

Secure Incentivization for Decentralized Content Delivery

Prateesh Goyal¹, Ravi Netravali², Mohammad Alizadeh¹, Hari Balakrishnan¹
MIT CSAIL¹, UCLA²

Abstract

Prior research has proposed using peer-to-peer (P2P) content delivery to serve Internet video at lower costs. Yet, such methods have not witnessed widespread adoption. An important challenge is incentivization: what tangible benefits does P2P content delivery offer users who bring resources to the table? In this paper, we ask whether monetary incentives can help attract peers in P2P content delivery systems. We first propose Gringotts, a system to enable secure monetary incentives for P2P content delivery systems. Gringotts provides a novel Proof of Delivery mechanism that allows content providers to verify correct delivery of their files, and shows how to use cryptocurrency to pay peers while guarding against liars and Sybil attacks. We then present results from an 876-person professional survey we commissioned to understand users' willingness to participate in Gringotts, and what challenges remain. Our survey revealed that 51% would participate for suitable financial incentives, and motivated the need for alternate payment forms, device security, and peer anonymity.

1 Introduction

Video streams now constitute over 70% of global Internet traffic [6, 27]. Most video is delivered to users today via Content Distribution Networks (CDNs) like Akamai and CloudFlare. Although CDN demand has nearly doubled since 2016 [6], they are too expensive for many content providers to use [25].

To combat these high prices, there have been many proposals to have peers cache and stream videos to each other, either as supplements to existing CDNs [13, 16, 34–37], or as decentralized peer to peer (P2P) systems [21, 28]. For example, peers within an Internet Service Provider (ISP) in a city could stream videos directly to one another, avoiding expensive Internet paths. By serving content using bandwidth and storage resources that would otherwise go unused, such systems can significantly reduce costs for content providers and CDNs.

Unfortunately, these proposals have not seen significant adoption in practice, primarily due to a lack of sufficient peer participation [4, 37]. A key reason is that traditional P2P systems have relied on tit-for-tat incentivization strategies [8, 23, 28], whereby clients (i.e., content consumers) must also contribute resources and act as peers (i.e., content distributors). Coupling clients and peers in this way limits adoption: it restricts content access for clients who do not wish (or are unable to) contribute resources, and fails to incentivize peers who wish to contribute resources without downloading content.

We believe that the steady increases in client-side uplink bandwidth and storage capacity [29] motivate revisiting P2P content delivery, but with a focus on new mechanisms to securely incentivize peers. In particular, we consider using mon-

etary payments as an incentive mechanism. Unlike tit-for-tat strategies, monetary incentives can decouple clients from peers in content delivery systems to increase potential participation.

Performing monetary payments securely is challenging. Such incentives can lead to numerous financially motivated attacks from any combination of malicious clients, peers, and content providers. For example, a content provider can refuse to pay a peer for content that was already delivered. Alternatively, an adversary can create fake colluding clients and peers to earn money without delivering any content (i.e., a Sybil attack). These problems persist regardless of whether payment is done using a central authority that everyone trusts (e.g., a bank or the CDN) or a decentralized approach. For instance, a central payment authority must reliably track all data transfers to determine payments, but existing tracking techniques rely on client- and peer-generated logs, making them vulnerable to collusion attacks [1]. Thus, the key challenge is to assure content providers that delivery is happening properly, despite the fact that they are not directly involved in each data transfer.

Our first contribution is a lightweight solution to this problem in which content providers, clients, and peers collectively produce a *Proof of Delivery Chain (PoDC)* that serves as a proof for the delivery of a file from a set of peers to a client (§2). PoDCs are unforgeable and tamper-proof: neither the peers nor the content providers can manipulate them to affect payments.

Second, we present the design of Gringotts (§2), a system that applies PoDC to decentralized P2P content delivery. Payments in Gringotts are made using a cryptocurrency. To distribute files, content providers advertise smart contracts for each file to serve. Smart contracts embed proof of delivery instructions, as well as information about where payments will come from (i.e., the content provider's cryptocurrency account). Peers are then guaranteed payment for delivered content by broadcasting PoDCs on a blockchain; to ensure sufficient scalability, Gringotts uses probabilistic payments to limit blockchain transactions. Gringotts is also robust to various forms of collusion and Sybil attacks with fake clients or peers.

