Analyzing the 2013 Bitcoin fork: centralized decision-making saved the day

JULY 28, 2015 BY ARVIND NARAYANAN

On March 11, 2013, Bitcoin experienced a technical crisis. Versions 0.7 and 0.8 of the software diverged from each other in behavior due to a bug, causing the block chain to "fork" into two. Considering how catastrophic a hard fork can be, the crisis was resolved quickly with remarkably little damage owing to the exemplary competence of the developers in charge. The event gives us a never-before-never-again look into Bitcoin's inner workings. In this post, I'll do a play-by-play analysis of the dramatic minutes, and draw many surprising lessons. For a summary of the event, see **here**.

First of all, the incident shows the necessity of an effective consensus process for the *human* actors in the Bitcoin ecosystem. The chronology of events was like this: it took about an hour after the fork started for enough evidence to accumulate that a fork was afoot. Once that was understood, things unfolded with remarkable speed and efficiency. The optimal course of action (and, in retrospect, the only one that avoided serious risks to the system) was first proposed and justified 16 minutes later, and the developers reached consensus on it a mere 20–25 minutes after that. Shortly thereafter — barely an hour after the discovery of the fork — the crisis response had effectively and successfully concluded. It took a few hours more for the fork to heal based on the course of action the developers initiated, but the outcome wasn't in doubt.

More surprisingly, it also shows the effectiveness of strong central leadership. That's because the commonsense solution to the fork — as well as the one programmed into the software itself — was to encourage miners running old versions to upgrade. As it turns out, the correct response was exactly the opposite. Even a delay of a few hours in adopting the downgrade solution would have been very risky, as I'll argue, with potentially devastating consequences. Without the central co-ordination of the Bitcoin Core developers and the strong trust that the community places in them, it is inconceivable that adopting this counterintuitive solution could have been successfully accomplished.

Further, two more aspects of centralization proved very useful, although perhaps not as essential. The first is the ability of a few developers who possess a cryptographic key to broadcast alert messages to every client, which in this case was used to urge them to downgrade. The second is the fact that the operator of BTC Guild, a large mining pool at the time, was able to singlehandedly shift the balance of mining power to the old branch by downgrading. If it weren't for this, it would have resulted in a messy "coordination problem" among miners, and we can imagine each one hesitating, waiting for someone else to take the leap.

Since most of the discussion and decision-making happened on the #bitcoin-dev IRC channel, it remains publicly archived and offers a remarkable window into the Core developers' leadership and consensus process. Consensus operated at remarkable speed, in this instance faster than consensus happens in the Bitcoin network itself. The two levels of consensus are intricately connected.

Let's now dive into the play-by-play analysis of the fork and the reaction to it. I've annotated the transcript of selected, key events from the **IRC log** of that fateful night. I've made minor edits to make the log easier to read — mainly replacing the nicknames of prominent community members with their real names, since their identities are salient to the discussion.

Signs of trouble

The first signs that something is wrong come from a miner with nickname thermoman, as well as Jouke Hofman, a Dutch exchange operator, who report strange behavior from their Bitcoin clients. Bitcoin core developer Pieter Wuille helps them debug these problems, but at this point everyone assumes these are problems local to the two users, rather than something on the network. But around 23:00, a variety of other services showing strange behavior are noticed

(blockchain.info, blockexplorer.com, and the IRC bot that reports the status of the network) making it obvious that something's wrong on the network. Luke Dashjr, a prominent developer, spells out the unthinkable:

```
23:06  Luke Dashjr            so??? yay accidental hardfork? :x
23:06  Jouke Hofman           Holy crap
```

Over the next few minutes people convince themselves that there's a fork and that nodes running the 0.8 and the 0.7 versions are on different sides of it. Things progress rapidly from here. A mere five minutes later the first measure to mitigate the damage is taken by Mark Karpeles, founder of Mt. Gox:

```
23:11  Mark Karpeles          I've disabled the import of bitcoin blocks for now
                              until this is sorted out
23:13  Luke Dashjr            I'm trying to contact poolops [mining pool operators]
```

It's pretty obvious at this point that the best short-term fix is to get everyone on one side of the fork. But which one?

