Analyzing Soft and Hard Partitions of Global-Scale Blockchain Systems

Kevin Bruhwiler, Fayzah Alshammari, Farzad Habibi, Juncheng Fang,
Yinan Zhou, Abhishek Singh, Ahmad Showail*, and Faisal Nawab
Department of Computer Science, University of California, Irvine
**Department of Computer Engineering, Taibah University, Madinah, Saudi Arabia*
{*kbruhwil, fayzaha, habibif, junchf1, yinanz17, abhishas, nawabf*}*@uci.edu*
**ahmad.showail@kaust.edu.sa*

2022 IEEE International Conference on Blockchain (Blockchain) | 978-1-6654-6104-7/22/$31.00 ©2022 IEEE | DOI: 10.1109/BLOCKCHAIN55522.2022.00049

***Abstract*—Partitioning attacks on blockchain systems are a serious threat, with the potential to cause significant harm on the individual level and to the system as a whole. Deliberate partitions can be used by attackers to defraud merchants using crypto-currencies and enrich themselves, while natural disasters and wars could disrupt blockchain systems for months, potentially destroying them entirely. Unfortunately, the exact effects of these partitions at large scale are unknown and difficult to model, making it challenging to implement preventative measures or plan for partition recovery. In this work, we examine a variety of real-world global scale partitioning events to develop a framework for categorizing and analyzing different partition types. We use this framework, along with a new metric introduced to finely measure the global coherence of a given blockchain, to quantify the impacts of soft and hard network partitions at varying scales. Our goal is to lay the groundwork for the introduction of new blockchain architectures to minimize the harm caused during partitions and facilitate rapid and uncontroversial recovery from them.**

***Index Terms*—SimBlock, Bitcoin, blockchain, simulator, experimental analysis,metrics, attacks, partitioning, recovery**

I. Introduction

Partitioning attacks have been a known threat since the invention of crypto-currencies. Attackers could deliberately fork the chain into two or more separate, seemingly valid, parts and spend currency on each piece, effectively multiplying their money. Such attacks are known as double-spending and are the primary reason most blockchain systems work to minimize the possibility of forks. There are, however, a number of possible causes of partitions that simply cannot be prevented by the blockchain itself. These attacks include targeting the routing infrastructure of the internet to prevent or reroute traffic and damaging the physical structure of the internet, such as submarine fiber-optic cables, to prevent network traffic entirely.

Apostolaki et. al. [1] were among the first to quantify the threats of partition attacks on Bitcoin, the most widely used public blockchain. They identified the concentration of mining power within a small handful of routing prefixes and demonstrated how an attacker could divert traffic from these addresses using routing attacks, with emphasis that identical attacks of larger scale have been successfully performed in the past. The authors suggest a number of ways to mitigate this risk, which they ultimately combined into a tool named SABRE [2] which users can run on top of their blockchain system. SABRE does not entirely prevent routing attacks, but makes them considerably more difficult to perform. However, routing attacks are not the only way in which large-scale blockchains can become partitioned.

In 2021, Jyothi [3] explored the possibility that a solar superstorm could damage the undersea fiber-optic cables that connect the Internets of different continents, and considered the most likely ramifications of the damage. She concluded that such an event would likely cause major connectivity issues across the northern hemisphere and disconnect much of North America's internet from the remainder of the world for potentially weeks. It's difficult to precisely quantify the effect this would have on blockchain systems such as Bitcoin, though the effects are likely to be serious. There is also concern that undersea cables could be deliberately destroyed as acts of terrorism or in the event of war [4] or by natural disasters such as earthquakes [5]. These may have similar destructive effects on blockchains, although their location and impact is more difficult to predict in advance. We will attempt to more precisely quantify how these events will impact real-world systems.

The risks of network partitioning on large-scale blockchain systems are not well understood due to the difficulties in modeling them and the lack of global-scale partitioning events that have occurred since blockchains exploded in popularity. In this work, we modify a topographical blockchain simulation to properly quantify the effects of a global-scale network partition on both individual miners and on the blockchain as a whole. We hope to provide the groundwork for preventative measures to be taken to minimize, or even entirely eliminate, the harm that such partitions might cause in the future. To this end, we propose four research questions that will guide the remainder of this work:

1) What types of global-scale blockchain partitions are possible?
2) How can the effects of partitions on blockchains be meaningfully measured?
3) What will be the impact of partitions at different scales?
4) How can these impacts be used to guide a partition detection and mitigation?

978-1-6654-6104-7/22/$31.00 ©2022 IEEE
DOI 10.1109/Blockchain55522.2022.00049

The remainder of the paper is organized as follows. Section II identifies and summarizes related work in simulating and measuring blockchains. Section III describes the manner in which we categorize global-scale partitions and provides real-world examples for each category. Section IV covers a variety of metrics useful for determining the impact of partitions on blockchains. Section V details simulations of each category of partition and presents results from the simulations. Finally, Section summarizes conclusions from the simulations in Section V and answers the research questions posed above.

