Fundamentals of Proof of Work

David Vorick
Jan 14, 2019

2019 is the year of the 51% attack. Once a problem only for cryptocurrencies of negligible value, high reputation and high market cap cryptocurrencies are now finding themselves victim to double spends, with exchanges taking the brunt of the damage.

As the attacks continue to grow in frequency and severity, exchanges are beginning to take steps to protect themselves. Originally this meant increasing the number of confirmations, however as

the attacks have expanded from tens of blocks to hundreds of blocks, the effectiveness of this strategy is being called into question.

Without a significant course correction, we can expect the damages to grow, even to the point where exchanges may begin to fold. These 51% attacks are successful because of fundamental weaknesses in protocols of the targeted cryptocurrencies, and exchanges will ultimately need to be much more restrictive when selecting which cryptocurrencies to support.

Game Theory and Threat Models

Many decentralized protocols assume that at least 51% of all participants will participate honestly. Bitcoin has been successful because the protocol designers realized that this assumption is inadequate for real world decentralized protocols. In the anonymous, unregulated Internet, participants are free to act as economic agents, often with few consequences for deviant behavior. Instead of assuming that greater than 51% of all actors will be acting honestly, Bitcoin

assumes that greater than 51% of all actors will be acting according to their best economic interest.

This threat model is substantially less forgiving. Instead of assuming that most participants will follow the protocol faithfully, Bitcoin developers assume that participants will proactively seek out ways to deviate from the protocol if those deviations can result in profit. This assumption greatly restricts the flexibility in protocol design choices, but has proven to be a crucial requirement for success out on the open Internet.

Bitcoin developers strive for something called incentive compatibility. If a protocol has incentive compatibility, it means that the optimal decision for each individual from their own perspective is also the optimal decision for the group as a whole. When protocols are incentive-compatible, individuals can be completely selfish because those selfish actions will benefit the group as well.

The game theory that keeps Bitcoin running securely is complex and often quite subtle. Many of the cryptocurrencies that have attempted to copy Bitcoin's protocol design have made changes that have broken the incentive compatibility that is critical to keeping Bitcoin secure. As a result, these cryptocurrencies are not secure, and the deluge of double spend attacks is a clean demonstration that not everything is in order.

Though altcoin designers have broken incentive compatibility in many ways, nothing has been more beneficial to the recent double spend attacks than the decision to use shared hardware as the means for blockchain security. When the same hardware is able to mine on multiple cryptocurrencies, critical incentive compatibilities break down.

There are two primary categories of cryptocurrencies with shared hardware. The first and most prominent category covers the ASIC resistant cryptocurrencies. ASIC resistant cryptocurrencies actually have a goal of using shared hardware; the belief is that security is increased because more widely available hardware will lead to greater hashrate decentralization. The second category of shared hardware cryptocurrencies is cryptocurrencies that are ASIC mined but share the same algorithm as some other cryptocurrency. When multiple cryptocurrencies share the same proof of work algorithm, the same hardware (even if that hardware is specialized) is able to target any of the cryptocurrencies and this disrupts the incentive compatibility in many of the same ways that ASIC resistance does.

What Has Changed Since 2017

Shared hardware has been a theme in cryptocurrency for many years, and yet only recently have high profile 51% attacks become a problem. Truthfully, these attacks have become possible recently for the simple reason that the industry has become more sophisticated. Better tools exist, smarter attackers exist, and in general there is just more and better infrastructure. While this infrastructure has largely benefited honest participants more than anyone else, it has also benefited attackers, and made it easier for sophisticated individuals to attack insecure cryptocurrencies.

We're going to be looking at a few of the developments which have been more important to 51% attacks, but even without these specific developments I believe that we would have eventually started to see high profile 51% attacks on shared hardware cryptocurrencies anyway. Shared hardware is simply a fundamentally insecure means to protect a blockchain against double spend attacks.

Hashrate Marketplaces

One of the key developments in enabling recent attacks has been the maturing of hashrate marketplaces. For shared hardware cryptocurrencies, knowing the most profitable cryptocurrency to mine at any particular moment requires a high degree of sophistication. Hashrate marketplaces allow hardware owners to rent their hardware out to more sophisticated miners, increasing the profits of all participants in the hashrate marketplace.