Our third contribution is an 876-person consumer survey that we commissioned to understand users' willingness to participate in Gringotts (§4). To our knowledge, these are the first published results on this question. Our key finding is that 51% of respondents would participate. Those who would not participate were primarily concerned with device security, content liability, and impacts on device performance. Our survey also revealed that 27% of users are already willing to accept cryptocurrency payments (with 40% unsure due to lack of familiarity). Thus, our findings motivate that, in order to fully realize P2P content delivery, techniques for alternate payment forms, device security, and anonymity for peers must be developed.

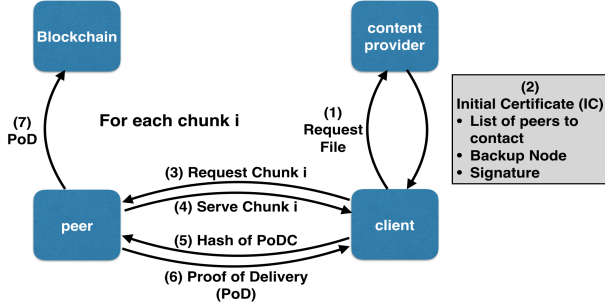


Figure 1: Downloading a file with Gringotts.

2 Secure and Practical Decentralized Incentivization

In this section, we describe how Gringotts handles content delivery and payments. We discuss solutions to potential security vulnerabilities and discuss practical overhead considerations.

2.1 Overview

Gringotts includes three major entities: a content provider that generates content (e.g., videos), a client who requests that content, and a peer who serves that content. To start, the content provider creates a *Smart Contract* for each content file that they would like to distribute. The Smart Contract includes information about payments, describing how much peers will earn by serving this file, where payments will come from (i.e., the content provider’s cryptocurrency account), and the rules that peers must follow to prove that they served a file for payment (i.e., *Proof of Delivery (PoD)*). Files are broken into chunks, and each chunk is placed on multiple peers who agree to the content provider’s terms for that file.

Figure 1 illustrates how Gringotts handles a client request for a single file. Client requests are initially forwarded to the content provider, who responds with an *Initial Certificate (IC)* that provides a guide on how to download all of the file’s constituent chunks. The IC includes a list specifying the peer (identified by IP address and public key) to download each chunk from, along with the address of a Backup Node, which is a trusted server (e.g., a traditional CDN) to contact in the event that a peer is unreachable. Backup Nodes operate identically to normal peers, generating a PoD for each request that they serve.¹ The content provider signs each IC with its private key to prevent forgery.

Upon receiving an IC, the client begins to download chunks from the listed peers. Clients maintain a sequential chain of PoD entries for each chunk they download, called the PoD chain (PoDC). The first element in the chain is the IC served by the content provider. After each successful chunk download, the client sends the peer a hash of the current PoDC. The peer then generates a new PoD (for the chunk they served) by signing the hash with its private key. The peer commits this PoD to the Blockchain, and also sends it to the client, which adds it to the PoDC. Thus, at the end of the file download, the

¹Backup Nodes can be deployed on programmable CDNs like Amazon CloudFront [2], which can execute arbitrary computations on each request.

Blockchain contains the entire PoDC for the file download. The PoDC effectively acts as a Blockchain, in that every link in the chain can be verified by a third party to ensure that the Smart Contract is not violated.

For simplicity, we described the client’s downloading of chunks to be serial. This ensures sequential PoDs in the PoDC, which is necessary for third party verification. However, for improved performance, chunks for a file can be downloaded in parallel streams, such that each stream generates an independent PoDC comprised of chunks that are downloaded serially. We also note that Gringotts can run atop existing blockchains using standard cryptocurrencies.

2.2 Overheads

Content Provider Overheads. With Gringotts, every client file request is first forwarded to the content provider which generates an IC for the download. This overhead mimics that of a video download with a traditional CDN. Clients traditionally begin a video streaming session by first downloading a Manifest file directly from the content provider; video content is then downloaded directly from CDNs [17]. In fact, with Gringotts, content providers can simply append ICs directly to Manifest files. Manifest files can be further modified to include hash values for each chunk in the file, allowing clients to verify the integrity of the data they receive.