II. Related Work

A. Simulating Large-Scale Blockchains

Blockchain is a distributed tamper-proof ledger maintained by independent nodes across a peer-to-peer network. Conducting evaluation and analysis in large-scale distributed ledger systems is a challenging and expensive task due to their complexity [17]. For example, running a simple experiment that includes even a small number of nodes in a public blockchain system requires both information to be collected about that area of the network along with the costs associated with deploying the nodes to the blockchain. Furthermore, conducting the same experiment in a private network is still not an easy task as it requires preparing nodes and modifying network conditions and configurations. On this basis, the simulated distributed ledger system emerged as an effective solution to this problem and several Blockchain simulators allow simulating the real world in different scenarios at a low cost.

Paulavicius et al. [17] provides a systematic review of the state of art software blockchain simulators. They identify 27 simulators that have been developed between 2013 and 2020. They also include SimBlock [8], the simulator we extend in this work. They describe the architecture and functionality of these simulators by using the multi-layered abstraction blockchain paradigm. This paradigm consists of the network, consensus, data, execution, and application layers. We refer interested readers to reviews [20] for the description of the mentioned layers. SimBlock [8], along with only two other simulators, implements all the three layers of consensus, data, and network. The network layer is the most relevant layer to our work as demonstrates in Section III. Moreover, they compare the implemented available input and output parameters (features) of all of the simulators. SimBlock is also one of two simulators that have the highest number of implemented features. Finally, they conducted an experimental validation analysis on three simulators, SimBlock [8], Bitcoin-simulator [19], and BlockSim [18], using data collected from the Bitcoin network. SimBlock was able to simulate the Bitcoin network in 2016 with a high level of accuracy [17]. In the following Subsection II-B, we provide some background about SimBlock particularly.

B. SimBlock

SimBlock is an event-driven simulator that simulates a blockchain network at the block level [8]. It establishes connections between nodes by creating point-to-point links with specific network latency and bandwidth. It supports two popular blockchain consensus protocols, Proof of Work (PoW) and Proof of Stake (PoS). It propagates information by using the Compact Block Relay (CBR) Protocol [16]. SimBlock is written in Java and it provides options to configure the network parameters, including the connection distribution for each node in the system, the number of regions, the speed of connections between each region, and the number of nodes present in each region. It also provides tools to modify the simulation parameters, which include block size, mining difficulty (in the case of proof-of-work), consensus protocol, and length of the simulation. We will use a combination of these parameters to formulate metrics to study the effects of partitioning.

SimBlock was originally developed as a testbed for exploring different topologies of blockchain node connections in order to shorten block confirmation time by improving the block propagation rate. Consequently the simulation contains many useful concepts regarding the connections of nodes in the local regions and connections of regions to each other, many of which are essential for modeling partitions of varying magnitudes. We primarily modify the connections between regions, either limiting the bandwidth or increasing the latency of inter-regional communication to simulate the disruptions in question (see Section III).

C. Blockchain Metrics

Quantifying the stability, efficiency, and security of a blockchain system is a difficult task. A good set of metrics must take into account a number of factors regarding a public blockchain such as the growth rate of the chain, the geographical distribution of nodes, the communication delays among the nodes, and the possibility of malicious behavior. A number of attempts to formulate these metrics have been made, such as Seike et. al. [9] who estimate the efficiency of a blockchain by formalizing fork-rate bounds, however, their formulas assume all nodes are connected to each other and possess similar mining power, which makes it unusable for partition analysis. Zheng et. al. [10] developed a lightweight tool to run along blockchain nodes that measures a variety of interesting metrics, such as the number of transactions on the chain per CPU cycle, the rate at which new peers are discovered, and the speed with which transactions are propagated to a certain number of nodes. While these metrics don't capture partitions themselves, measuring how these metrics vary on either side of a partition may reveal interesting differences between the behavior of hard and soft partitions.

III. Partition Model

We model blockchain partitions as belonging to an intersection of a "disruption" type and a "cut" type, illustrated in Table I. The disruption type specifies the severity with

which a connection between two different regions is damaged and the cut type indicates the number of disruptions and the relationship between the region in question and the rest of the network. It's important to clarify that the disruption types are do not necessarily correspond directly to physical connections. For example, it's rare that a submarine fiber-optic cable is damaged but not completely destroyed. However, if multiple such cables connect two regions then a single one of them being destroyed will slow communication between those regions without eliminating it entirely, resulting in a soft disruption. There is also a possibility for overlap in cases where only a single connection links a region to the rest of the network, such that any link cut is also an isolation cut. In these cases we consider it to be an isolation cut, taking the partition type with the higher number if any conflict exists. In the following sections, we list real-world examples of each partition type.