A side effect of hashrate marketplaces is that attackers now have a great pool of hardware that they can draw from quickly and temporarily when attempting an attack. Before hashrate marketplaces existed, attacking a cryptocurrency with 100,000 GPUs defending it more or less required owning 100,000 GPUs. Attacks of that scale would require many tens of millions of dollars to execute, which meant that heavily mined GPU coins were largely safe. After the development of hashrate marketplaces, the same 100,000 GPUs can be rented for several hours

at a cost of just tens of thousands of dollars. Hashrate marketplaces cut the security margin of shared hardware cryptocurrencies by multiple orders of magnitude.

We also have to expect that hashrate marketplaces for shared hardware will only continue to grow, because all participants benefit from joining a hashrate marketplace—hashrate marketplaces make mining more efficient.

These hashrate marketplaces don't make nearly as much sense for exclusive hardware cryptocurrencies. The benefit of a hashrate marketplace is that they help hardware owners avoid the complexity of deciding what to mine to make the most money. In an exclusive hardware cryptocurrency, there is only ever one thing to mine, which means there is not much to gain from joining a marketplace.

There is another critical game theory element at play with hashrate marketplaces. When a miner offers shared hardware up to a hashrate marketplace, there is a chance that the hardware will be abused to commit an attack. The shared hardware operator however is not incentivized to care, because the attacker is likely paying a small premium for the hardware (due to the need for burst access), and because the underlying hardware does not lose value if one of the cryptocurrencies that it targets is hit with a big attack—there are plenty of other sources of value for that hardware.

Exclusive hardware on the other hand can only derive value from the single cryptocurrency that it is able to target. Offering up exclusive hardware to an attacker is far riskier, because a successful attack has a more direct impact on the value of the hardware that is used. All hardware providers participating in a hashrate marketplace risk being wiped out by a successful attack on their sole source of income, and therefore are incentivized away from participating in marketplaces that reduce the security margins of the underlying cryptocurrency.

Large Mining Farms

The appearance of large mining farms has also played a big role in reducing the security of shared hardware cryptocurrencies. Many large mining farms exist that exceed 10,000 GPUs, multiple mining farms exist that exceed 100,000 GPUs, and the largest of the mining farms has well in excess of 500,000 GPUs.

From a security perspective, this means that any GPU mined cryptocurrency with less than 500,000 GPUs worth of hashrate on it can be single handedly 51% attacked by the largest mining farm. Cryptocurrencies with less than 100,000 GPUs mining on them are vulnerable to not just one farm, but multiple farms that are each capable of single-handedly launching a 51% attack and executing a double spend. Cryptocurrencies protected by less than 10,000 GPUs worth of hashrate are pretty much trivially vulnerable to attack.

Many of these GPU mining farms are purely motivated by profit, sharing little if any of the ideology of the cryptocurrency space. To some of these farms, if there is a way to make more money, then that is the best course of action, even if there is collateral damage to the underlying ecosystem.

Exclusive hardware addresses this in two ways. The first is that for exclusive hardware cryptocurrencies, there can fundamentally be at most only one mining farm that is capable of launching a 51% attack. Though it's not a fantastic guarantee by itself, exclusive hardware cryptocurrencies are guaranteed to have to trust at most one entity. This is contrasted against the vast majority of ASIC resistant cryptocurrencies—most ASIC resistant cryptocurrencies could be attacked at any time by any of a multitude of different mining farms.

The more significant advantage of exclusive hardware is incentive alignment. For profit maximizing mining farms, profit is generally not possible by attacking an exclusive hardware cryptocurrency because the attack is going to reduce the value of the mining farm's hardware. Even in the situation where one mining farm holds enough hashrate to commit a 51% attack, that mining farm is incentivized against executing that attack, because the total value of the hardware owned by the farm is greater than the total amount of money that the farm would be able to steal in an attack.