Blockchain Overheads. In the design presented in §2.1, each chunk download is recorded on the Blockchain by the corresponding peer. This transactional overhead may overwhelm the Blockchain. For example, 50,000 YouTube videos are watched every second, but existing Blockchains like Ethereum [33] can only support 25 transactions per second.

To limit the overhead on the Blockchain, we propose using probabilistic payments. A simple approach is for content providers to grant payments only for a PoD which is divisible by a number N (specified in the Smart Contract). Peers would only submit a PoD to the Blockchain if this condition is met, reducing the number of Blockchain transactions by a factor of N .

One challenge with probabilistic payments is that the PoDC will no longer be directly recoverable from the Blockchain, since not all PoD entries will be present. However, the PoDC is necessary to verify that a PoD is valid, i.e., that it corresponds to a file delivery that adhered to the rules specified by the Smart Contract. To overcome this, peers can request the current PoDC from the client each time they generate a payment-eligible PoD. The peer can then include the PoDC in its Blockchain entry. It is important to note that PoDC transactions only occur once every N chunk downloads. Further, we expect the size of each PoD to be several bytes, implying PoDC sizes under a few KB.

Setting the value of N for probabilistic payments entails a tradeoff: a large N implies low payment frequency, while a low N limits the savings on Blockchain transaction costs. Assume that N is selected such that an active peer gets paid (on average) 10 times per day, or 300 times per month. We

make two observations. First, given that existing Blockchains like Ethereum can handle 2.2 million transactions per day, Gringotts would be able to support 220,000 active peers at any time. Second, payment frequencies are relatively stable. Specifically, the number of payments per month will be Binomial distributed, with a mean of 300 and a standard deviation of 17.3. Of course, the appropriate value of N will change over time, as new Blockchains are created (e.g., Algorand [10] claims to support 100x more throughput than Ethereum).

2.3 Detecting and Thwarting Attacks

Using monetary incentivization in decentralized settings has inherent security risks, as any involved entity (content provider, peer, or client) can be malicious, either independently or in collusion. We primarily focus on financially-motivated attacks, but also discuss other common attacks (e.g., DoS attacks).

2.3.1 Single-entity Attacks

Malicious peer. To get paid, a peer must generate a PoD that is payment-eligible according to the Smart Contract. However, a malicious peer could attempt to generate a PoD *without* serving content to a client. Preventing such an attack in Gringotts is relatively straightforward. Recall that a peer generates a PoD using a hash of the PoDC that is provided by a client. If a malicious peer does not serve a chunk, the client can simply refuse to provide a PoDC hash (and subsequently download the chunk from the Backup Node). Further, since peers only learn the PoDC hash after serving a chunk, they cannot selectively choose to only service requests that will lead to a payment.

Malicious content provider. A content provider's goal is to maximize content distribution while minimizing cost. Since a content provider does not know a peer's private key (which is used to generate a PoD) in advance, it cannot predict what a PoD will look like with a new IC. However, a malicious content provider could generate ICs which have resulted in no payments to a peer during past downloads. To prevent this, ICs can be augmented to include a nonce, which serves as a unique identifier for each IC. Each content provider's nonces must be monotonically increasing (e.g., timestamps), and clients must include the appropriate nonce in each chunk request that they send a peer (i.e., Step 3 in Figure 1). In this way, peers can verify that incoming chunk requests are not intentionally designed to prevent payment by reusing a nonce value.

There is one complication with having peers verify that nonces are monotonically increasing. Given the decentralized nature of Gringotts, client network latencies can create race conditions for nonce verification at peers. For example, consider a scenario in which two clients simultaneously request a file from a content provider such that the difference in their nonce values is one unit. The first client's network delay to load the first n chunks could be greater than the corresponding delay for the second client, creating a scenario where the second client's request for chunk $n + 1$ reaches the appropriate peer before that of the first client. To handle

such scenarios, peers can maintain a sliding window of past nonce values. Incoming nonce values cannot match those in the window, and must be larger than the nonce value received immediately prior to the start of the window.