TABLE I
A CONCEPTUAL MODEL OF PARTITION TYPES

	Soft Disruption	Hard Disruption
Link Cut	**TYPE 1** Some connections between a region and the rest of the network are slowed or have reduced capacity	**TYPE 2** Some connections between a region and the rest of the network are completely severed
Isolation Cut	**TYPE 3** All connections between a region and the rest of the network are slowed or have reduced capacity	**TYPE 4** All connections between a region and the rest of the network are completely severed

A. Soft Link Partitions

Soft link partitions occur when some connections between a specific region and the rest of the network are damaged but not completely destroyed. These partitions are most likely to occur for well-connected regions when some of their physical connections, such as submarine fiber-optic cables, are damaged. Soft link partitions occur frequently, but rarely with the severity necessary to impact connectivity. One such partition that did have a significant impact on connectivity was the 2006 earthquake off the coast of Taiwan, which throttled communication between the United States and China by more than 90% [12].

B. Hard Link Partitions

Hard link partitions differ from soft link partitions in that they completely eliminate direct communications between two regions without disconnecting either from the network. This means that blocks mined in one region can still be propagated to the other indirectly through other regions, unlike the slow direct propagation in soft link partitions. Jyothi's work on solar superstorms [3] provides a number of potential outcomes for mass partial internet disconnectivity, primarily of submarine cables above the northern 40^{th} parallel, which would have the effect of eliminating direct communications between North America and Northern Eurasia while allowing reduced communications between the continents in the southern hemisphere. This combination of soft and hard link partitions will have currently unpredictable results on block propagation. Historical real-world examples of these sorts of partitions are rare. It's unusual that a region is well enough connected to the network to survive one link being completely severed while having low enough redundancy that all links between two regions can be damaged simultaneously.

C. Soft Isolation Partitions

Soft isolation partitions occur when all communications between a certain region and the rest of the network are slowed, but not entirely eliminated. Such events are most likely to occur in regions with poor connectivity to the rest of the network, although they can occur as virtual partitions for well connected regions. For example, for more than two hours in 2019 nearly all mobile network traffic in Europe was diverted through China Telecom [11] as the result of an accidental BGP leak in a Swiss datacenter. This had the result of significantly slowing some types of internet traffic between Europe and the rest of the network and would appear as a soft isolation partition to observers, even though no internet infrastructure was damaged.

D. Hard Isolation Partitions

Hard isolation partitions occur periodically in poorly connected regions where all traffic is routed across one physical connection. For example, the recent volcanic eruption near Tonga severed the nation of 100,000 from the internet for more than a month. In 2018 Mauritania was disconnected from the internet for two days after a fishing boat accidentally severed a submarine cable, and Armenia was disconnected for several hours when an underground cable was accidentally damaged with a spade. In 2019, Yemen, with a population of nearly 30 million, was disconnected from the global internet for four days in the midst of a civil war when an air raid damaged the sole submarine cable linking it to the internet [13]. These examples all focus on relatively small regions being disconnected from the network, but the potential exists for much larger hard isolation partitions, such as Russia's "Sovereign Internet," [14] which can supposedly be disconnected from the world-wide internet. Additionally, routing attacks, such as the one described in Section III-C, can create virtual hard isolation partitions by rerouting network traffic and declining to propagate it.

IV. PARTITION METRICS

Quantifying the impact of various partitions on blockchain systems is non-trivial as many of the metrics currently used to measure blockchain are either very slow to be impacted or are simply unaffected by partitions, such as the time to mine a new block or the rate at which transactions are confirmed. We identify two existing metrics which are valuable and introduce a third novel measure to compare and contrast the severity

of different partition types at different scales. These metrics are: the number of competing blocks of the same height, the average propagation time of blocks in the network, and the proportion of nodes who agree on the newest, most widely accepted block, introduced here as "percentage agreement."

A. Competing Blocks

The number of competing blocks at different heights is a useful metric to capture the general "badness" of a blockchain system, as the primary issue caused by network partitions is double spending, an attack enabled by having multiple, seemingly valid blocks at the same height. It can also be used as a proxy to identify the number of forks of the chain at a given time and additionally provides some temporal information, allowing us to see at what point in the chain conflicts started to appear and disappear. However, competing blocks are a coarse-grained metric, providing a concrete measure of when things are going wrong but providing little insight into why exactly the forks are forming and under what circumstances they disappear.

In the event of a hard isolation partition (described in Section III-D) the region being isolated may only contain a handful of miners who initially produce a very small number of competing blocks relative to the rest of the network. This will appear as a small fork in the competing blocks metric, likely appearing long after the partition has occurred, suggesting little to worry about. However, within that region, all systems interacting with the blockchain have slowed to a crawl or entirely frozen due to the lack of new blocks appearing. Additionally, all transactions on the few new blocks that are mined are vulnerable to double spending attacks by users who can spend on both sides of the partitions, either due to their ability to physically travel across the partition or because their wallet is shared by users on either side. These serious impacts will not be properly reflected by the competing blocks metric, and consequently other metrics must be considered.