Increased Attacker Budgets and Sophistication

One of the major differences between cryptocurrency in 2019 and cryptocurrency in 2017 is that the space is a lot more valuable, the theory is a lot better understood, and the number of experts is a lot higher.

In 2017, the number of people who understood that these vulnerabilities existed was not very high. Further, the value of a typical cryptocurrency was also not very high, meaning even for individuals who knew how to execute an attack, there wasn't much profit to be had by performing an attack.

In 2019, there are a lot more people out there who understand how cryptocurrencies work, and who understand how to attack cryptocurrencies that have fundamental flaws. Further, the potential payoff of committing a successful attack is much higher today, meaning that a larger

percentage of capable individuals are going pursue attacks. The increased rewards also mean that attackers can commit more time, money, and resources to engaging an attack.

This is a trend that is going to continue. Today we are seeing 51% attacks because they are the lowest hanging fruit with the highest payoff. However many of the major popular dapps today have fundamental weaknesses, and as they grow in value and as attackers grow in sophistication, those fundamental weaknesses are going to increasingly be exploited. In particular, I have concerns for most of the cryptocurrency projects involving (in order of concern): novel consensus algorithms, on-chain governance, oracles, stablecoins, prediction markets—among other things. It's often not the core ideas themselves that are broken, but rather the specific designs and implementations. This space currently suffers from a lack of peer review; many of the high profile projects deployed in our ecosystem have not been adequately reviewed and likely have significant active vulnerabilities.

Hardware Bear Markets

Hardware bear markets are a problem that impacts both shared hardware and exclusive hardware cryptocurrencies. If the value of mining hardware falls to the point where it is no longer profitable to mine, the hardware can become very cheap for an attacker to acquire.

The recent cryptocurrency bear market has substantially reduced the value of a lot of mining hardware, which simultaneously means that cryptocurrencies have a lower active total hashrate defending them and also means that attackers have much cheaper sources for renting or buying hardware.

The GPU marketplace is getting hit by a second big impact: there are now ASICs available for both Ethereum and Zcash. These two cryptocurrencies were previously driving most of the GPU hashrate, and that hashrate is slowly being pushed out by ASICs, which dramatically reduces the cost of renting GPUs to attack the lower value cryptocurrencies. As ASICs continue to come to market for the high value GPU cryptocurrencies, we can expect this effect to exacerbate, and 51% attacks will become increasingly common and inexpensive. I do not see this trend reversing, even with novel attempts at ASIC resistance on the horizon.

Bitcoin is also getting hit with a hardware bear market. It's estimated that as much as 1/3rd of the Bitcoin hashrate has been put up for fire sale by mining farms that are now insolvent. S9's are available today at prices far below the manufacturing cost, and while it doesn't seem like this is a security issue for Bitcoin yet, it may become an issue if the price falls another 2–4x.

The manufacturers themselves have been hit extremely hard by the bear market. It's estimated that Bitmain, Innosilicon, TSMC, and even Samsung all suffered substantial losses due to the sudden price drop, and because of that it's less likely that we will see heavy over-production in the future—we now see that heavy production is very risky, and Bitcoin is now at a scale where companies are unwilling to take such high risk positions. My guess is that this is the most severe hardware bear market Bitcoin will ever see.

Other exclusive hardware cryptocurrencies however are not as large as Bitcoin, and hardware manufacturers may be more willing to risk overproduction, which in turn could cause hardware bear markets for those cryptocurrencies in the event of a sudden price drop or other turmoil.

The Impact of the Block Reward

Because hardware is very financially expensive to obtain and operate, the security of a cryptocurrency against double spend attacks is highly dependent on its block reward. The total amount of protection that a cryptocurrency receives is proportional to the amount of hardware protecting it, and if a low block reward prevents any substantial amount of hardware from mining the cryptocurrency, the cryptocurrency will not have any substantial amount of security.

In general, we need to be thinking about security in terms of how many dollars a 51% attack would cost. If the total value of hardware mining a cryptocurrency is one million dollars, then we can expect that any trade over one million dollars is strictly vulnerable to a 51% attack, because the counterparty to that trade could have just spent a million dollars buying or manufacturing enough new hardware to commit the double spend attack.