2.3.2 Collusion Attacks

Collusion between clients and content providers. As noted above, a malicious content provider may want to prevent peers from receiving payments. While monotonically increasing nonces prevent the attack when a content provider acts alone, they are insufficient when content providers collude with clients. Specifically, a client (on behalf of a content provider) could fail to provide a PoDC to a peer that generates a payment-eligible PoD. Peers can certainly detect such behavior, and immediately stop serving content on behalf of the corresponding content provider. However, this poses a payment issue: peers receive probabilistic payments, so simply halting service for a content provider can yield significant amounts of wasted, uncompensated work. One way to mitigate this is for peers to request PoDC values from clients after every chunk they serve, rather than only doing so after they generate a payment-eligible PoD. In this way, peers would be able to quickly detect malicious clients. However, this solution adds significant bandwidth upload overheads to well-behaved clients. Instead, peers should probabilistically request PoDC values from clients. If the probability of requesting a PoDC is higher than the probability that a peer receives a payment for a chunk download, peers can identify malicious behavior without wasting significant resources.

Collusion between clients and peers (Sybil attack). In an effort to increase payments, peers can collude with clients to lie about content delivery, thereby earning money without expending any resources. As a first step towards prevention, Gringotts's content providers randomize the peers that are responsible for serving the chunks of a file, making it difficult for a colluding client to contact a colluding peer. However, this is insufficient since the cost of creating a client is zero. A malicious peer can spawn a large number of clients to ensure that a significant number of clients contact it. In this model, each malicious client can generate a PoD for the first chunk in a file if it is intended to be served by a colluding peer; otherwise, the client can terminate the connection without wasting any resources.

Unfortunately, existing anomaly detection approaches to prevent collusion are unable to detect such attacks [1]. Instead, Gringotts enforces that the first chunk in each file is downloaded from a secure, trusted node (e.g., a traditional CDN). This modification creates an overhead to client generation, since a client *must* expend bandwidth resources to download the first chunk of a file, before they can download subsequent chunks of the file from peers hoping to earn money. This overhead, in turn, makes Sybil attacks economically unfeasible.

To better understand this solution, consider a file download in which download bandwidth costs peers $\$x$ per chunk, content providers pay $\$y$ per chunk download (on average), and

peer upload costs are $\$z$ per chunk. Further, assume that the fraction of malicious peers owning the chunks of that file is m , and the file comprises l chunks. If peers and clients collude, the expected cost for a client to download the file up to chunk i is,

$$\text{Cost}[i] = 1 \cdot x + (i-1) \cdot (m \cdot 0 + (1-m) \cdot x) \quad (1)$$

This equation states that the first chunk will be downloaded from a secure node (costing x). Additionally, if the remaining $i-1$ chunks are routed to peers with equal probability, then a fraction m of these chunks will be downloaded with zero cost, while the remaining chunks will be downloaded from honest peers. Similarly, the expected value of payments to a malicious peer for the file download until chunk i is,

$$\text{Payment}[i] = (i-1) \cdot m \cdot y \quad (2)$$

The attack can be deemed economically unfeasible if the cost to the colluding client is greater than the payments received by the malicious peer, or,

$$\begin{aligned} \text{Payment}[i] &< \text{Cost}[i], & \forall i \in \mathbb{N} \\ \implies & y \leq x \left(\frac{1}{m} - 1 \right) \end{aligned} \quad (3)$$

Further, to make the system monetarily feasible for peers, their compensation ($\$y$) should be greater than the bandwidth costs of uploading a chunk. This enforces an additional constraint

$$y > z \quad (4)$$

These constraints bound payments based on the fraction of malicious peers serving the chunks in a file. For example, if the payment for a file is twice that of the download cost ($y=2 \cdot x$), then more than 33% of the peers holding the chunks for that file must collude to make the attack economically feasible.

Can a malicious peer spawn a large number of colluding peers to achieve high values of m for a particular file? Such an attack is unfeasible as a colluding peer will have to respond to requests from well-behaved clients, consuming peer bandwidth. Failing to respond to these requests can be easily detected by existing anomaly detection techniques. We do note, however, that the above approach works only if content providers ensure randomization of peer selection for chunk downloads *and* content providers periodically mandate churn in the list of peers considered for a given file. Otherwise, a small number of malicious peers might still potentially achieve a high value for m .

