TCS on the Web

BY

Stefan Neumann

TU Wien
Erzherzog-Johann-Platz 1
1040 Vienna, Austria
stefan.neumann@tuwien.ac.at
https://neumannstefan.com

Decentralized Thoughts[1] is a group blog on distributed computing, blockchains, and cryptography. Ittai Abraham started it in May 2019, and it has since grown into a broad collection of course notes, research-oriented posts, and overview articles. The site also contains a "Start Here" page[2] with material for a graduate course on Blockchains, Distributed Computing, and Cryptography, a course page[3], and a video collection[4].

In this conversation, we do not only talk about the blog itself. We also discuss how a good blog can convey intuition and collect scattered knowledge, something that is much harder to do in a paper. Decentralized Thoughts is especially interesting because it sits between several communities—distributed computing, cryptography, systems, and blockchains—and also between theory and practice.

Today I am very happy to talk to four of the authors of Decentralized Thoughts: Ittai Abraham[5], Kartik Nayak[6], Ling Ren[7], and Alin Tomescu[8].

[1] https://decentralizedthoughts.github.io
[2] https://decentralizedthoughts.github.io/start-here/
[3] https://decentralizedthoughts.github.io/course/
[4] https://decentralizedthoughts.github.io/videos/
[5] https://a16zcrypto.com/team/ittai-abraham/
[6] https://users.cs.duke.edu/~kartik/
[7] https://sites.google.com/view/renling
[8] https://alinush.github.io

Decentralized Thoughts: Blogging Consensus, Blockchains, and Cryptography

A Conversation with Ittai Abraham, Kartik Nayak, Ling Ren, and Alin Tomescu

Ittai started Decentralized Thoughts in May 2019 on his own. After that, the blog soon started having multi-author posts on diverse topics. How has the blog changed over time and what are your motivations for blogging?

Ittai: I had been thinking about starting a blog for a long time, but I did not really have the courage to do it. This goes back to my time at Microsoft Research, where I was involved in pushing for what later became the *Windows on Theory* group blog[1]. Every year I would think, "now I will start a blog," and then I would not. What finally pushed me in 2019 was the feeling that, as the saying goes, someone was wrong on the internet. There were too many public discussions about blockchains that did not make sense to me from the point of view of distributed computing. So I started writing. That helped me overcome the hesitation of doing something outside the formal publication format.

The multi-author part was natural to me because I had the group-blog model in mind. We are in the same community, we collaborate, and people can contribute when they have something useful to explain. I sometimes think of it as a kind of proof of work: if you want to be part of the blog, you write posts. Some people contribute a lot for a while and then less; others join later. The blog changes with the people who put work into it.

Kartik: For me, part of the motivation was that consensus and blockchains are difficult to enter from any single community. The classic textbooks are very good, but they frame the area in a particular way. If you come from cryptography, PODC-style distributed computing, systems, or blockchains, people use different language and focus on different questions. When I started working in the area, many of my first insights came from conversations with people like Ling, Ittai, Dahlia, and Sasha, and from building my own mind map of the literature. The blog gradually became a way to put such a mind map in writing.

[1] `https://windowsontheory.org`

Ling: The field of consensus and blockchains is in dire need of good educational resources. The classic textbooks on distributed computing are amazing, but they cover a much broader field. And over the years, the terminology has evolved, and the focus areas have shifted. For someone who wants to learn consensus and blockchains today, the classic textbooks are no longer fully adequate. Decentralized Thoughts has become one of the leading resources for that purpose.

Alin: What I like about the format is that it lets us write down what researchers often know only by reading between the lines. There are many things people explain to each other in conversations, but they are not written down in papers. A blog post gives you a place to spell out that understanding directly.

Your site contains not only a blog but also a large amount of content, including a graduate course, cheat sheets, and videos. Who do you write for, and how do you hope people use the material?

Ittai: There is no single audience. It is a group blog, so different people write for different reasons. Sometimes I write course notes for students. Sometimes I write because there is a widely held view on the internet that I think is wrong or incomplete. Sometimes I want to talk about the philosophy of blockchains, and sometimes I want to explain a new paper. The posts that get the strongest feedback are often not the ones for five specialists, but the ones that give a broader, higher-level understanding of an area.