It's difficult to appraise the total value of hardware mining a cryptocurrency, and difficult to appraise the cost of manufacturing a new set of hardware that's worth enough to perform a 51% attack, but as a general rule of thumb it's between 6 and 24 month's worth of block reward. The open competitiveness of hardware mining generally ensures that it will be in that range.

This helps us to apply a maximum safe transaction value to a cryptocurrency, however before picking a value we need to talk about the phrasing 'double spend'. The truth is that a double spend could really be a triple spend or quadruple spend, or whatever multiple of spend that allows an attacker to be successful. A single double spend attack could simultaneously double spend a dozen different exchanges all at once. So it's actually not enough to consider a single transaction when contemplating security against a double spend attack, we need to also consider that other attacks may be happening simultaneously.

The actual upper bound for transaction value is going to be specific to each cryptocurrency, and depends on many factors that go beyond just the block reward. But as a general rule of thumb, I would start to get nervous about transactions that are larger than 1 month worth of block rewards for exclusive hardware cryptocurrencies, and I would get nervous about transactions that are larger than one hour worth of block reward for cryptocurrencies with large established hashrate marketplaces.

Cryptocurrency Shorts

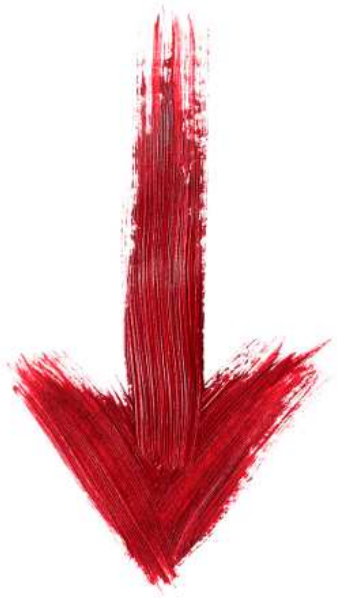

A short is essentially a loan. When you take out a short on a cryptocurrency, you are taking out a loan for a number of coins where you agree to return the same number of coins (usually plus some interest) in the future. Typically, when a person takes out a short they sell the coins immediately and then hope that the price drops so that they can buy them back cheaper and return them, having made a profit in the process.

Shorts require two sides. There is the person taking out the short or the loan, and then there is the person providing the loan. When it comes to cryptocurrencies, there is an important bonus element of tension between the person taking out the loan and the person providing the loan: the person taking out the loan may be using that money to attack the cryptocurrency and crash the price. An attack may be a double spend, or an attack may simply be a denial of service, where the attacker mines empty blocks forever. Or, depending on the cryptocurrency, there may be other advanced attacks that are being planned.

I bring this up for two reasons. The first is to warn exchanges and market participants against enabling short markets. If you are offering cryptocurrency loans, you are potentially funding attackers who will devalue the very asset you hope to get back in the future. Offering shorts for cryptocurrencies is substantially more risky than offering shorts for traditional markets.

The other reason is that a large short market increases risk for other parties depending on the security of that cryptocurrency. If a large short market exists for a cryptocurrency, then a potential attacker has a big source of capital that they can use to fund an attack, and if the attack is successful they will not need to return much of that capital. Therefore exchanges and other users should be particularly wary / avoiding of cryptocurrencies that have large short markets.

Limitations of Increasing the Confirmation Time

A common response to network turmoil is to increase the confirmation time for deposits. And in a lot of cases, this is good advice: increasing the confirmation time is sometimes very helpful in avoiding certain types of risks. However, sometimes increasing the confirmation time is not useful at all, and offers no additional practical protection.

One of the biggest areas that increased confirmation times help is with turmoil in the peer to peer network. If for some reason blocks are propagating slowly, or if the network splits in half, or if some peers are trying to withhold blocks or commit routing layer attacks, then increasing the number of confirmations can be very beneficial. Changing from 60 minute confirmation times to 24 hour confirmation times means that the longest chain has more time to propagate, the network split has more time to heal, or the routing layer attack has more time to be addressed.