Finally, a key property of the proposed solution is that it does not impose any overhead on well-behaved clients, as bandwidth is only used to download chunks in the requested file. However, requiring that the first chunk of each file is downloaded from a trusted node reduces the potential savings of using peer resources for content delivery, since infrastructure-backed nodes must now serve $1/l$ fraction of overall traffic.

2.3.3 Non-financial Attacks

Malicious client. Though clients are not involved in the financial aspects of content delivery, they can still perform attacks on the system. Most notably, a group of malicious clients can perform a Denial of Service (DoS) attack on a

content provider by failing to support payments to peers for serving that content provider's chunks. Peers that detect such behavior can decide to not serve the content provider's files, preventing content distribution. However, peers can only detect that a client is preventing payments after serving a chunk and requesting a PoDC from the client. A PoDC is requested with probability p for each chunk download, implying that a client will have to download $1/p$ chunks from a peer before receiving a request for a PoDC. Since p is intentionally kept to small values, the client will have to expend significant bandwidth resources (downloading chunks) in order to carry out such a DoS attack, making it impractical.

3 Availability and Quality of Service

Beyond secure incentivization, content providers also expect high Quality of Service (QoS) and availability. This requires efficient file placement strategies (i.e., deciding what peers serve content). Gringotts supports prior P2P file placement strategies [22, 36], and provides content providers with more flexibility in influencing content delivery pricing and performance.

Availability. Content providers must ensure that their files are hosted on a sufficient number of peers to tolerate bursts in request volume. Smart contracts provide a flexible way for content providers to influence the replication and availability of their files. Specifically, payment policies in Smart Contracts can be easily modified to reflect current demands of a file. For example, a content provider can pay more for a file during peak demand times, akin to surge pricing with services like Uber.

Quality of Service (QoS). To ensure fast content downloads [14], each Blockchain record can be modified to include client-observed QoS metrics for the corresponding chunk download. Content providers can then factor in QoS decisions into file placement and payment policies (via Smart Contracts). For example, a content provider can select a peer for a chunk download only if the predicted download speed could support the lowest video bitrate. If no peers can meet the QoS goal, the content provider can redirect requests to the backup node (we discuss routing in §6). We note that misreported QoS information by malicious clients can be easily detected (and mitigated) by considering all published QoS values for each peer.

4 Consumer Survey

To understand the expectations and requirements for user participation in a P2P content delivery service, we commissioned a third-party professional organization to undertake aN 876-person consumer survey in the US. 95% of respondents were between the ages of 18-60; 52% were female; 45% of annual household incomes were between \$25K-\$100K. Each participant was asked 11 questions, relating to payments, resource availability, and participation concerns. Our key findings are:

- 51% of users said they would participate.
- Of those who would participate, 70% expect to earn no more than 50% of their monthly Internet bill.

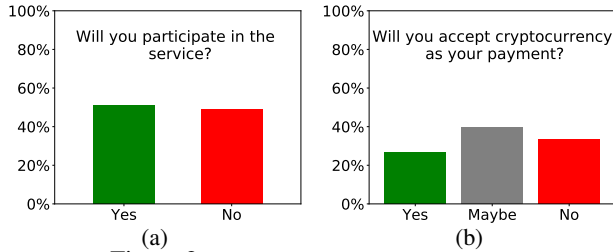


Figure 2: Key consumer survey findings.

Concern	Fraction concerned
Security and Privacy	82.8%
Liability for Illegal Content	50.5%
Performance Impact on Device	47.1%
Payment Concerns	42.1%
Personal Ethics	29.4%

Table 1: Concerns for users who specified that they would *not* participate in the service.

Concern	Fraction concerned
Don't Know How to Use/Sell	59.7%
Not Setup to Receive	54.7%
Volatility and Risk	49.1%
Don't Know What it is	21.9%
Other	6.9%

Table 2: Concerns for users who specified that they would *not* be willing to accept payment in cryptocurrency.

- 27% of users are willing to accept payment in the form of cryptocurrency; 40% were unsure about cryptocurrency payments, while 33% were unwilling.
- Users who would not participate were most concerned with security and privacy (83%), liability over content (50%), and impact on device performance (47%).

