

Proof of Backhaul: Trustfree Measurement of Broadband Bandwidth

Peiyao Sheng^{*†}, Nikita Yadav^{†‡}, Vishal Sevani[†], Arun Babu[†], SVR Anand[†], Himanshu Tyagi^{†‡} and Pramod Viswanath^{†§}

^{*}University of Illinois at Urbana Champaign, USA

Email: psheng2@illinois.edu

[†]Witness Chain, USA

Email: {vishal.sevani, arun.babu, svr.anad}@kaleidoscope-blockchain.com

[‡]Indian Institute of Science, India

Email: {nikitayadav, htyagi}@iisc.ac.in

[§]Princeton University, USA

Email: pramodv@princeton.edu

Abstract—Recent years have seen the emergence of decentralized wireless networks consisting of nodes hosted by many individuals and small enterprises, reawakening the decades-old dream of open networking. These networks have been deployed in an organic, distributed manner and are driven by new economic models resting on performance-based incentives. A critical requirement for the incentives to scale is the ability to prove network performance in a decentralized “trustfree” manner, i.e., a Byzantine fault tolerant network telemetry system.

In this paper, we present a Proof of Backhaul (PoB) protocol which measures the bandwidth of the (broadband) backhaul link of a wireless access point, termed prover, in a decentralized and trustfree manner. In particular, our proposed protocol is the first to satisfy the following two properties: (1) *Trustfree*. Bandwidth measurement is secure against Byzantine attacks by collusion of challenge servers and the prover. (2) *Open*. The barrier-to-entry for being a challenge server is low; there is no requirement of having a low latency and high throughput path to the measured link. At a high-level, our protocol aggregates the challenge traffic from multiple challenge servers and uses cryptographic primitives to ensure that a subset of challengers or, even challengers and provers, cannot maliciously modify results in their favor. A formal security model allows us to establish guarantees of accurate bandwidth measurement as a function of the maximum fraction of malicious actors.

We implement our protocol with challengers spread across geographical locations and release the code [5]. Our evaluation shows that our PoB protocol can verify backhaul bandwidth of up to 1000 Mbps with less than 10% error using measurements lasting only 100 ms. The measurement accuracy is not affected in the presence of corrupted challengers. Importantly, the basic verification protocol lends itself to a minor modification that can measure available bandwidth even in the presence of cross-traffic.

Finally, the security guarantees of our PoB protocol output are naturally composable with “commitments” on blockchain ledgers, which are commonly used for decentralized networks.

I. INTRODUCTION

Provisioning of robust and fast internet bandwidth is a critical need of the information age. Bandwidth testing is a cornerstone of the field of classical computer networking, and continues to be relevant to this day. The contemporary Internet relies heavily on consistent and trustworthy bandwidth, spurred by the emergence of bandwidth-intensive technologies and applications such as Ultra High Definition videos, high-speed 5G networks and Augmented Reality/Virtual Reality (AR/VR). Commercial bandwidth testing services such as SpeedTest [6] and Fast [2] endeavor to provide measurement accuracy by transmitting large files over a long duration through powerful servers. Recent scientific studies [57], [56] aim to improve the efficiency of measurement both spatially and temporally via a statistical sampling framework. In all these scenarios, both the network operation and bandwidth measurements are largely centralized operations, involving coordination and integration and trust among network and measurement entities.

Separately, decentralized networks have been in the making for decades. Starting with Software Defined Networking [36], [35] to simplifying the hardware and open software [45] to facilitate application development, to real-world deployments of decentralized Internet Service Providers (ISPs) [4] and decentralized Mobile Network Operators (MNOs) [25] have emerged. These decentralized networks have been made possible by the convergence of several engineering, business, and policy developments: the availability of cheap hardware for WiFi access points (and now even cellular base stations); the availability of cloud-native orchestration and AAA software [39]; and the availability of lightly licensed spectrum (CBRS, standing for citizen band radio service has 150MHz of spectrum between 3.55 and 3.7 GHz) for cellular communication [23]. CBRS-driven “private 5G” network offerings allow the participation of smaller operators (including small business complexes) and are emerging as efficient and complementary avenues to the traditional vertically-integrated wide area cellular network service [10], [11], [12]. Apart from private 5G networks augmenting traditional cellular carriers, a token-driven incentive ecosystem known as the Helium network [25] has led to hundreds of thousands of “hotspots” being hosted by individuals, although yet to be integrated into a multi-RAT

(radio access technologies) network.

With the growth of these decentralized networks, a central and new engineering challenge emerges – we need to design *secure and decentralized* network telemetry. In centrally managed networks, network telemetry is used for performance measurement and subsequent optimization. In contrast, network telemetry plays a more pivotal role in decentralized networks: it is now needed to ensure incentive compatibility, i.e., the network nodes provide the service that they are being paid for. For this purpose, there are two new requirements for decentralized network telemetry:

- **Trustfree.** The protocol is secure against Byzantine attacks by the parties involved.
- **Open.** The barrier-to-entry for servers participating in decentralized telemetry is low. In particular, any node with a “reasonably good” internet connection should be able to participate.

The measurements that we get as the output of such trustfree and open network telemetry protocols can be viewed as a cryptographically secure proof of network performance.