Another place that increased confirmation times can help is during times of selfish mining, or during times of a rogue <50% hashrate miner. When there is heavy selfish mining, or if for some reason a large miner is mining weird or incorrect blocks, the chance of large reorgs goes up substantially. Instead of typically seeing 2–3 block reorgs, you might start seeing reorgs that are as many as a dozen blocks deep. However, because there is no 51% attack, it's highly unlikely that you will see reorgs that go beyond a few dozen blocks. The network will generally still move in one direction.

For actual 51% attacks, increasing the confirmation time has a much lesser impact. Raising the confirmation time from 60 minutes to 6 hours will increase the amount of hashrate that an attacker needs to rent, or will increase the amount of time that a mining farm needs to spend on an attack, however this is really only going to be an effective tactic for cryptocurrencies right on the threshold of being attackable.

Something important to keep in mind is that when a cryptocurrency gets hit with a 51% attack, the attacker gets the full block rewards for all of the blocks that they mine. If the price only falls a bit following the attack, the attack will actually fund itself. This is one of the key reasons that increasing confirmation times does not help for small GPU mined cryptocurrencies. An attacker may be able to mine a whole week's worth of blocks with just a few hours of hashrate rented from a marketplace, especially if that cryptocurrency is very small or has a low block reward.

Limitations of Address Blacklisting

One thing that has thwarted attackers previously is emergency blacklists applied to exchanges. When an attacker performs a double spend, they have to extract the money somehow. This usually involves transferring the money to another exchange and then trading further. Exchanges have been able to stop thefts and double spends in the past by blacklisting any addresses involved in a double spend attempt—one exchange will tell the others which addresses are problematic, and then the exchanges work together to ensure the money is returned.

Although this is sometimes effective, attackers will be increasingly able to get around this security measure. Whether it is by using privacy coins, or whether it's by delaying the actual

double spend until the stolen cryptocurrency has been moved to a wider set of wallets, or whether it's by using decentralized exchanges instead of centralized exchanges to extract value, address blacklisting will get increasingly ineffective as attackers get more sophisticated.

This doesn't mean that exchanges should stop using address blacklisting. It's a good technique that has recovered lots of stolen funds. But exchanges shouldn't be depending on address blacklisting to save their funds in the event of an attack, because many times address blacklisting will fail to recover funds.

Recommendations to Limit Risk

Though the situation is grim, especially for exchanges, there are a few things we can do to at least temporarily mitigate risk for some of the larger shared hardware cryptocurrencies. Ultimately, these mitigations can all be circumvented by a sufficiently sophisticated attacker, and fundamental developments in the space such as decentralized exchanges and decentralized hashrate marketplaces are also going to eventually nullify these mitigations. The only established long term solution is to require all cryptocurrencies to switch to exclusive hardware—every cryptocurrency on an ASIC friendly algorithm, and every cryptocurrency on a **different** ASIC friendly algorithm. But perhaps we can buy a little bit of time with reduced risk while everyone is given a chance to migrate.