B. Propagation Time

Propagation time is a measure of the average length of time it takes for a block to be propagated to the rest of the nodes in the network. Outside of partition studies, it's of interest when compared to the mining time of new blocks; should blocks be mined faster than they can be propagated then many forks will form and the stability in the network will fall. Here, it is used to measure the severity of semi-partitions in which the network remains completely connected but is hampered by slow or missing connections between regions. Propagation time provides a much more fine-grained measurement of badness than competing blocks due to the fact that it can be measured at arbitrary points in time (not only when blocks are mined) showing how the functioning of the system changes as the partition subsists or disappears. Slow propagation time can also result in wasted work on a blockchain, as miners spend longer mining on old nodes not yet knowing that a new node has been mined.

However, propagation time suffers from many of the same issues that competing blocks does when it comes to hard isolation partitions. Blocks are not propagated across hard isolation partitions and consequently they may show a decrease in propagation time as blocks have fewer total nodes to reach in their partition. In semi-partitions with competing forks, block propagation time suffers a similar issue in which a block doesn't propagate to distant nodes because they have a different version of the chain for which that block is invalid, resulting in apparently short propagation time.

Propagation time is also unique among the metrics we are using because it does not require global knowledge; nodes can independently estimate propagation time by determining how long it took themselves to receive a block and sharing that information with their neighbors. This means that insights gained by looking at block propagation time could potentially be used by nodes to guide their own behavior. Section V-D gives an example of how propagation time may be useful to individual nodes.

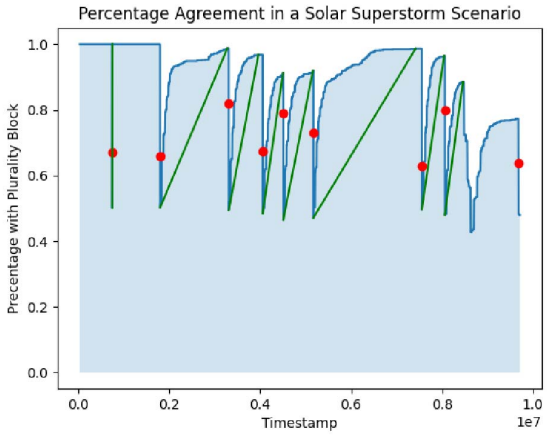

Fig. 1. An example of the percentage agreement overtime during a simplified solar superstorm scenario. Red dots indicate when a new block is mined, and green lines show how the agreement increases. The partition begins at 200 seconds.

C. Percentage Agreement

Finally, we introduce percentage agreement as a measure of the total stability of the network. Percentage agreement identifies the most recent block accepted by a majority of nodes in the network and computes the proportion of nodes that have accepted it at any given time. This provides two key pieces of information: in the event of hard isolation partitions it shows a drop in the percentage agreement proportional to the size of the partitioned region, and in the event of a semi-partition the rate at which the percentage agreement falls and rises as new blocks are mined can be used to precisely quantify how effectively new blocks are being propagated. Agreement increase can be visualized as the slope of the green lines in Figure 1, connecting the points of lowest and highest agreement after a new block has been mined. The average

agreement, a metric used extensively in the evaluation in this work, is measured as the area underneath the curve in Figure 1 divided by the total running time of the simulation.

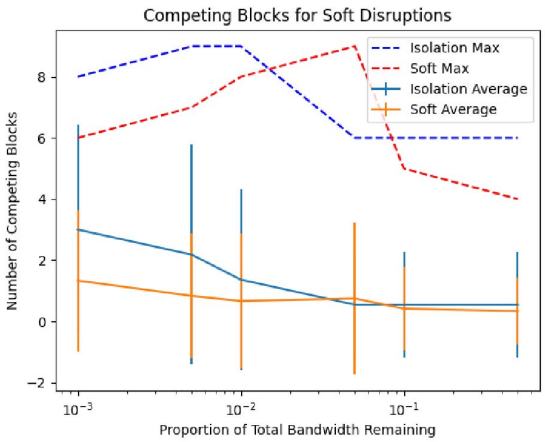

Fig. 2. The average number of competing blocks during soft disruptions, average of 100 simulation runs. In both link and isolation cuts the number of competing blocks is highly variable, although isolation partitions are more likely to generate competing blocks as the partition worsens. Note that the x-axis is shown in log scale.