In this paper, we focus on measuring a specific network performance parameter which is of central importance in decentralized wireless network deployments. In such deployments, users are required to get a broadband connection with appropriate bandwidth, as a backhaul for the wireless access point. But how do we know that the user has indeed set up a good backhaul connection? Can we apply any existing techniques from the extensive, decades-spanning literature on bottleneck link-throughput measurement? It turns out none of the existing tools is applicable for our setting; below we point out shortcomings of prominent techniques and clarify our contributions.

Comparison with speedtest. Speedtest (speedtest.net) is a state-of-the-art bandwidth testing tool widely used globally. Whenever a user (the “prover”) sends a measurement request, a nearby server is selected from a centralized challenger server pool. The selected server generates traffic continuously until the target link is saturated. This necessitates the challenger server to maintain a high bandwidth, low latency, and low packet loss link to the prover, which presents a significant barrier to entry for potential challengers. Furthermore, the measurements rely on the rates of sending packets from challengers and acknowledgements from the prover – an untrustworthy prover or challenger can adversely impact the measurement. Speedtest and similar architectures are unsuited for trustfree network telemetry.

Traffic aggregation. One way to allow more challengers to participate in the telemetry (and thus being more open) is to aggregate traffic from multiple challengers. Such aggregation removes the requirement of high capacity for a single server to measure high-bandwidth links, by uniting a group of servers to generate sufficient traffic in *parallel*. While this technique can improve the accuracy (e.g., recent works [14], [2], [57], [56]), the method is not trustfree: a Byzantine prover can readily manipulate the measurement results with no check or balance.

Interactive Measurement. To eliminate the need of trust on the prover, challengers should interact with other parties in

the network to generate measurements. Popular interactive telemetry tools such as traceroute [28] and pathchar [29] use the timing information obtained by combining the Internet control message protocol’s (ICMP) time-to-live (TTL) and packet dropped messages to estimate link performance over the Internet. In particular, challengers estimate the round-trip time (RTT) to the two endpoints of the link to be measured. The throughput is then derived by dividing the packet size by the difference in RTT. To further increase accuracy, packets of different sizes can be transmitted and the measurements can be aggregated through linear regression [29], [22]. Secure measurement, resistant to collusion between the prover and challengers, is not guaranteed in these protocols.

Our contributions. We present the first multichallenger PoB protocol for measuring backhaul bandwidth that satisfies the aforementioned trustfree and open properties, (details in §III). Broadly, the protocol is built by implementing the following ideas:

- 1) *Traffic aggregation.* We simultaneously send challenge traffic from multiple challengers to the prover. The duration of the challenge is chosen to be sufficiently high to ensure that traffic from all the challengers queues at the prover’s backhaul link.
- 2) *Unforgeable probe.* We randomly select challengers from a larger pool, each of which sends digital signatures as traffic. This prevents any other party from forging the measurement probe. Furthermore, we limit the amount of challenge traffic from each challenger to curtail any single party’s influence.
- 3) *Short witness.* The prover can send a short message to the challengers to prove that it has received appropriate amount of data. As will be seen below, our security considerations require us to use a partially verifiable hash. For this purpose, we use a Merkle tree [43].
- 4) *Robust timing measurement.* We estimate the overall challenge’s round-trip time (RTT) by taking the median of the RTTs measured by various challengers.

We implement these steps and experimentally validate the design choices to identify the best performing configuration; see Figure 1 for a depiction. The main contribution of this work is the trustfree property of the proposed protocol – it is secure under a rigorous threat model that we outline in §IV. Our proposed protocol is the first one that can measure bandwidth of hundreds of Mbps without requiring any specialized server with high throughput and low latency for challengers in the trustfree setting. We further extend this protocol to measure available bandwidth in the presence of cross-traffic, making it a truly distributed “speedtest.”

We analyze the security of our multichallenger PoB protocol under a formal threat model which allows any subset of parties (up to $1/3$ challengers collaborating with the prover) to maliciously deviate from the protocol. Given the unforgeable nature of our probe, even a corrupted prover is still required to obtain probe packets from the challengers. However, corrupted challengers, too, can modify the packet flow using two attacks: (i) the *withholding* attack where a corrupted challenger does not send probe packets; and (ii) the *rushing* attack where a corrupted challenger coordinates with the corrupted prover to

Technique	Secure	Challenger BW < Backhaul BW	Accuracy
Pathchar [22], [29], [40]	✗	✓	Low
Packet dispersion based [17], [18], [38], [50]	✗	✗	-
Secure BW estimation [33], [53], [59]	✓	✗	-
Multichallenger PoB	✓	✓	High

(a)

Backhaul BW (Mbps)	Challenger BW (Mbps)	Challenge Data (MB)	Attack	Measured BW (Error %)	Guaranteed BW (Mbps)
250	25	3.44	-	246 (1.6%)	184
500	20	6.86	-	475 (5%)	356
750	75	10.31	-	705 (6%)	529
1000	100	13.75	-	921 (8%)	691
250	32	3.44	Rushing	331 (0.6%)	249
250	32	3.44	Withholding	241 (3.6%)	181

(b)

TABLE I: (a) Comparison of our multichallenger PoB protocol with prior-art techniques. (b) Summary of our performance results with 10 challengers. We perform attacks with 2 corrupted challengers.