Tracking Global Hardware Availability

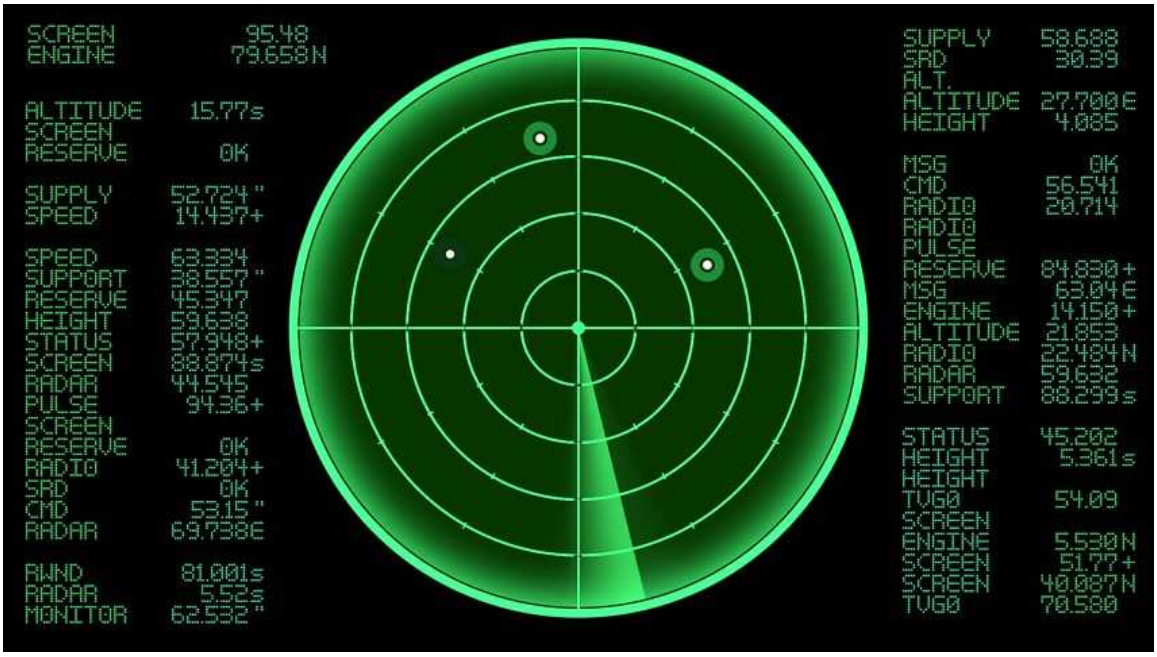

One of the things that can help exchanges to manage risk is to keep an eye on the global hardware availability for each cryptocurrency. The percentage of useful hardware that is mining on a particular cryptocurrency is a good indicator of how much security that cryptocurrency has.

For exclusive hardware cryptocurrencies, the only thing that you really need to watch out for is low block rewards and hardware bear markets. If, for example, the majority of the hardware that once targeted a cryptocurrency is now no longer mining due to low profitability, then the cost of an attack is likely very low because hardware can likely be purchased by an attacker at a very low price. For all other situations, exclusive hardware cryptocurrencies are likely secure against hashrate attacks.

For shared algorithm cryptocurrencies that have ASICs or other highly specialized hardware, the key thing to look at is how much hashrate is mining each cryptocurrency. For cryptocurrencies that have more than 70% of the total hashrate actively mining, I would say there's not much to worry about. For cryptocurrencies with between 10% and 70% of the total hashrate, I would say that 24 hour confirmation times are prudent. Even at 70% hashrate, there are games that larger mining farms could play to commit attacks and potentially succeed at executing double spends. With 24 hour confirmation times, these attacks become a lot less feasible. The shared algorithm cryptocurrencies with less than 10% of the total hashrate are likely insecure. The decision to halt deposits and withdrawals is of course always dependent on risk tolerance and other factors, however my general recommendation would be to halt deposits and withdrawals on these cryptocurrencies until the hashing algorithm is changed to something more secure.

For GPU mined cryptocurrencies, risk management really requires understanding the current state of the hashrate marketplaces and the state of the large mining farms that are in operation.

Though I have not spent a ton of time or rigor with these values, my estimate is that there is currently a total of between 100 million and 250 million dollars worth of GPUs available on hashrate marketplaces today. This number is critical for determining whether a cryptocurrency is vulnerable to a 51% attack. This alone is not sufficient however, as there have been reports which strongly suggest that certain large mining farms have also been participating in 51% attacks against smaller cryptocurrencies. In particular, at least one of the farms in the 10 million to 100 million dollars of GPUs range has seemed willing to attempt attacks.

Given the above, my recommendation for today would be to require 24 hours of confirmations for all GPU mined cryptocurrencies that have between 50 and 250 million dollars of hardware actively mining on them, and to disable deposits for all cryptocurrencies below this threshold. Below 50 million dollars of hardware, the cost and difficulty of mounting an attack just does not seem to be very high.

As the ecosystem evolves and the state of both large mining farms and hashrate marketplaces changes, the risk analysis for cryptocurrencies of various sizes and algorithm types will be changing. Exchanges who stay on top of these changes will have more accurate risk analyses and will be more able to make the best business decisions.