Up or down? The critical decision

At 23:18, Pieter Wuille sends a **message** to the bitcoin-dev mailing list, informing them of the problem. But he hasn't fully grasped the nature of the fork yet, stating "We risk having (several) forked chains with smaller blocks" and suggests upgrading as the solution. This is unfortunate, but it's the correct thing to do given his understanding of the fork. This email will stay uncorrected for 45 minutes, and is arguably the only slight misstep in the developer response.

```
23:21  Luke Dashjr            at least 38% [of hashpower] is on 0.8 right now
                              otoh, that 38% is actively reachable
```

Dashjr seems to suggest that the 0.8 → 0.7 downgrade is better because the operators of newly upgraded nodes are more likely to be reachable by the developers to convince them to downgrade. This is a tempting argument. Indeed, when I describe the fork in class and ask my students why the developers picked the downgrade rather than the upgrade, this is the explanation they always come up with. When I push them to think harder, a few figure out the right answer, which Dashjr points out right afterward:

```
23:22  Gavin Andresen         the 0.8 fork is longer, yes? So majority hashpower is 0.8....
23:22  Luke Dashjr            Gavin Andresen: but 0.8 fork is not compatible
                              earlier will be accepted by all versions
```

Indeed! The behavior of the two versions is not symmetric. Upgrading will mean that the fork will persist essentially indefinitely, while downgrading will end it relatively quickly.

(Lead developer) Gavin Andresen still protests, but Wuille also accepts Dashjr's explanation:

```
23:23  Gavin Andresen         first rule of bitcoin: majority hashpower wins
23:23  Luke Dashjr            if we go with 0.8, we are hardforking
23:23  Pieter Wuille          the forking action is a too large block
                              if we ask miners to switch temporarily to smaller blocks again,
                              we should get to a single chain soon
                              with a majority of miners on small blocks, there is no risk
23:24  Luke Dashjr            so it's either 1) lose 6 blocks, or 2) hardfork for no benefit
23:25  BTC Guild              We'll lose more than 6
```

BTC Guild was a large pool at the time, and its operator happened to be online. They are correct — the 0.8 branch had 6 blocks at the time, but was growing much faster than the 0.7 branch and would continue to grow until the latter gradually caught up. Eventually 24 blocks would be lost. BTC Guild will turn out to be a key player, as we will soon see.

More explanation for why downgrade is the right approach:

```
23:25  Pieter Wuille          all old miners will stick to their own chain
                              regardless of the mining power behind the other
23:25  Luke Dashjr            and the sooner we decide on #1, the fewer it loses
23:26  Pieter Wuille          even with 90% of mining power on 0.8
                              all merchants on an old client will be vulnerable
23:26  Luke Dashjr            if we hardfork, all Bitcoin service providers have an emergency situation
23:30  Pieter Wuille          and we _cannot_ get every bitcoin user in the world
                              to now instantly switch to 0.8
                              so no, we need to rollback to the 0.7 chain
```

Achtung!

Many Bitcoin users are not aware of the centralized ability in Bitcoin for a few developers to send out alert messages. Not only does it exist, it becomes crucial here, as Core developer Jeff Garzik and another user point out:

```
23:31  Jeff Garzik            alert is definitely needed
23:31  jrmithdobbs            also, if this isn't an alert message scenario
                              I don't know what is, where's that at? :)
```

A bit of a comic interlude from the operator of ozcoin, another mining pool:

```
23:31  Graet                  ozcoin wont even get looked at for an hour or so
23:31  Graet                  no-one is avalable and i need to take kids to school
```

The situation is urgent (the more clients upgrade, the harder it will be to convince everyone to downgrade):

```
23:32  phantomcircuit         0.7 clients are already displaying a big scary warning
                              Warning: Displayed transactions may not be correct!
                              You may need to upgrade, or other nodes may need to upgrade.
                              unfortunately i suspect that warning will get people to upgrade
```

Other complications due to custom behavior of some miners' software:

```
23:35  lianj                  oh, damn. anyhow please keep us guys updated which code change is made
                              to solve the problem. our custom node does .8 behavior
```