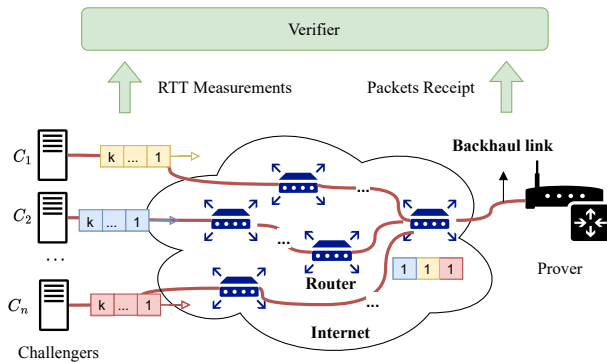


Fig. 1: The multichallenger PoB Protocol.

send the packets or their information quickly without using the challenged link. To compensate for the withholding attack, we must send more packets than the link bandwidth to have sufficient traffic even after withholding attack. To compensate for the rushing attack, we multiply the actual measured bandwidth with a correction factor derived from the bound of bandwidth inflation to arrive at the *guaranteed bandwidth*. In addition, corrupted challengers can modify their outputs needed for verification. Specifically, they may report wrong RTT or they may claim modified packet data. We circumvent the former attack by taking a median of the measurements. To circumvent the latter attack, we employ a Merkle tree. This allows us to verify the consistency of the hash response from the prover with the data of uncorrupted challengers, without the need for correct data from the corrupted parties.

Overall, denoting the maximum fraction of corrupted challengers by β , we show that for $\beta < 1/3^1$, our protocol does not allow any prover to inflate the bandwidth and allows an honest prover to establish at least a fraction $(1 - 2\beta)/(1 - \beta)$ of the true bandwidth.

Design scope exploration. Our main protocol assumes availability of digital signatures and nothing more. Different variants of this protocol can be implemented depending on the availability of resources; we explore the design scope in § V and present the results in Table II. Specifically, we

¹The adversarial threshold can be $1/2$ if the verifier has access to a timer. See the discussion in §IV-B.

observe that (1) the protocol can achieve a higher accuracy of $(1 - \beta)$ using TCP packets if a fairness property holds, as this prevents corrupted challengers from reporting deviating bandwidths higher than the majority; (2) the accuracy is reduced to $(1 - 3\beta)/(1 - \beta)$ when digital signatures are not used, due to the deniability of the packets delivery; (3) by limiting the access of corrupted challengers to additional links with higher bandwidth, the feasibility of rushing attacks is eliminated. Additionally, by implementing a shuffled packet creation phase, the effects of other attacks can be minimized and optimal accuracy can be achieved.

Implementation and evaluation. To convert the idealized protocol into a practical tool, we implement a variant of our protocol designed to address real-world issues (§VI) and thoroughly evaluate its performance (§VII). For instance, measuring links with 100 Mbps and higher bandwidth, which is commonplace in broadband services, requires latency measurements with an accuracy that is hard to achieve due to jitter in the Internet; we elaborate on overcoming this challenge in Appendix B.

Our evaluation focuses on the loss of measurement accuracy that occurs when using multiple challengers and traffic aggregation; in particular, we consider the loss of accuracy due to: (i) time synchronization errors and network jitter; (ii) computation time delays due to the use of digital signatures, hash computation, verification, and Merkle trees; and (iii) geographically spread challengers with heterogeneous capabilities. We also implement rushing and withholding attacks to illustrate that the security guarantees of the theory hold in practice. Our main experimental results are summarized in Table I. We report both the actual measured bandwidth and the output of our protocol – the guaranteed bandwidth – which is obtained by applying the correction factor $(1 - 2\beta)/(1 - \beta)$.

II. BACKGROUND AND RELATED WORK

Bandwidth estimation. In data networks, bandwidth signifies the data volume a network path can transfer over time. Two metrics related to bandwidth are extensively investigated in the literature, the maximum possible data rate called *capacity* and the maximum available data rate called *available bandwidth* [50]. Packet dispersion techniques [18], [17], [52], [34], [38] are widely used to measure bottleneck link capacity. Available bandwidth measurement techniques are outlined in [30], [31], [51], [54], [27], [16], [42], [41], [6], [2],

[57], [56]. Of these, Pathload [30], [31] and Pathchirp [51] generate brief traffic loads with varying stream rates to adjust estimations based on one-way delay differences. The state-of-the-art commercial tools [2], [6] use global high-bandwidth servers to generate saturating TCP traffic on target links, while recent research like FastBTS[57] reduce the data usage and bandwidth requirements through concurrent connections. Swiftest[56] leverages UDP to address TCP limitations. Given bandwidth measurements’ importance, certain works focus on securing these against inflation attacks [33], [59]. Others, such as [53], propose secure estimation against malicious peers, using consensus vectors from network-wide bandwidth evaluations. This scheme, while opportunistic, doesn’t fully appraise high-bandwidth nodes. In another direction, Torpath [24] suggests a proof system for network telemetry, remunerating Tor network relays based on transmitted data volume. In comparison, in this paper, our PoB protocol targets backhaul bandwidth measurement of end nodes, with no bandwidth requirements on the measuring nodes.