V. Partition Simulations

We identify a number of different partition scenarios to simulate based on real-world partitioning events from each category in Table I. We are primarily interested in the differences between different partition types at different scales and in determining their potential impact. All simulations are performed using mining and propagation data collected from the Bitcoin network. However, it should be considered that some blockchains may be impacted in different ways. For example, chains with a rapid rate of block generation will suffer much more from reduced propagation time than chains with slow generation such as Bitcoin. We also perform several experiments using a modified version of SimBlock on artificial partitioning scenarios to determine the precise breakpoints at which soft disruptions begin to cause issues and identify how some local metrics may be useful in identifying when a hard partition has begun.

A. Soft Link Partition Scenarios

To test the impact of soft link partitions we use the real-world example of the 2006 earthquake near Taiwan, which reduced internet communication between North America and China by 90% [12]. Table II shows the impact, averaged over 10 simulation runs, of varying the reduction in bandwidth between the two regions on the percentage agreement. The agreement increase, the rate at which the agreement of the nodes changes, is shown for the simulation both before and during the partition along with its standard deviation. We provide the increase

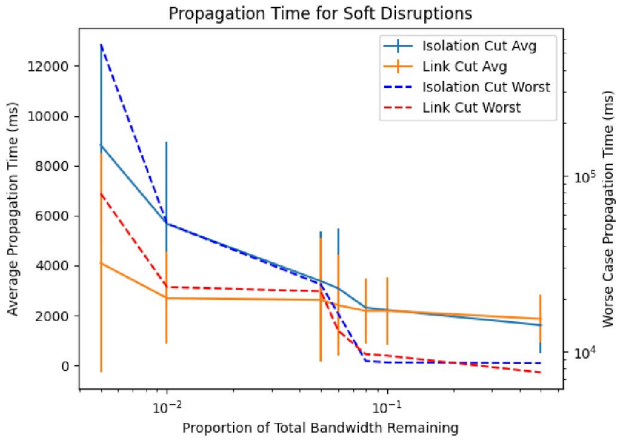

Fig. 3. The average and worst-case propagation time during soft disruptions. In both link and isolation cuts the propagation time does not start to be significantly impacted until the bandwidth drops to 5% of its initial capacity. Note both the y-axis and right x-axis are shown in log scale.

TABLE II
Agreement Variation During Soft Link Partitions Shown With Standard Deviation

Soft Link Bandwidth Throttling	Agreement Increase Before Partition (%/sec)	Agreement Increase After Partition (%/sec)	Average Agreement (%)
0.5	15.96 ± 3.29	12.06 ± 3.01	99.46
0.1	16.52 ± 2.53	14.80 ± 1.53	98.79
0.05	13.68 ± 0.42	9.91 ± 4.14	99.02
0.01	11.23 ± 2.25	7.15 ± 6.11	98.83

before the partition because the semi-random way in which nodes are connected means that propagation time can vary somewhat between simulation runs.

Table II shows that soft link partitions between two major regions do not have a significant impact on the coherence of the blockchain system as a whole. The average agreement of the system dips very slightly only when the soft disruption nearly entirely eliminates communication across a connection. What is impacted more significantly is the rate of agreement increase, which drops significantly as the disruption worsens. The average agreement shows that this low rate of agreement increase does not significantly impact the coherence of the system, however, slower propagation will mean that transaction verification times increase and, at least for proof-of-work blockchains, some mining time will be wasted for miners who take longer to receive the new block.

Figure 2 shows the average and maximum number of competing blocks during both soft link and soft isolation partitions, averaged over 100 simulation runs. The maximum number of competing blocks can be considered a ”worst case” situation, which is valuable when doing risk analysis for blockchain forks. At each partition scale, multiple forks occur in some simulation runs while being absent in others. Figure 2

TABLE III
AGREEMENT VARIATION DURING SOFT ISOLATION PARTITIONS SHOWN WITH STANDARD DEVIATION

Soft Isolation Bandwidth Throttling	Agreement Increase Before Partition (%/sec)	Agreement Increase After Partition (%/sec)	Average Agreement (%)
0.5	11.79 ± 3.53	11.48 ± 2.75	99.48
0.1	11.16 ± 2.64	7.65 ± 3.07	99.18
0.05	12.74 ± 1.05	4.64 ± 2.79	99.14
0.01	12.94 ± 2.35	1.13 ± 0.58	97.3

TABLE IV
AGREEMENT VARIATION DURING HARD LINK PARTITIONS SHOWN WITH STANDARD DEVIATION

Number of Hard Link Cuts	Agreement Increase Before Partition (%/sec)	Agreement Increase After Partition (%/sec)	Average Agreement (%)
1	14.32 ± 5.83	7.48 ± 1.90	94.03
2	15.82 ± 2.61	5.34 ± 2.92	93.02
3	14.06 ± 1.05	4.03 ± 1.15	90.14

primarily shows that more severe soft link partitions are, on average, not likely to create significantly more competing blocks. However, the variation in the number of competing blocks increases substantially as the severity of the partition increases. Interpreted another way, soft link partitions do not appear likely to damage the stability of the network but they are likely to waste work in mining blocks on older parts of the chain in somewhat unpredictable ways (this conclusion is supported by Table II). Also interesting to observe is that the "worst case" situation is only loosely tied to the severity of the partition. Significant numbers of wasted blocks are possible during less severe partitions.