Feasibility of participation. We asked users about the devices they have to store/serve content, the free storage space of those devices, and the amount of time they would use those devices in the service. 83% of users owned a laptop, and 82% of laptop owners reported free storage space of more than 20 GB. For context, consider that storing an average 4-minute YouTube video that supports 5 bitrates requires 200 MB of space. Peers with more than 20 GB of free storage can store more than 100 average videos. Further, consider that the average uplink capacity in the US is 22 Mb/s [20]. Streaming HD video requires an average of 5 Mb/s throughput [19], suggesting that peers should be able to stream up to 4 HD videos in parallel. Experiments with an Apache web server on a desktop with a 2.8 GHz processor reveal that mean and peak CPU utilization (single core) are 0.36% and 13%, respectively, for 4 concurrent HD videos, and 1.7% and 33% for 16 concurrent video streams.

The remaining considerations are daily participation times and earnings. 65% of laptop owners stated that they use their computers for more than 2 hours a day. Serving content for 2 hours a day on a 22 Mb/s link amounts to 580 GB of data served per month. If we assume that peers earn \$0.05/GB [3], they will make \$29 per month, which is more than the re-

quirement of covering half of their monthly Internet costs (the average US plan costs \$50 per month [11]). These numbers are conservative as 66% of users would keep their laptops on for longer times, even when not in use, to serve (and earn) more.

Concerns. Figure 2a shows that 49% of users would not participate in the service. Table 1 lists the concerns shared by those users. As shown, a significant fraction of users were concerned about the impact that the service would have on their device, both with respect to security and privacy, and device/network performance. Thus, a practical deployment must ensure host machine isolation and caps for resource consumption. Many users would not participate due to content liability concerns, stemming equally from ethical considerations, legality, and privacy. Consequently, the majority of users were willing to serve movies/shows/news, but only 11% were willing to serve adult content, promoting distributed content filtering and the inclusion of content information in Smart Contracts.

The other primary user concern was with respect to payments. 27% of respondents were willing to accept payments in cryptocurrency, 40% were unsure, and 33% were against it (Figure 2b). Table 2 shows that this concern was largely due to unfamiliarity, as many users did not support cryptocurrency payment because they either did not know how to use/sell them, or they were not setup to receive them. Further, a significant fraction of users were concerned with the volatility and risk of cryptocurrencies. We expect these numbers to decrease in upcoming years as cryptocurrencies evolve and become more widely used, but our results do motivate exploration of secure alternative payment forms.

5 Related Work

Secure Incentivization. Past incentivization strategies for P2P file-sharing systems have focused primarily on preventing “free-riding” using tit-for-tat mechanisms [8, 24]. However, early approaches could not effectively prevent against malicious parties [23]. More recently, there have been proposals to use virtual currencies to solve the free-riding issue [22, 28, 32]. In these solutions, users are granted and debited virtual currency for each transaction they participate in; a central authority is responsible for verifying all transactions. Unlike Gringotts, all of these solutions incentivize only consumers of a service to become resource contributors, limiting adoption.

Peer-to-Peer Systems. Numerous prior systems such as BitTorrent have motivated the potential of P2P file delivery services [7, 8, 15, 26]. Gringotts borrows promising techniques from these systems including policies on file placement and routing strategies (§3). More recently, Filecoin [9] is a P2P file storage service, which uses cyptocurrency-based incentivization to attract peers to store files. Filecoin relies on a “Proof of Storage” (not “Proof of Delivery” like Gringotts) technique which cannot be directly applied to content delivery. Moreover, Filecoin does not let third parties (e.g., content providers) sponsor a client’s download of a file.

6 Discussion

Feedback. This paper primarily focuses on secure incentivization, a key part to P2P content delivery. However, we are seeking feedback about the additional challenges that must be overcome to realize P2P content delivery in practice, both related to our user survey results and not (e.g., from a content provider’s perspective). In addition, we hope to receive suggestions for other domains where secure incentivization and PoDCs can be used beyond content delivery. For instance, could Gringotts be used to build a decentralized overlay network for scaling blockchain transactions [12]?