Per-hop capacity estimation. Of all the bandwidth estimation techniques in literature, [28], [29], [22], [40], [37], [47], [26] are closest to our work. These techniques can measure capacity for any link in an end-to-end path and so can be used to measure the prover backhaul, which is our goal. Traceroute [28] and pathchar [29], [22] make use of time-to-live (TTL) information in ICMP packets to control the packet drop at different intermediate hops to measure capacity of any link. [37], [47], [26] improve the approach used by pathchar [29], [22] with variable packet sizes. However these techniques require precise timing measurements of the order of packet transmission times. For bandwidth in 100s of Mbps, the packet transmission times are of the order of tens of microseconds. Given the jitter in latency over the Internet, our experiments reveal that such precise timing measurements are difficult. Indeed, previous work [26] reports errors over 20% for measuring bandwidths of 500 Mbps or more.

Decentralized networks. Every decentralized network proposal features a proof system verifying network performance parameters, with participants benefiting from substantiating their contribution. Helium [25] exploits blockchain’s potential to establish a decentralized data network via a proof-of-coverage tokenized incentive. Hotspots receive compensation for reliable coverage confirmed by regular challenge requests and beacon responses. Althea [55] operates as a distributed ISP, fostering last-mile connectivity through a competitive platform, using route meters and accuracy scores to ensure quality connections. Route Bazaar [19] creates end-to-end connectivity agreements among untrusted parties, resolving contractual and routing difficulties. The path performance is ensured via periodic forwarding proofs on blockchains, encapsulating encrypted path tags, traffic samples, and performance metrics.

III. THE MULTICHALLENGER POB PROTOCOL

In this section, we formulate the PoB problem (§III-A), introduce main techniques (§III-B) and describe our multichallenger PoB protocol in details (§III-C).

A. Problem Statement

We consider a system consisting of a group of end nodes such as base stations, WiFi access points and remote servers over the Internet willing to assist with backhaul bandwidth measurement. All nodes are connected to the network core through one backhaul link, simply referred to as backhaul from hereon, with an internal state θ representing the bandwidth of the link. We model the network core as a single point since fiber cables usually provide extremely high bandwidth, e.g., 100 Gbps. A PoB protocol allows a trusted verifier to use a subset of available nodes for securely measuring the backhaul bandwidth for a specific node called a *prover*, denoted P . The verifier can not observe the internal state θ_P of the prover directly. Instead, it needs to interact with the system by issuing a challenge request to the rest of the parties. We assume that n participants serve as *challengers*, denoted as $\{C_1, \dots, C_n\}$, among which up to $f = \beta n$ challengers are corrupted, where β represents the fraction of adversarial challengers. These challengers are responsible for generating and sending probes to the prover and output the measurements to the verifier. The output of PoB protocol is an estimation of the backhaul bandwidth of the prover. It guarantees the following two security properties:

- **Approximate completeness:** When the prover is uncorrupted, if the protocol outputs θ'_p , the actual bandwidth of the prover θ_p satisfies $\theta'_p \geq \alpha \theta_p$ for a constant accuracy ratio $\alpha \in [0, 1]$.
- **Soundness:** The protocol will not output a bandwidth higher than θ_p , even when the prover is corrupted.

Design landscape. To maximize the accuracy of the protocol while ensuring soundness, we explore the design landscape which takes into account a set of fundamental parameters. Firstly, we explore the cryptography primitives that can be utilized, specifically debating between digital signatures and pseudorandom generators. Each of these cryptographic tools offers its own set of advantages; digital signatures provide non-repudiation and integrity assurance, while pseudorandom generators can supply a degree of randomness with less overhead. Secondly, our protocol design considers the types of packets that will be utilized for data transmission. We weigh the pros and cons of TCP and UDP packets, including slow start and fairness. Lastly, we investigate the threat model within which our protocol will operate. Of particular interest is the potential for adversaries to launch “rushing attacks”, in which the adversary attempts to gain an unfair advantage by rushing to transmit information before honest parties. The details and implications of such attacks are explored in subsequent sections (§ IV-A). These considerations (summarized in Table II) form the bedrock of our protocol design and implementation, guiding us towards a balanced solution that combines robust security and performance.

Other assumptions for theoretical analysis. Except for signatures and pseudo-random generator, our protocol makes use of collision resistant cryptographic hash functions. These primitives are assumed to be perfectly secure. When using digital signatures, parties verify every received signature by default and ignore those signed invalidly. We assume the network is synchronous and every challenger has access to a synchronized clock. Each node knows the public address and public key of

TABLE II: Comparison of different protocols in design landscape

	Packets	Crypto primitive	Rushing attack	Accuracy
PoB	UDP	signature	Yes	$(1 - 2\beta)/(1 - \beta)$
PoB-TCP	TCP	signature	Yes	$1 - \beta$
PoB-PRG	UDP	pseudorandom generator	Yes	$(1 - 3\beta)/(1 - \beta)$
PoB-shuffle	UDP	signature	No	$1 - \frac{(1+\delta_b)\beta^t}{(1-\delta_g)(1-\beta)^t}$ [ⓐ]

[ⓐ] $0 < \delta_b, \delta_g \leq 1$

others. We suppose there exists a trusted verifier such as a blockchain to broadcast information to the system. It is crucial to note that the assumptions made in our theoretical analysis are for the purpose of simplification and easy understanding. However, when evaluating our implementation, we take into account the possible deviations from these assumptions that may occur in real-world scenarios.