Gavin Andresen and Jeff Garzik still aren't convinced (they seem to be thinking about getting 0.8 nodes to switch to the other branch, rather than the more blunt solution of asking miners to downgrade the client)

```
23:34  Jeff Garzik            and how to tell bitcoind "mine on $this old fork"
23:35  Gavin Andresen         exactly. Even if we want to roll back to the 0.7-compatible chain,
                              I don't see an easy way to do that.
```

This shows the usefulness of the developers having direct channels to the pool operators, another benefit of central co-ordination:

```
23:38  Luke Dashjr            FWIW, Josh (EclipseMC) has to be on a plane in 20 minutes,
                              so he needs this decided before then :/
```

As time goes on the solution only gets harder, as illustrated by a new user wading into the channel.

```
23:39  senseless              So whats the issue?
                              I got the warning you need to upgrade. So I upgraded.
```

Gavin Andresen, notable for his cautious approach, brings up a potential problem:

```
23:40  Gavin Andresen         If we go back to 0.7, then we risk some other block
                              triggering the same condition.
```

Happily, as others pointed out, there's nothing to worry about — once majority hashpower is on 0.7, other blocks that have the same condition will be harmless one-block forks instead of a hard fork.

The BTC guild operator offers to basically end the fork:

```
23:43  BTC Guild             I can single handedly put 0.7 back to the majority hash power
                             I just need confirmation that thats what should be done
23:44  Pieter Wuille         BTC Guild: imho, that is was you should do,
                             but we should have consensus first
```

So much for decentralization! The fact that BTC Guild can tip the scales here is crucial. (The hash power distribution at that time appears to be roughly 2/3 vs 1/3 in favor of the 0.8 branch, and BTC Guild controlled somewhere between 20% and 30% of total hash power.) By switching, BTC Guild loses the work they've done on 0.8 since the fork started. On the other hand, they are more or less assured that the 0.7 branch will win and the fork will end, so at least their post-downgrade mining power won't be wasted.

If mining power were instead distributed among thousands of small independent miners, it's far from clear that coordinating them would be possible at all. More likely, each miner on the 0.8 branch would wait for the 0.7 branch to gain the majority hash power, or at least for things to start heading clearly in that direction, before deciding to downgrade. Meanwhile, some miners in the 0.7 branch, seeing the warning in their clients and unaware of the developer recommendation, would in fact *upgrade*. The 0.8 branch would pull ahead faster and faster, and pretty soon the window of opportunity would be lost. In fact, if the developers had delayed their decision by even a few hours, it's possible that enough miners would have upgraded from 0.7 to 0.8 that no single miner or pool operator would be able to reverse it singlehandedly, and then it's anybody's guess as to whether the downgrade solution would have worked at all.

Back to our story: we're nearing the critical moment.

```
23:44  Jeff Garzik           ACK on preferring 0.7 chain, for the moment
23:45  Gavin Andresen        BTC Guild: if you can cleanly get us back on the 0.7 chain,
                             ACK from here, too
```

Consensus is reached!

Time for action

Right away, developers start giving out advice to downgrade:

```
23:49  Luke Dashjr           surge_: downgrade to 0.7 if you mine, or just wait
23:50  Pieter Wuille         doublec: do you operate a pool?
23:50  doublec               yes
23:50  Pieter Wuille         doublec: then please downgrade now
```

BTC Guild gets going immediately…

```
23:51  BTC Guild             BTC Guild is going back to full default block settings and 0.7 soon.
00:01  BTC Guild             Almost got one stratum node moved
```

… even at significant monetary cost.

```
23:57  BTC Guild             I've lost way too much money in the last 24 hours
                             from 0.8
```

One way to look at this is that BTC Guild sacrificed revenues for the good of the network. But these actions can also be justified from a revenue-maximizing perspective. If the BTC Guild operator believed that the 0.7 branch would win anyway (perhaps the developers would be able to convince another large pool operator), then moving first is relatively best, since delaying would only take BTC Guild further down the doomed branch. Either way, the key factor enabling

BTC Guild to confidently downgrade is that by doing so, they can ensure that the 0.7 branch will win.