Controversial aspects. Gringotts’s use of peer resources for content delivery presents a tension with the goals of ISPs, since having peers serve files increases upstream traffic (and costs) for ISPs. Prior work has demonstrated that one way to resolve this tension is to use intelligent peer selection algorithms where peers in the same autonomous system as a client are preferentially selected [5, 37]. However, it is not clear how this would work alongside other ISP mechanisms for revenue generation.

Likely discussion. We hope that this paper generates discussion along four axes. First, for which Internet services should we revisit P2P support? Second, how can we address the concerns of peers highlighted in our user survey (§4)? Third, can Gringotts’s secure incentivization strategy be applied to more generalized resource sharing settings? A key challenge to generalization is creating a policy which can prove that a service-specific task was performed correctly. For example, Gringotts requires “Proof of Delivery,” Filecoin [9] uses “Proof of Storage,” and offloading GPU-based tasks to peers (rather than costly cloud services) would require “Proof of Faithful Computation.” Fourth, in this paper, we only considered peers to be personal computers. To increase participation, could other devices be used, such as continually-powered WiFi routers, IoT devices, or unused machines in colocation datacenters?

Open issues. In addition to the challenges from our user survey (§4), a key open issue relates to routing. With Gringotts, content providers send clients a list of peers to contact for each file chunk. Generating this list in a centralized manner [8] adds computational overhead to the content provider, who must keep track of the set of live peers that is hosting its files. Decentralized routing strategies (e.g., Distributed Hash Tables (DHTs) [18, 30]) can help overcome overheads of a centralized routing strategy but they suffer from several challenges. First, randomizing peer selection with DHTs is difficult, since malicious nodes can directly influence routing decisions, enabling Sybil attacks [31]. Second, discovering the peer to contact for each chunk takes non-negligible time, harming QoS.

Circumstances for failure. If Gringotts fails to attract enough peers, P2P content delivery might not be able to achieve desired availability and QoS requirements. So, while PoDC approach seems robust to vulnerabilities, it does not guarantee adoption. Though this is tolerable when P2P is used to augment existing CDNs, it may preclude standalone P2P systems.

References

- [1] P. Aditya, M. Zhao, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, and B. Wishon. Reliable client accounting for p2p-infrastructure hybrids. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 8–8. Usenix Association, 2012.
- [2] Amazon. Amazon CloudFront – Content Delivery Network (CDN) for Video Streaming. <https://aws.amazon.com/cloudfront/>. July 14, 2018.
- [3] Amazon. Amazon cloudfront pricing. <https://aws.amazon.com/cloudfront/pricing/>.
- [4] N. Anjum, D. Karamshuk, M. Shikh-Bahaei, and N. Sastry. Survey on peer-assisted content delivery networks. *Comput. Netw.*, 116(C):79–95, Apr. 2017.
- [5] D. R. Choffnes and F. E. Bustamante. Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems. In *ACM SIGCOMM Computer Communication Review*, volume 38, pages 363–374. ACM, 2008.
- [6] Cisco. Cisco visual networking index: Forecast and methodology, 2016-2021. <https://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf>. June 6, 2017.
- [7] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing privacy enhancing technologies*, pages 46–66. Springer, 2001.
- [8] B. Cohen. Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer systems*, volume 6, pages 68–72, 2003.
- [9] Filecoin. Filecoin: A decentralized storage network. <https://filecoin.io/filecoin.pdf>. July 19, 2017.
- [10] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 51–68. ACM, 2017.
- [11] H. S. Internet. How much should i be paying for high-speed internet? <https://www.highspeedinternet.com/resources/how-much-should-i-be-paying-for-high-speed-internet-resource/>. Mar 28, 2018.
- [12] U. Klarman, S. Basu, A. Kuzmanovic, and E. G. Sirer. bloxroute: A scalable trustless blockchain distribution network whitepaper.