B. Protocol Overview and Main Results

Heuristically, the protocol proceeds by randomly selecting a set of n challengers from all the participants to send a train of probes to the prover (Figure 1). The protocol enforces packets from different challengers to arrive at the link to be measured around the same time. This traffic aggregation strategy effectively combines the group of challengers to an equivalent challenger with larger bandwidth and thereby renders the prover’s backhaul the bottleneck link.

Formally, suppose that the protocol starts at time t_0 , and each challenger C_i starts to send a sequence of k packets of size b each at time t_{i1} , $1 \leq i \leq n$. We have the following two requirements:

- 1) *Aggregation condition.* There is a $\theta_0 \leq \min(\theta_1, \dots, \theta_n)$ such that the bandwidths θ_i of C_i satisfy

$$t_0 + \frac{b}{\theta_0} = t_{11} + \frac{b}{\theta_1} = \dots = t_{n1} + \frac{b}{\theta_n}. \quad (1)$$

- 2) *Bandwidth condition.* The quantity θ_0 satisfies

$$(n - f) \cdot \theta_0 \geq \theta_p. \quad (2)$$

The “aggregation condition” coordinates the arrival time of packets sent from various challengers, allowing multiple traffic flows to be effectively aggregated and merged into one stream at an appropriate rate. In this way, at least $(n - f)b$ bits of data are transmitted within the transmission time of one packet for a single challenger (b/θ_0). Therefore, the equivalent bandwidth of the challenger group is enlarged by at least a factor of $(n - f)$. The “bandwidth condition” ensures that the prover’s backhaul becomes the bottleneck link. We assume there always exist enough potential challengers to meet both conditions.

While honest participants are supposed to correctly report their own bandwidth and send packets on time, corrupted parties can violate the conditions in arbitrary ways. For instance, a corrupted challenger can *rush* the packets through extra links or refuse to send any packets. We therefore require the prover to send back a response to all challengers on receiving $(n - f)k$ packets as a transmission receipt, since we can not expect more

packets in the case of a *withholding* attack (detailed in §IV-A). Then challengers measure the time it takes to transmit all these packets. Since corrupted challengers can claim arbitrary values, the median of all reported time is used to avoid manipulations and provide robust timing measurement.

Cryptographic primitives. To save bandwidth use during verification, the prover transmits a concise witness: a hash of received packets to terminate measurements. We utilize a hash function $Hash(\cdot)$ that accepts any string and delivers a deterministic fixed-length random string. If input is a set of messages, it’s serialized to a string for hashing. For verification, each challenger confirms only its sent packets and leverages the Merkle tree to perform inclusion checks with partial data. Using $MerkleRoot(\cdot)$, a sequence of hashes aggregates to a single cumulative hash, reducing verification overhead per challenger to $O(\log n)$, significantly cutting response traffic for large n . For our protocols using digital signatures to create unforgeable probes and ensure bad behavior traceability, we use a key generation function $keyGen$ outputting a secret-public key pair, a signing function $sign(sk, msg)$ enabling anyone to sign any message with a secret key sk , and a verification function $verify(pk, msg, \sigma)$ to confirm if signature σ is derived by signing the given message msg with the secret key paired with public key pk . For protocols involving a pseudo-random generator, a function $gen(seed)$ yields the next random number from the *seed*.

Blockchain as a verifier. Our PoB protocol is triggered by a challenge request issued from a verifier, who is also responsible for the broadcast of public parameters such as protocol start time t_0 and bandwidth requirement θ_0 . Generally, any trusted entity can play the role of a verifier. In tokenized decentralized settings, smart contracts supported by blockchains are a good fit to transparently generate, broadcast protocol parameters and coordinate measurement reports from multiple challengers. Implementing a version of our protocol with blockchain as a verifier (deploying appropriate smart contracts) is beyond the scope of this paper.

Main results. Table II presents different variants of our protocol designs with their corresponding parameters. Our main protocol PoB uses UDP packets for probing, operates under the assumption that every packet is signed using digital signatures and adversarial nodes have access to extra high-bandwidth links to send packets. We describe the full protocol in the subsequent subsection and analyze its security and accuracy in § IV-B. Detailed discussions of these variants are found in § V, revealing intriguing elements in these explorations.

We observe, firstly, that the protocol can enhance its accuracy to $(1 - \beta)$ when it utilizes TCP packets, a version

referred to as PoB-TCP. This heightened accuracy is achievable only when a fairness property is maintained. Fairness ensures that nodes sending traffic at the same speeds simultaneously to a link receive the same rate, therefore we can rely on the results reported by the majority of nodes to deter malicious inflation. Nonetheless, TCP flows come with inherent drawbacks, such as a slow start, rendering it less than ideal for implementation.

Secondly, we observe a reduction in the protocol’s accuracy to $(1 - 3\beta)/(1 - \beta)$ in the absence of digital signatures. To ensure packet unforgeability without signatures, the protocol employs a pseudorandom generator, thus earning it the name PoB-PRG. This method is more cost-effective to calculate. In particular, challengers send packets containing a random number, generated from a unique seed committed to the verifier before initiating the protocol and revealed afterwards. This procedure enables the prover to verify the correctness of the received packets. However, senders can deny having sent the packets, which precipitates the accuracy drop.