Figure 3 shows the average and worst-case block propagation time for both soft link and soft isolation partitions, discussed in further detail in Section V-B. Worst case block propagation quantifies the time it took the slowest block to propagate across the network. We display it separately because the region where a block is mined has a significant effect on how that block propagates during a partition. We can see that the break points occur when the bandwidth is reduced to about 5% of its initial capacity, which aligns with the change in agreement increase seen above. From these two metrics we can conclude that, although soft link partitions have the potential to reduce propagation time, they only have a significant impact if the reduction in communication is so severe as to nearly eliminate it entirely. The Taiwanese earthquake used as an example would have had little impact on the functioning of a global blockchain.

B. Soft Isolation Partition Scenarios

Although there are few real-world examples of large-scale soft isolation partitions, we use the accidental BGP leak in a Swiss datacenter [11] which rerouted all mobile traffic from Europe through China as inspiration and measure the impact of soft isolation partitions of various scales on Europe. Table III examines the change in agreement, and in the rate of agreement increase, as the intensity in the soft disruption increases. Section V-A provides further context as to why we use these metrics.

Table III shows that, similar to soft link partitions, soft isolation partitions do not significantly impact the stability of the system, as the global percentage agreement only starts to meaningfully fall after the disruption is so severe the communication is nearly eliminated. However, it does have a much more dramatic impact on the rate at which agreement increases, meaning that it takes the system much longer to learn about new blocks as they are mined. In a blockchain protocol where new blocks are mined more frequently this may result in many more forks. It should still be observed that the rate of agreement increase only falls significantly when the soft disruption leaves less than $1/10^{\text{th}}$ of the initial bandwidth, which is a severe disruption, and one for which no real-world scenarios have been observed.

Figure 2 provides some insight into the impact of reduced block propagation time on the amount of work wasted and the number of forks. While the standard deviation of the number of competing blocks is extremely high (many runs of the simulation produce no competing blocks, even during severe soft isolation partitions) there is a noticeable increase in the number of competing blocks at the same 5-10% mark observed in Table III. This is most likely a result of blocks being mined by nodes who have not yet received the newest blocks due to the slower propagation. The relatively small number of competing blocks produced also lines up with the results of Table III which shows that, on average, the system maintains a high level of stability.

Figure 3 also shows the average and worst case propagation times for soft isolation partitions. Its results match those of the percentage agreement, showing that severe impacts are only seen once the soft disruption has reduced the bandwidth to $< 5\%$. At $< 5\%$ available bandwidth, the worst case propagation time has increased to 20 seconds, nearly twice the block generation rate of Ethereum [21].

C. Hard Link Partition Scenarios

Hard link partitions cannot be evaluated in quite the same way as soft link partitions due to their "all-or-nothing" nature. However, we can experiment by varying the number of hard link cuts to see their relative impact. To do so we randomly cut a number of connections between a single region and the rest of the network and examine their impact in Table IV. The results are averaged over 10 simulation runs. The average agreement can be interpreted as the total stability of the blockchain system, while the agreement increase is the rate at which the system recovers stability after a new block is

TABLE V
THE AVERAGE AND WORST-CASE BLOCK PROPAGATION TIMES RELATIVE TO THE PROPORTION OF NODES OUTSIDE OF A HARD ISOLATION PARTITION SHOWN WITH STANDARD DEVIATION

Proportion of Nodes in Hard Isolation Partition	Average Block Propagation Time (ms)	Worst Case Block Propagation Time (ms)
40%	2361.20 ± 702.26	5558
30%	1933.09 ± 1070.69	6168
20%	1710.61 ± 812.57	5772
10%	1169.79 ± 555.31	3877

mined.

Table IV illustrates that hard link partitions have a much greater impact on system stability than soft disruptions. A single hard link partition will ensure that, on average, approximately 6% of the nodes in the network will not share the most recent block. The effects become more pronounced as the number of cuts grows. Additionally, the rate at which the agreement increases is also significantly reduced due to the fact that fewer paths exist to propagate blocks between regions. Collectively, these impacts will mean significantly increased transaction confirmation time and a significant amount of wasted computation in the case of proof-of-work chains.

D. Hard Isolation Partition Scenarios

Hard isolation partitions are generally trivial to quantify; the percentage agreement drops to the size of the largest partition, the agreement increase remains unchanged, and the number of forks eventually becomes equivalent to the number of partitions. Of interesting note, however, is an observation made in Table V regarding the relationship between hard partitions and block propagation time *only for nodes outside of the partition*. Intuitively, it is unclear whether a hard partition would cause the propagation time to decrease due to the smaller number of nodes that a block must be propagated to or increase due to the reduced connectivity of the network.

