in complexity, remember those we've tragically lost and the centenaries of the founders of our field.

Weblogs have become the places that brings our ever growing academic community together in a way our conferences no longer can. As this blog approaches 20 years, 3000 posts, 25000 comments and 10 million page views, I'm glad Computational Complexity is able to play its part in that effort.

2 A Kolmogorov Complexity Proof of the Lovász Local Lemma

When I was a graduate student at MIT in the late 1980s, Joel Spencer traveled up from New York to teach a course on the Probabilistic Method, one of the most useful courses I had for future research. The probabilistic method is a general approach to show that certain combinatorial objects, such as Ramsey Graphs, exists by showing that they occur with some positive probability over some distribution on graphs. In most cases you can find such graphs by just trying random examples.

In that class I first learned about the Lovász Local Lemma, that if you had a set of binary random variables with high enough probability and enough independence, then there was a positive probability they all were true. However the proof did not lead to a randomized constructive algorithm.

In 90's I attended a talk on an algorithmic proof of the Lovász Local Lemma. I had difficulty following the talk and couldn't even understand how the parameters related to the lemma. I tried asking the speaker afterwards but was dismissed with a "You don't understand."

As a best paper and best student paper winner at STOC 2009, I felt I should attend Robin Moser's talk on his paper "A Constructive Proof of the Lovász Local Lemma" [5] but I didn't have high hopes. How wrong I was. Not only did Moser present an amazing result but also an incredibly inventive simple proof that he came up with while preparing the talk.

Smartly Moser focused on an application of the Lovász Local Lemma to satisfiability instead of the full lemma during the talk. He had a beautiful short can't-believe-that-works construction with a simple information-theoretic argument, which I converted in my head to Kolmogorov complexity, because I think better computationally. As I watched enthralled, I said to Eric Allender sitting next to me "Are we really seeing a Kolmogorov proof the the Lovász Local Lemma?" Indeed we were.

After the talk I quickly wrote up the Kolmogorov proof which I posted early the next morning.

Theorem 1. *Suppose we have a k-CNF formula ϕ with n variables and m clauses and each clause shares a variable with at most r other clauses. Then there is a constant d such that if $r < 2^{k-d-1}$ then ϕ is satisfiable. Moreover we can find that assignment in time polynomial in m and n.*

The full algorithm consists of two short routines Solve and Fix below.

Algorithm 1 Solve(ϕ)

Pick a random assignment of ϕ
while There is an unsatisfiable clause C **do**
  Fix(C)
end while

Assume Fix(C) always terminates. Every clause that was satisfied before we called Fix(C) will still remain satisfied and C will also now be satisfied. So Solve makes at most m calls to Fix.

Algorithm 2 Fix(C)

Replace the variables of C with new random values
while While there is an unsatisfied clause D that shares a variable with C **do**
  Fix(D)
end while

We need to show all the Fix(C) terminate. Suppose the algorithm makes s Fix calls including all the recursive ones. We will show s is bounded and thus the algorithm terminates.

Fix a Kolmogorov random string x of length $n + sk$ (random relative to ϕ, k, s, r, m and n) and assume the algorithm uses the first n bits as the initial assignment and k bits each to replace the variables in each Fix call.

If we know which clause is being fixed, we know the clause is violated so we know all the bits of this clause and thus we learn k bits of x. We then replace those bits with another part of x.

So we can describe x by the list of clauses we fix plus the remaining n bits of the final assignment. We can describe the C such that Fix(C) is called by Solve by $m \log m$ bits and the remaining fixed clauses by $\log r + O(1)$ bits because either it is one of r clauses that intersects the previous clause or we indicate the end of a recursive call (keeping track of the recursion stack).

Since x was random we have

$$m \log m + s(\log r + O(1)) + n \geq n + sk$$

and thus

$$s(k - \log r - d) \leq m \log m$$

for some constant d.

By assumption we have $r \leq 2^{k-d-1}$ or equivalently $k - \log r - d \geq 1$ and thus s must be bounded by $m \log m$.

We choose the x randomly which with high probability will be Kolmogorovly random and the algorithm will run in polynomial time. QED

In follow-up work with Gábor Tardos [6], Moser builds on these techniques to give a constructive version of the full Lovász Local Lemma with the original parameters which yields $r < 2^k/e$ for satisfiability. The Moser-Tardos paper received the Gödel Prize in 2020.

References

[1] Lance Fortnow. Turles. *Computational Complexity Weblog*, March 2007. https://blog.computationalcomplexity.org/2007/03/turtles.html.

[2] Lance Fortnow. A Kolmogorov complexity proof of the Lovász Local Lemma. *Computational Complexity Weblog*, June 2009. https://blog.computationalcomplexity.org/2009/06/kolmogorov-complexity-proof-of-lov.html.

[3] Lance Fortnow. The story of the blog. *Computational Complexity Weblog*, June 2009. https://blog.computationalcomplexity.org/2009/06/story-of-blog.html.

[4] Jason Kottke. Finding the order of the multiplicative group (mod n)...zuh? *kottke.org*, September 2002. https://kottke.org/02/09/order-multiplicative-group.

[5] Robin A. Moser. A constructive proof of the Lovász Local Lemma. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 343–350, New York, NY, USA, 2009. Association for Computing Machinery.

[6] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2), February 2010.

[7] Newsweek Staff. A blog about writing about blogs. *Newsweek*, August 2002. https://www.newsweek.com/blog-about-writing-about-blogs-144269.