Moreover, we point out that by restricting corrupted challengers’ access to additional high-bandwidth links, we can eradicate the possibility of rushing attacks. These attacks would enable challengers to complete measurements in a shorter time. Given this assumption, we suggest the PoB-shuffle protocol. This variant introduces a shuffled phase, wherein challengers are randomly selected from the pool based on bandwidth to co-sign the packets. This strategy minimizes the effects of collusion between challengers and provers, as more challengers would need to collude to launch attacks. This enhancement optimizes accuracy in a probabilistic way, provided that security parameters are aptly set and the pool size is sufficiently large.

C. Full PoB Protocol

The full PoB protocol contains two phases, a *measurement phase* described in Algorithm 1, where challengers generate and send packets, and a *verification phase* described in Algorithm 2, where the prover constructs proofs for the verifier. Finally, the verifier outputs the final results after verification.

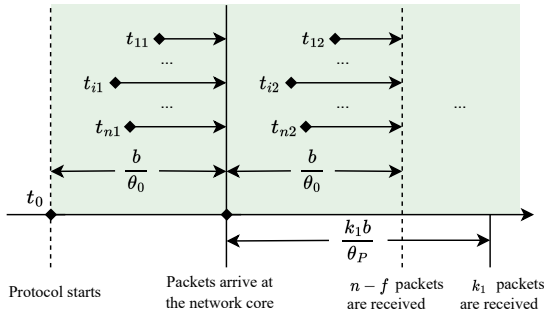


Fig. 2: The measurement phase of PoB protocol. t_{iq} is the time for challenger C_i to send the q -th packet. k_1 is the actual number of packets received with sequence number 1.

Measurement phase. At the beginning of the measurement phase, the verifier produces three public protocol parameters (t_0, m_0, θ_0) and broadcasts it to all challengers, where t_0 is the start time of the protocol, m_0 is a random message, θ_0 is the global minimum bandwidth. To participate in the measurement

Algorithm 1 The Measurement Phase of PoB Protocol

```

1: as a challenger  $C_i$ 
2:    $t_0, m_0, \theta_0 \leftarrow$  generated and broadcast by verifier
3:   measure its own bandwidth  $\theta_i$  (require  $\theta_i \geq \theta_0$ ).
4:   generate  $(pk_i, sk_i) \leftarrow keyGen$ , send  $pk_i$  to verifier.
5:   for sequence number  $q = 1, \dots, k$  do
6:      $t_{iq} \leftarrow t_0 + q \cdot b/\theta_0 - b/\theta_i$ 
7:      $\sigma_{iq} \leftarrow sign(sk_i, (q, m_0))$ ,  $m_{iq} \leftarrow (i, q, \sigma_{iq})$ 
8:     send packet  $m_{iq}$  to  $P$  at  $t_{iq}$ 
9:   upon receiving  $(h_{1i}, h_2, \sigma_i)$  from  $P$  do
10:    if  $verify(pk_P, (h_{1i}, h_2), \sigma_i)$  outputs 1 then
11:      record round trip time  $\Delta_i \leftarrow curTime_i - t_i$ 
12:
13: as a prover
14:    $\forall i \in [1, n], \mathcal{M}[i] \leftarrow \emptyset, (pk_p, sk_p) \leftarrow keyGen$ 
15:   upon receiving packet  $M'$  from  $C_i$  do
16:      $(i, q, \sigma) \leftarrow M'$ 
17:     add  $(q, \sigma)$  to  $\mathcal{M}[i]$ 
18:     if  $\sum_{j=1}^n |\mathcal{M}[j]| = (n - f)k$  then
19:        $\forall j \in [1, n], h_{1j} \leftarrow Hash(\mathcal{M}[j])$ 
20:        $h_2 \leftarrow MerkleRoot(\{h_{1j}\}_{j \in [1, n]})$ 
21:        $\forall j \in [1, n], \sigma_j \leftarrow sign(sk_p, (h_{1j}, h_2))$ 
22:        $\forall j \in [1, n]$ , send  $(h_{1j}, h_2, \sigma_j)$  to  $C_j$ .

```

process, challengers must first measure their own bandwidth θ_i and generate a key pair consisting of a public key pk_i and a private key sk_i . The public key is then sent to the verifier. The time to start sending the first packet t_{i1} is determined by Eq. (1). The challenger C_i generates a sequence of k packets by signing the public message m_0 together with a sequence number q and sends them one by one to the prover with a fixed duration b/θ_0 . The process is depicted in Figure 2, where the q -th packet of C_i is sent at time t_{iq} (see line 6 of Algorithm 1). On receiving the packets from the challengers, the prover separates the messages from different challengers and adds them to corresponding sets. When the total number of received packets reaches $(n - f)k$, the prover generates a response to broadcast to all challengers. This terminates the measurements phase. The response contains (1) a receipt h_{1i} to each challenger C_i , which is the hash of all packets received from the same sender; and (2) a Merkle root h_2 constructed from all receipts. All challengers record the round trip time Δ_i between the start time t_{i1} and the time $curTime_i$ at which a valid response is received.