Table V shows that reduced connectivity dominates the effect on block propagation time, with larger hard isolation partitions significantly increasing the time it takes to propagate a block through the network. This is a valuable insight because estimating the average propagation time for a block does not require global knowledge, unlike the number of forks or the global agreement. This could be used by individual nodes to identify when the network may be suffering from a hard isolation partition and allow them to take steps to protect themselves from future forks. It can also be observed that, for each additional 10% of the network that is partitioned, the average block propagation time increases by approximated 300-500ms, enabling a rough estimation of the partition size from the change in propagation time alone.

E. Solar Superstorm Scenario

Finally, we set up a simplified version of the solar superstorm scenario, detailed in [3], in which all submarine fiber-optic cables north of the northern 40° parallel are disabled, and all remaining cables north of the tropics are reduced randomly to between 1% and 10% capacity. Though this is unlikely to be a very accurate representation of what would really happen during a solar superstorm, it does provide an interesting hypothetical situation in which multiple regions are suffering from a mix of soft and hard link partitions.

Figure 1 plots the agreement over time, with the partition beginning at 200 seconds. Red dots identify when new blocks are mined and the slope of the green lines, drawn from the point of lowest agreement after a block is mined to the point of highest agreement before the next block, is the rate at which the agreement increases. It's interesting to observe that, although the partitions severely slow the rate at which agreement rises, the system is generally able to reach $> 90\%$ agreement before a new block is mined. In some cases, where a new block is mined relatively quickly, the agreement doesn't quite reach 90%. This makes it likely that a modest proportion of nodes may, at some point, be two blocks behind. However, the consistent propagation of blocks means that this scenario is not likely to completely destroy the system.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we have examined real-world global scale internet partitioning events and used them to determine four different categories into which partitions fall. We introduced percentage agreement as a measure of the total stability of a blockchain system and examined how the rate at which it changes can be used to quantify the severity of a variety of different partition types. We also examined the average propagation time of blocks in the network and identified a soft disruption tipping point of 5% original bandwidth at which the block propagation time begins to rapidly increase. We also demonstrated how average propagation time can be used locally by individual nodes to estimate the existence and size of hard isolation partitions. Below we answer the four research questions posed at the beginning of this work.

1) *What types of global-scale blockchain partitions are possible?*

While hard partitions are possible, global network partitions do not necessarily completely isolate a region, they may only slow communications between a region and the rest of the network, or completely sever communications between two regions. We identify a distinction between link cuts, which only damage communications between two regions, and isolation cuts, which affect all communications between a region and the network. We also identify a difference between soft disruptions, which slow or throttle connections without completely eliminating communication, and hard disruptions, which prevent communication entirely.

2) *How can the effects of partitions on blockchains be meaningfully measured?*

The effects of partitions on blockchains are complex, but we identify three different metrics to quantify the damage done to the network. First, we look at the number of wasted blocks,

blocks mined on forks that are eventually abandoned. This metric is a good coarse-grained measurement of how well the blockchain is performing, however, it does not have the resolution to effectively compare partitions. We also examine how block propagation time can be used to compare the relative impact of soft-disruptions at different scales and use it to identify a tipping-point at which soft-disruptions begin to have a serious impact. Finally, we introduce percentage agreement as a method for measuring the total cohesiveness of a blockchain and determining how it changes over time.

3) What will be the impact of partitions at different scales?

We simulated partitions of each category at a number of scales and attempted to identify tipping points at which impacts become potentially disastrous. For soft disruptions we identify a clear inflection point at approximately 5% original bandwidth where block propagation times begin to increase exponentially. We also observe that soft disruptions, even extreme ones, have little impact on the total cohesion of the network, while hard disruptions quickly begin to seriously reduce the cohesion of the network. We also observe the somewhat non-intuitive fact that hard isolation partitions reduce block propagation time in approximately linear relation to their size.

4) How can these impacts be used to guide a partition detection and mitigation?

We observe that block propagation time, a metric which can be estimated locally by nodes, could be useful in estimating when a partition is occurring and how severe it is, allowing nodes to take protective measures. We also discuss how reduced block propagation times may be considerably more impactful on blockchains with rapid block generation, providing a basis for slowing block generation in the event of a partition.

As future work we intend to continue evaluating partitions on various blockchains, especially those which use different proof-of protocols. Our simulations were purely for proof-of-work protocols, and the impact on blockchains which use proof-of-stake, collective signing, or other more communication dependent protocols are likely to endure significantly different impacts. We also plan to model some blockchain protocols with block generation rates that vary dynamically when a potential partition is detected.

We will also pursue work in developing a partition tolerant protocol which can survive and recover from severe network partitions, currently termed PeloPartition [22]. Such a protocol depends not only on developing rules to recover fairly and efficiently from a partition such that minimal work is wasted and value lost but also on protecting users during a partition from attackers without significantly limiting the functioning of the blockchain.