Alin: For me, the best reason to write is that something is interesting and under-explained. In cryptography, there may be many papers on a topic, but still no clear explanation of what is actually known, which assumptions matter, or how one would use the primitive securely. The first posts I wrote with Ittai were like that: the papers existed, but a newcomer still could not easily understand the whole story. A blog post is a chance to write down that missing understanding. Also, it is simply fun to write.

Kartik: A useful blog post can give the reader a mind map. In consensus, a protocol is a way to control a huge space of possible executions. Papers usually explain the contribution and the difference from prior work. A blog post can go bottom-up: what are the rules, why are they there, what goes wrong if we remove one of them, and how does this protocol relate to nearby ones? That kind of explanation is often what a reader needs in order to actually understand the area.

Ling: I hope the material helps people enter the field. It can be useful for a student, a researcher from a neighboring area, or a practitioner who wants to understand the assumptions and results that matter. We try to use terminology that unifies the classic consensus literature and recent advances in blockchains, and that makes it a practical entry point.

Several of you contrasted blog posts with papers. What can a blog do that a paper usually cannot?

Ittai: A paper is written under very specific incentives. In practice, the incentive is often to get accepted by three reviewers. That is not the same as teaching the reader what is important. A paper is not always written to explain what really happened or what the reader should remember. A blog post has a different goal: it can simply try to improve understanding. It can say, "here is the idea, here is why it matters, here is how it relates to other ideas."

A blog can also be updated. If you later find a better proof, a better explanation, or a useful external resource, you can edit the post and add links. That is much harder with a paper, and even a book update is a big event. A blog post can remain useful for years precisely because it can be improved.

Alin: Blogs are good at collecting scattered knowledge. For a primitive such as ECDSA[2], there may be many papers, but not one place that explains the scheme, public-key recovery, batch verification, malleability, why Bitcoin used it, why Ethereum used it, and what the practical tradeoffs are. A blog can be a focused survey of one idea or one primitive. It does not have to claim novelty; it can focus on clarity.

Kartik: That is important because academic writing often does the opposite. To get a paper accepted, you emphasize the difference from previous work. But when you want to understand an area, it is equally important to know when two things are really the same idea in different languages. A blog post can make those similarities explicit.

The blog also bridges theory and practice. How do you think about that gap?

Alin: On the cryptography side, I see a big gap between the research that is done in academia and what is actually done in practice. Academic papers and practical implementations have different concerns. In a paper, the Fiat–Shamir transform can look almost trivial: hash the transcript. In a real system, this is much more difficult. I see a big gap in terms of concerns and in terms of appreciation for how practical these ideas are when implemented—what is actually possible to do in two months in a team and what is impossible to do. There is no academic interest in how to write secure cryptography—to actually do it in practice where users' funds may be lost if you get the details wrong. I think the blog posts are a great opportunity to discuss some of that.

Ling: On this topic, I'd like to mention a recent article by Moshe Vardi that resonates with me a lot.[3] He argues that theoretical computer science should take

[2] `https://en.wikipedia.org/wiki/Elliptic_Curve_Digital_Signature_Algorithm`

[3] `https://cacm.acm.org/opinion/what-is-theoretical-computer-science/`

more inspiration from physics than from mathematics. In physics, there are beautiful theories, but only those that explain the real world are promoted, while those that do not are ultimately discarded. Much of the success of computer science rests on its relevance to practical systems. We should promote theories that better connect to the real world.

Kartik: In distributed computing, a protocol is trying to shepherd an enormous number of possible executions into a set of good states. It does that by imposing rules. When I read a paper, it is often not clear what each rule is doing, or what breaks if the rule is removed. But that is exactly the understanding one needs in practice. A blog post can spend the time to explain these failure modes and the intuition behind the rules.

Ittai: Practitioners often need the lay of the land. They may not want every technical detail of one paper; they want to know what the space looks like, which assumptions exist, which tools are simple, and where to look next. A blog can give a more neutral overview than a paper whose goal is to sell one particular contribution.

The blog often connects classic distributed computing with modern blockchains. Has the blockchain perspective changed which theoretical questions seem important?

Ling: Yes. Bitcoin proposed an ingenious and highly unconventional solution to consensus. It introduced features that traditional models overlooked: permissionless participation, participants becoming inactive and rejoining, and systems where one may not even know how many other participants are active. These features have opened many new theoretical directions in consensus.