Verification phase. In the verification phase, the prover is responsible for proving to the challengers the content of received packets. To that end, it constructs another response revealing the indices B_i of packets received from C_i and showing the inclusion of each receipt in the Merkle tree with a Merkle proof P_i . It also sends the Merkle root to the verifier. On receiving the Merkle proof P_i from the prover, the challengers reconstruct the receipt hash and the Merkle root. The challenger C_i forwards Δ_i and the number of packets sent by it to the verifier, after making sure that both the hashes are consistent. At the end of the second phase, the verifier aggregates these measurements from all the challengers about how long the measurement phase takes and how many packets are indeed received by the prover. It also forwards the reports

Algorithm 2 The Verification Phase of PoB Protocol

```
1: as a prover
2:    $B_i \leftarrow$  a bitmap of size  $k$  where the  $q$ -th bit in the
   bitmap is set if some packet  $(q, *) \in \mathcal{M}[i]$ 
3:    $P_i \leftarrow$  the Merkle proof of  $h_{1i}$ 
4:   send  $(B_i, P_i)$  to  $C_i$ , output (REPORT,  $h_2$ ) to verifier.
5:
6: as a challenger
7:   upon receiving  $(B_i, P_i)$  from  $P$  do
8:     Add all  $(q, \sigma_{iq})$  to a set  $\mathcal{M}$  if the  $q$ -th bit is set to
     1 in  $B_i$ . Check whether  $Hash(\mathcal{M}) = h_{1i}$ .
9:     Reconstruct the Merkle root  $h'$  using  $P_i$  and  $h_{1i}$ .
     Check whether  $h_2 = h'$ .
10:    If both two checks are passed, output (REPORT,
     $P, h_2, \Delta_i, |\mathcal{M}|$ ) to verifier.
11:
12: as a verifier
13:    $\mathcal{M} \leftarrow \emptyset, cnt \leftarrow 0$ 
14:   upon receiving (REPORT,  $h$ ) from prover  $P$  do
15:     record  $h$ 
16:   upon receiving (REPORT,  $P, h_2, \Delta_i, k_i$ ) from  $C_i$  do
17:     Check  $h_2 = h$ , add  $\Delta_i$  to  $\mathcal{M}$ ,  $cnt \leftarrow cnt + k_i$ 
18:     if  $cnt \geq (n - f)k$  and  $|\mathcal{M}| \geq n - f$  then
19:        $\Delta' \leftarrow Median(\mathcal{M})$ 
20:        $\theta'_p \leftarrow \frac{cnt \cdot b \cdot (n - 2f)}{\Delta' \cdot (n - f)}$ 
21:       Output (POB,  $P, \theta'_p$ )
22:
```

from challengers to the prover, who checks the consistency and submits the packets and the Merkle proof in case disputes exist. Once the verifier has received “sufficiently many” valid reports. Specifically it waits to receive a confirmation from at least $n - f$ challengers with at least $(n - f)k$ packets in total. We assume there is no packet drop and address the packet drop problem in §VI. Then it calculates the final output bandwidth by dividing the total size of received packets by the median of the reported RTTs; see line 20 of Algorithm 2.

IV. SECURITY MODEL AND ANALYSIS

The primary challenge in trustfree networking is the inherent security vulnerability, since any party can depart from the protocol at will and even collude with other parties to manipulate the results. In this section, we formalize a broad threat model underlying measuring bandwidth, systematically examine the security issues to which the system is exposed (§IV-A), and analyze the security guarantees (§IV-B).

A. Threat Model

We consider a static adversary allowed to corrupt at most f among n challengers before the protocol starts, the rest of uncorrupted challengers are referred to as honest. The prover can also be corrupted. In addition to the backhaul link indicated in the model in §III-A, we allow the adversary to access external communication channels. Specifically, the adversary has access to additional links with arbitrarily high bandwidth connecting to all the participants. The corrupted parties can act arbitrarily in order to either inflate or deflate the measured bandwidth; we discuss prominent attacks below.

Withholding attack. The measurement of bandwidth requires the challengers to send probes and measure the time it takes for the prover to receive the probes. The corrupted challengers who have been bribed by the consumers or the competitors of a prover might be motivated to deflate the bandwidth estimation to reduce service costs. They can delay the sending of the packets to increase the observed RTT or even withhold the packets for the entire protocol. During the verification phase, corrupted challengers can also refuse to report verification results. Moreover, the prover can also bribe the challengers to withhold packets during the measurement phase but report that the maximum number of packets have been sent in the verification phase.

Rushing attack. Since a reasonable incentive system will allow the participants to be compensated in proportion to their bandwidth, provers can collude with challengers to inflate bandwidth to get more rewards. During the measurement phase, instead of the backhaul link which is filled with the packets from uncorrupted challengers, corrupted challengers transmit packets through an extra link with an extremely high bandwidth to finish the measurements within a shorter time.

Information sharing attack. Besides the rushing attack, another way for the prover to get more information about the data from the challengers than that was transmitted through the backhaul link is to exploit the information structure. In the verification phase, to facilitate the verification of whether the packets received by the prover are indeed those sent by the challengers, the challengers are required to provide the information related to packet generation. If the information to generate packets is much smaller than the actual packet data and is shared to the prover directly, the prover can also terminate the measurements much earlier since it can generate a fraction of packets by itself. For instance, in our protocol, corrupted challengers can send their secret keys to the prover.