- [13] G. Kreitz and F. Niemela. Spotify—large scale, low latency, p2p music-on-demand streaming. In *Peer-to-Peer Computing (P2P), 2010 IEEE Tenth International Conference on*, pages 1–10. IEEE, 2010.
- [14] B. Krishnamurthy, C. Wills, and Y. Zhang. On the use and performance of content distribution networks. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, pages 169–182. ACM, 2001.
- [15] N. Leibowitz, M. Ripeanu, and A. Wierzbicki. Deconstructing the kaza network. In *Internet Applications. WIAPP 2003. Proceedings. The Third IEEE Workshop on*, pages 112–120. IEEE, 2003.
- [16] Z. Liu, Y. Ding, Y. Liu, and K. Ross. Peer-assisted distribution of user generated content. In *Peer-to-Peer Computing (P2P), 2012 IEEE 12th International Conference on*, pages 261–272. IEEE, 2012.
- [17] H. Mao, R. Netravali, and M. Alizadeh. Neural Adaptive Video Streaming with Pensieve. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17*, pages 197–210. ACM, 2017.
- [18] P. Maymounkov and D. Mazieres. Kademia: A peer-to-peer information system based on the xor metric. In *International Workshop on Peer-to-Peer Systems*, pages 53–65. Springer, 2002.
- [19] Netflix. Internet connection speed recommendations. <https://help.netflix.com/en/node/306>.
- [20] Ookla. Speed test. <http://www.speedtest.net/reports/united-states/#fixed>. Sep 7, 2017.
- [21] peer5. Peer5 - the serverless p2p cdn for live video streaming. <https://www.peer5.com>. July 10, 2018.
- [22] R. S. Peterson and E. G. Sirer. Antfarm: Efficient content distribution with managed swarms. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation, NSDI'09*, pages 107–122. Berkeley, CA, USA, 2009. USENIX Association.
- [23] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani. Do incentives build robustness in bit torrent. In *Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation, NSDI'07*, pages 1–1, Berkeley, CA, USA, 2007. USENIX Association.
- [24] D. Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *ACM SIGCOMM computer communication review*, volume 34, pages 367–378. ACM, 2004.
- [25] D. Rayburn. Cdn market update: Web performance, diy, and cdn pricing trends. <http://www.streamingmediablog.com/wp-content/uploads/2017/06/rayburn-cdn-pricing-2017.pdf>. May 15, 2017.
- [26] M. Ripeanu. Peer-to-peer architecture case study: Gnutella network. In *Peer-to-Peer Computing, 2001. Proceedings. First International Conference on*, pages 99–100. IEEE, 2001.
- [27] Sandvine. global internet phenomena-latin american & north america.
- [28] M. Sirivianos, J. H. Park, X. Yang, and S. Jarecki. Dandelion: Cooperative content distribution with robust incentives. In *2007 USENIX Annual Technical Conference on Proceedings of the USENIX Annual Technical Conference, ATC'07*, pages 12:1–12:14. Berkeley, CA, USA, 2007. USENIX Association.
- [29] Speedtest. 2017 United States Speedtest Market Report. <http://www.speedtest.net/reports/united-states/>. September 11, 2017.
- [30] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. *ACM SIGCOMM Computer Communication Review*, 31(4):149–160, 2001.
- [31] G. Urdaneta, G. Pierre, and M. V. Steen. A survey of dht security techniques. *ACM Computing Surveys (CSUR)*, 43(2):8, 2011.
- [32] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems*, volume 35, 2003.
- [33] G. Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151:1–32, 2014.
- [34] H. Yin, X. Liu, T. Zhan, V. Sekar, F. Qiu, C. Lin, H. Zhang, and B. Li. Livesky: Enhancing cdn with p2p. *ACM Trans. Multimedia Comput. Commun. Appl.*, 6(3):16:1–16:19, Aug. 2010.
- [35] G. Zhang, W. Liu, X. Hei, and W. Cheng. Unreeling xunlei kankan: Understanding hybrid cdn-p2p video-on-demand streaming. *IEEE Transactions on Multimedia*, 17(2):229–242, 2015.
- [36] Y. Zhang, C. Gao, Y. Guo, K. Bian, X. Jin, Z. Yang, L. Song, J. Cheng, H. Tuo, and X. Li. Proactive video push for optimizing bandwidth consumption in hybrid cdn-p2p vod systems. In *IEEE INFOCOM*, 2018.

[37] M. Zhao, P. Aditya, A. Chen, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, B. Wishon, and M. Ponc. Peer-assisted content distribution in akamai netsession. In *Proceedings of the 2013 Conference on Internet*

Measurement Conference, IMC '13, pages 31–42, New York, NY, USA, 2013. ACM.