Now that the decision has been taken, it's time to broadcast an alert to all nodes:

```
00:07  Gavin Andresen        alert params set to relay for 15 minutes, expire after 4 hours
```

The alert in question is a model of brevity: *"URGENT: chain fork, stop mining on version 0.8"*

At this point people start flooding the channel and chaos reigns. However, the work is done, and only one final step remains.

At 00:29, Pieter Wuille **posts to bitcointalk**. This essentially concludes the crisis response. The post said, in its entirety:

> *Hello everyone,*
>
> *there is an emergency right now: the block chain has split between 0.7+earlier and 0.8 nodes. I'll explain the reasons in a minute, but this is what you need to know now:*
>
> - *After a discussion on #bitcoin-dev, it seems trying to get everyone on the old chain again is the least risky solution.*
> - *If you're a miner, please do not mine on 0.8 code. Stop, or switch back to 0.7. BTCGuild is switching to 0.7, so the old chain will get a majority hash rate soon.*
> - *If you're a merchant: please stop processing transactions until the chains converge.*
> - *If you're on 0.7 or older, the client will likely tell you that you need to upgrade. Do not follow this advise – the warning should go away as soon as the old chain catches up*
> - *If you are not a merchant or a miner, don't worry.*

Crucially, note that he was able to declare that the 0.7 branch was going to win due to BTC Guild switching to it. This made the downgrade decision the only rational one for everyone else, and from here things were only a matter of time.

What would have happened if the developers had done nothing?

Throughout the text I've emphasized that the downgrade option was the correct one and that speed of developer response was of the essence. Let's examine this claim further by thinking about what would have happened if the developers had simply let things take their course. Vitalik Buterin **thinks** everything would have been just fine: "if the developers had done nothing, then Bitcoin would have carried on nonetheless, only causing inconvenience to those bitcoind and BitcoinQt users who were on 0.7 and would have had to upgrade."

Obviously, I disagree. We can't know for sure what would have happened, but we can make informed guesses. First of all, the fork would have gone on for far longer — essentially until every last miner running version 0.7 or lower either shut down or upgraded their software. Given that many miners leave their setups unattended and others have custom setups that aren't easy to upgrade quickly, the fork would have lasted days. This would have several effects. Most obviously, the psychological impact of an ongoing fork would have been serious. In contrast, as events actually turned out, the event happened overnight in the US and had been resolved the next morning, and media coverage praised the developers for their effective action. The price of Bitcoin dropped by 25% during the incident but recovered immediately to almost its previous value.

Another adverse impact is that exchanges or payment services that took too long to upgrade their clients (or disable transactions) might find themselves victims of large double-spend attacks. As it happened, OKPay **suffered** a $10,000 double spend. This was done by a user trying to prove a point and who revealed the details publicly; they got lucky in that their payment to OKPay was confirmed by the 0.8 branch but not 0.7. A longer-running fork would likely have exacerbated the problem and allowed malicious attackers to figure out a systematic way to create double-spend transactions. [1]

Worse, it is possible, even if not likely, that the 0.7 branch might have continued *indefinitely*. Obviously, if this did happen, it would be devastating for Bitcoin, resulting in a fork of the currency itself. One reason the fork might keep going is because of a "**Goldfinger attacker**" interested in de-stabilizing Bitcoin: they might not have the resources to execute a 51% attack, but the fork might give them just the opportunity they need: they could simply invest resources into keeping the 0.7 fork alive instead of launching an attack from scratch.

There's another reason why the fork might have never ended. Miners who postponed their decision to switch from 0.7 to 0.8 by, say, a week would face the distasteful prospect of forgoing a week's worth of mining revenue. They might instead gamble and continue to operate on the 0.7 branch as a big fish in a small pond. If the 0.7 branch had, say, 10% of the mining power of the 0.8 branch, the miner's revenue would be multiplied tenfold by mining on the 0.7 branch. Of course, the currency they'd earn would be "Bitcoin v0.7", which would fork into a different currency from "Bitcoin v0.8", and would be worth much less, the latter being considered the legitimate Bitcoin. We analyze this type of situation in Chapter 7, "**Community, Politics, and Regulation**" of our **Bitcoin textbook-in-progress** or the corresponding sections of the **video lecture**.