Relationships With Mining Farms and Hashrate Marketplaces

Some of the total risk might be able to be reduced by having exchanges form relationships with the large mining farms and the prominent hashrate marketplaces.

The hashrate marketplaces have been the source of most of the attacks. Centralized hashrate marketplaces have the ability to put limits on the total amount of hashrate that can be rented at once, and can even do things like Know Your Customer (KYC) for anyone attempting to buy substantial amounts of hashrate, and may reduce the risk of attack for smaller cryptocurrencies. At the very least, a hashrate marketplace may be able to warn exchanges when a bunch of hashrate is suddenly being pointed at a particular cryptocurrency.

A highly sophisticated attacker may be able to leverage Sybil attacks or even account compromises to circumvent these controls. And of course, the more controls that centralized marketplaces put in place, the more users will be driven towards decentralized solutions, where no such controls will be able to exist. So these controls will be at best a temporary solution, however a temporary solution may buy enough time for cryptocurrencies to migrate to better solutions.

Forming relationships with many of the larger mining farms is also likely to be highly beneficial. If nothing else, these relationships are likely to give insights into the current state of mining for various cryptocurrencies, and could give exchanges an idea for which cryptocurrencies might be more or less vulnerable. In terms of risk mitigation, I believe these relationships would have a larger than expected impact for the amount of effort required.

Automatically Halting Trading And Blacklisting Addresses

When a large reorg is detected on a cryptocurrency, trading should automatically be halted on that cryptocurrency, and if a double spend is detected the addresses involved in that double spend should be automatically blacklisted. This should happen across as many exchanges as possible, not just the exchanges impacted by the double spend attacks.

Though halting trading immediately won't help with the fact that money has been stolen, it does substantially reduce the number of options that an attacker has for handling the stolen money. Also, attackers can often predict price movements following large attacks and make large trades against those price movements. If trading is frozen, that source of profitability is reduced for potential attackers.

Blacklisting addresses has a similar effect: it reduces options for attackers. Shutting down more options for attackers means more opportunities to recover the money, and also means fewer attacks in the first place, even if there are ways to circumvent all of these controls.

We can say from experience that attackers often aren't that sophisticated, and often do make big mistakes. Even when there's nothing you can do against a theoretically perfect attacker, real attackers are far from perfect. Actively pursuing attackers and hoping that they make a critical mistake can be incredibly effective.

Scorched-Earth Counterattacks

There is a more advanced, and a more risky, option to handle double spend attacks, which is to launch a counter-attack. When an attacker mines a double spend on a cryptocurrency, the impacted exchange can potentially buy up a bunch of hashrate to extend the original chain, cementing the original transaction from the attacker.

The attacker can of course counter attack as well, responding to the extension of the original chain with an extension of the attack chain. The difficult thing here is that at every point in time, it makes sense for the exchange to spend more money extending the original chain, and it makes sense for the attacker to spend more money extending the attack chain. Even when the attacker and the exchange have both spent far more money than the theft is worth, it still makes sense for them to keep extending their respective chains in an attempt to get the money back.

Imagine that an attacker steals $50,000 from an exchange by spending $10,000 on proof of work. At this point, the attacker is +$40,000, and the exchange is $-50,000. The best move for the exchange here is to spend $10,000 themselves to restore the original chain as the longest chain, which means the attacker is now $-10,000, and the exchange is also -$10,000. If we let this game play out, we get the following:

Stage:	Attacker	Exchange
Stage 1a:	+$40,000	-$50,000
Stage 1b:	-$10,000	-$10,000
Stage 2a:	+$30,000	-$60,000
Stage 2b:	-$20,000	-$20,000
Stage 3a:	+$20,000	-$70,000
Stage 3b:	-$30,000	-$30,000
Stage 4a:	+$10,000	-$80,000
Stage 4b:	-$40,000	-$40,000
Stage 5a:	+$0	-$90,000
Stage 5b:	-$50,000	-$50,000
Stage 6a:	-$10,000	-$100,000
Stage 6b:	-$60,000	-$60,000

By the time that the attacker no longer stands to profit from the attack as a whole, the exchange has lost the same amount of money defending themselves that they would have lost if they had just let the attacker go in the first place. At no point in time is the exchange ever up, the exchange only stands to lose greater and greater amounts of money in the best case.

