The Barriers to Overcoming Internet Feudalism

HotNets-XVI Dialogue

David Andersen and Xiaowei Yang

DGA: When I started reading this paper, I was nervous we were going to see a reprise of the P2P craze of the early 2000s.

XWY: This paper is about more than P2P—it examines many efforts to decentralize today's Internet components, and lays out a set of research challenges using many of the technological improvements we've seen since those earlier efforts. Anyway, did you change your mind by the end of the paper?

DGA: Halfway—now it feels like P2P with a big dose of blockchain. And as we all know, blockchain makes everything better. No, but seriously, in the end I appreciated that the authors weren't being breathlessly exuberant about the possibilities of decentralization and took a more pragmatic view of how the technical options had changed since the last time the community looked at these issues. For ourselves and our audience, though, let's make sure we're clear on what "Feudalism" means here.

XWY: Definitely. Feudalism is an interesting word. I think they just mean Internet services are becoming interestingly monopolized or centralized. I may choose to use a different word myself. What do you reckon?

DGA: Well, they're clearly hearkening to the idea that there are a few mega-providers and they're very powerful—but I agree, there's some subtlety there. Centralized or monopolized seems to indicate just one, but we've got a handful of providers. Importantly, we're not limited to using only one. I use Google Drive, Dropbox, Google Cloud, AWS, Facebook, and more, and try to pick and choose to get the best of all of them. But I haven't suggested a better term yet...

XWY: Good point. I did miss the point on competition.

DGA: But competitive and decentralized are of course not the same thing. Given that the "decentralize all the things" approach was proposed when we were building DHTs (distributed hash tables) fifteen years ago, and didn't result in a lot of decentralization, what's really changed since then that makes the approach worth re-evaluating?

XWY: That's a great question. One thing is the mobile revolution. Since the first iPhone debuted in 2007, mobile phones have changed the way people use computers. Another emerging technology is blockchains. Who knew we could have completely decentralized currencies, ranging from completely public to surprisingly anonymous, fifteen years ago?

DGA: Well, at least the intriguing theoretical underpinnings—Bitcoin hasn't exactly demonstrated that it can take over the world yet from a scaling perspective, but I'm willing to give it another decade. But you're right; from a research perspective, it's conceivable that we could have decentralized, automatically verifiable payments, and more importantly, global or appropriately-scoped verifiable ledgers for contracts and transactions of all sorts. That does seem like a pretty compelling suite of advances for justifying a re-examination, even if not all of the pieces have arrived yet (and may never arrive) in the blockchain space. But what do the cell phones do for us? Do they make it easier, or *harder*?

XWY: I bet they make it easier. They add more storage, computational power, communi...

DGA: Wait wait wait. What about the fact that they're intermittent? So power-limited? With such expensive

bandwidth?

XWY: ...cation. But you raise a good point. Let me think. But they are also ubiquitous, and storage is getting cheap. Maybe we can use storage to overcome some of their limitations? WiFi is getting more and more available. With that, bandwidth isn't that expensive any more. And they are getting better with power too. But you are right. It is still puzzling: what are the chances we can get it right this time? I wonder if centralization may be the only economically viable solution to cost-effective and high performance services.

DGA: I think I agree based upon the numbers. I take issue with their estimates about the storage needed to create decentralized versions of Google, Facebook, Microsoft, Baidu, Tencent, Apple, Amazon, and the rest. One stat I dug up said that in 2015, global datacenter storage was at 150 Exabytes. Cisco projected that by 2017, the public cloud would have about 400 EB of data. That's a pretty big leg up from the 80 EB estimated in the paper (and larger than the 210 EB they estimate that user devices have). The same thing goes for bandwidth—wide-area bandwidth may only be 200 Tbps, but intra-datacenter bandwidth has absolutely exploded with the recent rise of full-bisection networks and cheap+fast networking gear manufacturers such as Mellanox. And applications such as search use that bandwidth!
But I think they have some compelling use cases. I really like the example of DNS, in particular. It doesn't change *that* fast (the root and TLDs, at least), so the transactional volume might be manageable with blockchains. And it's very heavily read-only; it's already a cache-based system. It even has the beginnings—if not yet deployed very well—of recursive security in the form of DNSSEC. But does decentralizing DNS buy us anything big (other than the potential lack of a central authority to help handle abuse)?

XWY: Well, the "standard" benefits of decentralization include resilience (no single point of failure), scalability, and freedom of choice? But DNS, as a mature system, has overcome many of the performance issues. Root DNS servers are distributed and scale well. But avoiding a central authority itself is already a highly desirable feature, IMHO. I would be able to get whatever name I want for my server without going through a central authority.

DGA: Ah. That raises an important question about the goals of decentralization. One part is philosophical, at heart: Is it better to suffer the tyranny of a central authority, or to subject ourselves to the chaos of rule-by-code? I'm not sure I want to get into that debate. I've read too many posts on the Internet, and it seems to mostly lead to a lot of shouting. Another is whether there are quantitative benefits we can achieve in the decentralized version—and I'd be happy to include "competition" and "having many alternatives" as things we can quantify, just to be sure I'm not shutting down all of the advantages of decentralization in one wave of my hands. Does this paper promise us things beyond a cypherpunk ethos?

XWY: Indeed! What about privacy? It's challenging to have any privacy at all nowadays. Google and Apple now own every bit of my digital life. They probably know me better than myself. Decentralization prevents any centralized party from knowing too much about a user. I think this is valuable and goes beyond an ideological preference.

DGA: An excellent point. Which, I think, brings us back to economics and technical capabilities. I'm pretty sure that the GoogleBookSofts of the world would be happy to store all of our data encrypted with user-managed keys in strong, hardware-enforced sandboxes.... If they could still make the same money and provide the same services to users. Users haven't seemed willing to pay for services instead of accepting advertising, but perhaps they'd be more willing to give up a bit of their already-paid-for device capacity.

XWY: And the paper points out that research progress has given us the (potential) capability to implement some more of these things. Perhaps particularly when you combine them with other technical advances

they didn't discuss, such as trusted execution / attestation.

DGA: And perhaps that's the true value of exploring this question about decentralization: Any of the challenges we overcome in building truly decentralized systems also help us build better, cheaper, and potentially more secure/privacy-friendly distributed (but centrally managed) systems, and that's good for everyone, much as we can see the lingering influence of DHTs on many of today's distributed storage systems. But the paper still made me appreciate the benefits of the feudal model, because the problems they point out seem *really* hard!