While the exact course of events that would have resulted from inaction is debatable, it is clear that the downgrade solution is by far the less risky one, and the speed and clearheadedness of the developers' response is commendable.

All this is in stark contrast to the **dysfunctional state** of the consensus process on the block size issue. Why is consensus on that issue failing? The main reason is that unlike the fork, there is no correct solution to the block size issue; instead there are various parties with differing goals that aren't mutually aligned. Further, in the case of the fork, the developers had a well-honed process for coming to consensus on technical questions including bugs. For example, it was obvious to everyone that the discussion of the fork should take place on the #bitcoin-dev IRC channel; this didn't even need to be said. On the other hand, there is no clear process for debating the block size issue, and the discussion is highly fragmented between different channels. Finally, once the developers had reached consensus about the fork, the community went with that decision because they trusted the developers' technical competence. On the other hand, there is no single entity that the Bitcoin community trusts to make decisions that have economic implications.

Conclusion

In summary, we have a lot to learn from looking back at the fork. Bitcoin had a really close call, and another bug might well lead to a different outcome. Contrary to the view of the consensus protocol as fixed in stone by Satoshi, it is under active human stewardship, and the quality of that stewardship is essential to its security. [2] Centralized decision-making saved the day here, and for the most part it's not in conflict with the decentralized nature of the network itself. The human element becomes crucial when the code fails or needs to adapt over time (e.g., the block size debate). We should accept and embrace the need for a strong leadership and governance structure instead of treating decentralization as a magic bullet.

[1] This gets into a subtle technical point: it's not obvious how to get a transaction to get into one branch but not the other. By default any transaction that's broadcast will just get included in both branches, but there are several ways to try to subvert this. But given access to even one transaction that's been successfully double-spent, an attacker can amplify it to gradually cause an arbitrary amount of divergence between the two branches.

[2] To underscore how far the protocol is from being fixed for all time by a specification, the

source code of the reference implementation is the only correct documentation of the protocol. Even creating and maintaining a compatible implementation has proved to be near-infeasible.

Thanks to Andrew Miller for comments on a draft.

FILED UNDER: CRYPTOCURRENCIES TAGGED WITH: BITCOIN, GOVERNANCE

Comments

Anonymous says:
July 28, 2015 at 3:19 pm

You missed the two most important lines of that IRC chat:

This one line helps understand which side the “economic majority” is on:
http://bitcoinstats.com/irc/bitcoin-dev/logs/2013/03/11#l1363043241.0
23:07 Luke-Jr MagicalTux: is MtGox on the 0.8 or pre-0.8 side of the fork?

and this one shows the majority of mining capacity abandon its lead on the fork to get on the same side as where the economic majority is:

http://bitcoinstats.com/irc/bitcoin-dev/logs/2013/03/11#l1363044661.0
23:31 Eleuthria Working on rolling back BTC Guild to 0.7 now.

Christopher Camp says:
July 30, 2015 at 10:24 am

Another great piece on the bitcoin governance issue.

Such an interesting tension within the bitcoin community – the distrust of government and the need for governance. A real challenge and not clear how it will play out. But certainly interesting to watch.

Beasley says:
August 1, 2015 at 1:15 am

This is the classic argument for big government creep: “if only we had a trusted party in charge, they could take care of all our problems.” Life is fraught with risk. Crap happens.
You said it yourself: “the human element becomes crucial …”
GOVERNMENT HUMAN

> **Beasley says:**
> August 1, 2015 at 1:17 am
>
> To clarify …
> Intent of the last line was GOVERNMENT is not equal to HUMAN

> **unsystemizer says:**
> August 1, 2015 at 11:38 pm
>
> Well said.
>
> By the way, this Committee to Save the World didn’t seem to benefit Wall St, not even the Wall St of BTC (despite that interesting comment about the economic majority, above; don’t forget that MtGox itself wasn’t the economic majority, but their customers were and that is different from the leveraged lunacy on top of a fractional reserve system).
>
> Yet the fakers from the Govt were celebrated in the media (a feature), while this was merely a bug.