Ittai: For me, blockchains also highlight questions around incentives, economics, and censorship resistance. These topics sit at the boundary of computing, game theory, and economics. I have been saying for years that I want more of that content on the blog, and I keep failing to do it. But I do think it is very important, and censorship resistance is already a good example of a question at that intersection.

Kartik: The blockchain perspective also forces us to compare models and assumptions more carefully. Results that look similar may be stated in different languages by different communities. Conversely, a small model difference may be the crucial issue. The blog helps us make those distinctions visible.

Alin: From the cryptography side, modern blockchain systems also expose the practical meaning of assumptions. A primitive is not only a theorem; it is something that must survive implementation choices, adversarial incentives, and real users. That changes which explanations are useful.

Has writing for Decentralized Thoughts changed your own research process?

Ittai: Writing for a wider audience forces clarity. I sometimes think of Timothy Gowers's essay[4] about the two cultures of mathematics: one more algebraic, building from axioms, and one more combinatorial, with islands of clever constructions. Writing posts has helped me move toward a more structured, almost algebraic view of distributed computing: a landscape of models, protocols, and lower bounds that are connected to each other. Sometimes seeing those connections suggests new theoretical questions.

Kartik: I would not say that writing one post immediately leads to one paper. But it changes how you understand the material. When you try to explain the mind map clearly, you notice missing pieces and start asking questions differently.

Alin: I have seen a blog post directly inspire research. In a project on publicly verifiable secret sharing, we needed a batched zero-knowledge range proof and were stuck. I remembered a short blog post by William Borgeaud about a simple range proof. It was not batched and not zero-knowledge, but after reading it again it became clear that the idea could be batched and adapted to our setting. That few-paragraph explanation helped us find the right direction.

What are you currently most excited about, in your own research and for the future of Decentralized Thoughts?

Kartik: For the blog, it is still on my agenda to write some chapter posts. We have many posts, but a new reader may not know in which order to read them. I would like to see more chapter posts or index posts: pages that gather several posts and turn them into a coherent path through a topic. That would also help with the courses we teach.

Ittai: Recently I spent time editing older posts on purpose, rather than only writing new ones, because some of the older posts are used a lot. That is one advantage of the blog. It can be archival, like a post on selfish mining that remains useful for years, and it can also be timely, almost like social media. When Simplex suddenly became important, for example, we could write about it quickly and relate it to Tendermint while the topic was active.

I would also like to see more material on incentives, economics, and censorship resistance. And I have a hidden goal of connecting more cryptography material to Decentralized Thoughts. I think in the future, people will start learning cryptography with post-quantum primitives. So I think there is an opportunity here to educate people about cryptography again.

Alin: I am very optimistic about high-quality technical writing online. It is searchable and open. If you search for a somewhat obscure topic, a clear blog post is

[4]Gowers, William Timothy. "The two cultures of mathematics." Mathematics: Frontiers and Perspectives 65 (2000): 65.

often much easier to find than the relevant PDF. In the LLM age, avoiding paywalls and writing knowledge in a form that is easy to index may become even more important.

Ling: The insights from the recent blockchain boom give us a lot of new research directions in distributed computing. As I mentioned before, Bitcoin has many interesting and attractive new features. But at the same time, it also has significant downsides, including long latency and energy inefficiency. Achieving the new features of Bitcoin without incurring its downsides has been one of my main research interests in recent years.

Finally, is there anything else you want our readers to know?

Alin: I would encourage young researchers to write. Start a blog, write about what you are learning, or explain your own research. Writing things down is a great way to see whether you really understand them and where the gaps are. It also creates connections with people who care about the same ideas.

Kartik: When you write a blog, the feedback is not always public. Often there are no comments and little visible reaction, but then someone writes or says that they know your work because they read a post. That is another good reason for researchers, especially young researchers, to make their ideas accessible.

Ittai: We are happy to receive high-quality guest posts. In that sense, Kartik and I often act as lightweight editors: someone with expertise sends a post, we help shape it, and the community gets a useful explanation. More broadly, attention is one of the scarce resources in research. Blog posts help focus attention on ideas, and researchers should practice communicating their ideas to a wider audience.

Ling: As much as I enjoy writing on Decentralized Thoughts, I do not think it can fully replace a good textbook. A blog can do much of the work of orientation, but a carefully written textbook is needed for a rigorous and complete understanding of the field. Writing a textbook is a daunting task, but I think I will try.

Thank you for this interview, Ittai, Kartik, Ling, and Alin!