VII. Acknowledgements

This research is supported in part by the NSF under grant CNS-1815212.

References

[1] Apostolaki, Maria, Aviv Zohar, and Laurent Vanbever. "Hijacking bitcoin: Routing attacks on cryptocurrencies." 2017 IEEE Symposium on Security and Privacy (SP). IEEE, 2017.
[2] Apostolaki, Maria, et al. "SABRE: Protecting bitcoin against routing attacks." arXiv preprint arXiv:1808.06254 (2018).
[3] Jyothi, Sangeetha Abdu. "Solar superstorms: planning for an internet apocalypse." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021.
[4] Carter, Lionel, et al. "Insights into submarine geohazards from breaks in subsea telecommunication cables." Oceanography 27.2 (2014): 58-67.
[5] Pope, Ed L., Peter J. Talling, and Lionel Carter. "Which earthquakes trigger damaging submarine mass movements: Insights from a global record of submarine cable breaks?." Marine Geology 384 (2017): 131-146.
[6] Alharby, Maher, and Aad van Moorsel. "Blocksim: a simulation framework for blockchain systems." ACM SIGMETRICS Performance Evaluation Review 46.3 (2019): 135-138.
[7] Memon, Raheel Ahmed, Jian Ping Li, and Junaid Ahmed. "Simulation model for blockchain systems using queuing theory." Electronics 8.2 (2019): 234.
[8] Aoki, Yusuke, et al. "Simblock: A blockchain network simulator." IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS). IEEE, 2019.
[9] H. Seike, Y. Aoki and N. Koshizuka, "Fork Rate-Based Analysis of the Longest Chain Growth Time Interval of a PoW Blockchain," 2019 IEEE International Conference on Blockchain (Blockchain), 2019, pp. 253-260.
[10] P. Zheng, Z. Zheng, X. Luo, X. Chen and X. Liu, "A Detailed and Real-Time Performance Monitoring Framework for Blockchain Systems," 2018 IEEE/ACM 40th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP), 2018, pp. 134-143.
[11] A. Improta and L. Sani, BGP Route Lear Incident Overview, (Apr. 16, 2020), https://www.catchpoint.com/blog/bgp-route-leak
[12] Winston Qiu, Submarine Cables Cut after Taiwan Earthquake in Dec 2006, SUBMARINE CABLE NETWORKS, (Mar. 19, 2011), https://www.submarinenetworks.com/news/cables-cut-after- taiwan-earthquake-2006
[13] Khalili, Joel. "Internet Access Hangs by a Thread for Hundreds of Millions." TechRadar, TechRadar Pro, 29 Feb. 2020, https://www.techradar.com/news/internet-access-hangs-by-a-thread-for-hundreds-of-millions.
[14] Person. "Russia Disconnects from Internet in Tests as It Bolsters Security - RBC Daily." Reuters, Thomson Reuters, 22 July 2021, https://www.reuters.com/technology/russia-disconnected-global-internet-tests-rbc-daily-2021-07-22/.
[15] Quiroz-Gutierrez, Marco. "Crypto Is Fully Banned in China and 8 Other Countries." Fortune, Fortune, 5 Jan. 2022, https://fortune.com/2022/01/04/crypto-banned-china-other-countries/.
[16] Nagayama, Ryunosuke, et al. "Identifying Impacts of Protocol and Internet Development on the Bitcoin Network." ArXiv:1912.05208 [Cs], June 2020. arXiv.org, http://arxiv.org/abs/1912.05208.
[17] Paulavic˘ius, S. Grigaitis, and E. Filatovas, 'A Systematic Review and Empirical Analysis of Blockchain Simulators', IEEE Access, vol. 9, pp. 38010–38028, 2021, doi:10.1109/ACCESS.2021.3063324.
[18] M. Alharby and A. van Moorsel, "Blocksim: A simulation framework for blockchain systems," ACM SIGMETRICS Perform. Eval. Rev., vol. 46, no. 3, pp. 135–138, 2019.
[19] A. Gervais and V. Glykantzis. (2016). Bitcoin and Blockchain Simulator. [Online]. Available: https://github.com/arthurgervais/Bitcoin-Simulator [73] Y. Aoki, K. Otsuki, T. Kaneko, R. Banno, and K. Shudo. (2019).
[20] C. Fan, S. Ghaemi, H. Khazaei, and P. Musilek, "Performance evaluation of blockchain systems: A systematic survey," IEEE Access, vol. 8, pp. 126927–126950, 2020.
[21] Buterin, Vitalik. "Ethereum: A next-generation smart contract and decentralized application platform." (2014).
[22] Fang, Juncheng, et. al. "PeloPartition: Improving Blockchain Resilience to Network Partitioning", IEEE, Blockchain-2022