And, this game doesn't really have an ending state. At all points in time, it makes sense for each party to keep trying to get the original $50,000 back, because at each step you are spending a new $10,000 to recover $50,000. That is why this strategy is called 'scorched earth'—nobody wins, and lots of money gets destroyed.

The value to this strategy is that the exchange can, at least in theory, prevent the attacker from making money. If an attacker knows ahead of time that an exchange is willing to commit to a scorched earth strategy, then the attack doesn't make any sense and the exchange is unlikely to be attacked beyond the first few times.

There is another big complication with this strategy. The attacker has a big advantage in terms of preparation. An attacker can spend weeks or months preparing an attack, and an exchange needs to respond to the attack almost immediately. And, if the attacker is willing to engage the exchange like this, it may very well be the case that the attacker has some large advantage. For

example, if the attacker is using code that is more heavily optimized, the attacker may only be spending $5,000 each round, while the exchange is spending the full $10,000 each round. The exchange has no way to tell whether or not the attacker has an advantage in this situation either.

There could also be issues with this strategy if multiple exchanges attempt to perform it simultaneously. The exchanges may end up getting to a hashrate war with eachother instead of the attacker, and that could get extremely expensive depending on the budgets for each exchange.

And a final consideration for this strategy is that it could have massive collateral damage on the ecosystem. Many cryptocurrencies aren't really able to handle a large number of consecutive reorgs. Nodes may crash, other transactions may be lost or double spent in the middle of the war, and generally speaking users will be at much greater risk for the full duration of this scorched earth battle.

For all of the above reasons, I do not recommend that exchanges pursue this strategy to fight double spends.

Developer Arbitration

The final strategy I wanted to bring up was developer arbitration, because it is a strategy that has been successful for cryptocurrencies in the past. When a theft occurs, the developers can always launch a hardfork that returns the stolen coins. This introduces a very high level of centralization around the developers, and also the developers are imperfect human beings who could potentially be tricked into misreading an attack, and instead of returning stolen coins, the developers may end up taking legitimate coins from a user and giving them to an attacker.

Developers could also begin signing blocks. Once a block is signed by the developers, that block is permanent, and the transactions in the block cannot be double spent. This has been done by cryptocurrencies numerous times through history, but itself is very perilous. If the developer key gets stolen, all sorts of problems can happen. And, the fact that developers are effectively deciding which transactions are allowed on the network potentially puts them in the unforgiving sights of financial regulators.

Developers should be genuinely cautious of doing things like this, because if a developer does make the wrong decision when returning funds, signs the wrong block, or allows a known terrorist group to make a transaction, there could be serious legal repercussions. Especially now that there is a lot more regulator attention on this space, I don't recommend this avenue, even ignoring the usual centralization concerns.

Conclusion

As the cryptocurrency space continues to develop, we are going to continue seeing sophisticated attacks. In the next 6–12 months, most of these attacks are likely to be focused around double spends of cryptocurrencies with poor proof-of-work security, but increasingly the vulnerable decisions of developers are going to be exploited. Secure cryptocurrency design is difficult, and most cryptocurrencies and decentralized applications have not succeeded at ensuring their projects are secure.

That's being felt to the tune of millions of dollars in thefts today resulting from shared hardware hashrate attacks, but these attacks are only the first wave of high profile attacks that are going to being hitting the cryptocurrency community.

To prevent further losses, steps need to be taken in the short term to protect exchanges from shared hardware hashrate attacks. In some cases jumping to 24 hours of confirmations should be sufficient, and in others deposits should probably just be disabled until the cryptocurrency is able to fork to a more secure paradigm. In the long term, exchanges are going to need to be more conservative with their risk models and more proactive about diligence with the coins that they choose to list.

Special thanks to Ethan Heilman for review and feedback.

Thanks to Zach Herbert, Ken Carpenter, and Matthew Sevey.