DoS and related attacks. Attacks are possible where a challenger sends traffic with invalid signatures, sends duplicate packets, or reports fewer packets in the verification phase. These attacks result in the same outcome: the prover is unable to report its actual bandwidth because unwanted extra traffic was present at the time of challenge. One can view all these attacks as variants of Denial-of-Service (DoS) attack. DoS or distributed Dos (DDoS) attacks are possible in our setting, especially if one implements the protocol by setting public IP address. Specifically, a challenger or a group of them with a high bandwidth link can flood the prover backhaul with invalid packets, preventing the valid packets sent by uncorrupted challengers from reaching the prover. Even a challenger who has not been selected for a particular challenge but knows the time of the challenge can disrupt the challenge similarly. It will be challenging to alleviate the problem using standard filtering techniques [48] without cooperation from the ISP. In fact, if ISP cooperates or the deployment is on a local network, a promising direction for countermeasure will be to adopt a network architecture with accountability[15] and incentivize ISPs to enforce challenge-specific security policies[46]. The detection and dis-incentivization of such attacks (for instance, using crypto-economic “slashing”) is a topic of future work. Furthermore, we assume that there is no attack during the network synchronization phase; in future we can adopt Byzantine resistant time synchronization schemes such as those in [21]

to further enhance security.

B. Security Properties

Theorem 1 (Soundness). *When $f < n/3$, the prover cannot inflate the bandwidth.*

Proof: According to the protocol, all packets with sequence number q sent by uncorrupted challengers will arrive at the network core at $t_0 + q \cdot b/\theta_0$ and be added to the backhaul link queue Q . Because it takes at least b/θ_0 to finish transmitting all packets with the same sequence number (according to bandwidth condition in Eq.2), the queue will never be empty during the measurement phase. Before sending the response, the prover waits for $K \geq (n - f) \cdot k$ packets, among which at most fk packets come from corrupted challengers. These packets can be sent through an external link (rushing attack) or generated by prover directly if the secret keys are shared in collusion (information sharing attack). In either case they will not actually consume the bandwidth of the prover's backhaul link. Even so, there are still at least $K - fk \geq (n - 2f) \cdot k$ packets sent by uncorrupted challengers. Since packets from uncorrupted challengers are not forgeable by anyone else, the earliest time at which the prover can send response is the time at which $(K - fk)$ packets from Q get delivered, which is at least $t_R = t_0 + b/\theta_0 + (K - fk)b/\theta_p$, whereby the uncorrupted challengers will receive the response and time $t_R^i > t_R$.

Since the verifier needs to collect at least $n - f$ time measurements, of which $n - 2f$ must be reported by uncorrupted challengers, the median Δ' of the RTTs must be bounded by the minimum of honest measurements since $f < n/3$, in this way the estimated time will not get affected by individual misreports. Then, denoting the set of honest challengers as H , we have $\Delta' \geq \min\{t_R^i\}_{i \in H} - t_0 - b/\theta_0 > (K - fk)b/\theta_p$ and

$$\theta'_p = \frac{K \cdot b \cdot (n - 2f)}{\Delta' \cdot (n - f)} \leq \frac{K \cdot (n - 2f) \cdot \theta_p}{(K - fk) \cdot (n - f)} \leq \theta_p. \quad \blacksquare$$

Theorem 2 (Approximate completeness). *When $f < n/3$ and the prover is uncorrupted, the protocol will always output bandwidth with accuracy $\alpha = \theta'_p/\theta_p \geq (n - 2f)/(n - f)$.*

Proof: When the prover is uncorrupted, it waits for $(n - f) \cdot k$ packets to generate the response. Even under withholding attacks, where corrupted challengers never send their packets, $(n - f) \cdot k$ packets generated by honest challengers will arrive at the prover before $t_R = t_0 + b/\theta_0 + (n - f)kb/\theta_p$. Then all uncorrupted challengers receive the response at the same time and output to the verifier. Assuming that the size and the latency of the response is negligible, we have the median RTT $\Delta' = (n - f)kb/\theta_p$. If corrupted challengers try to misreport the number of packets received by the prover, claim the proof sent by the prover is incorrect, or even withhold the measurement results, the prover can send the genuine packets it has received from the challenger to the verifier together with the Merkle proof. The verifier will reconstruct the Merkle root from the submitted partial data and Merkle proof to solve disputes. Thus, even under attacks, the total number of packets

are no less than $(n - f)k$. Consequently, the protocol will output

$$\theta'_p = \frac{(n - f)kb \cdot (n - 2f)}{\Delta' \cdot (n - f)} = \frac{n - 2f}{n - f} \theta_p. \quad \blacksquare$$

Remark. (Adversarial threshold.) Our protocol can tolerate up to a fraction 1/3 of Byzantine challengers. This threshold of 1/3 arises from the requirement to ensure that a majority of $(n - f)$ RTT measurements are collected from uncorrupted challengers. This allows the verifier to terminate the collection responsively (or “lazily”) when receiving enough reports without the requirement of a timer. However, if the verifier has access to a timer with desired accuracy (roughly 100ms for us), it can wait for a certain period (determined by maximal network delay and backhaul links transmission delays) to collect the measurements, by which all challengers who fail to send the report are recorded as 0. In this case at least $n - f$ reports of total n reports are from honest challengers, the majority requirement $n - f > f$ means that the protocol is able to tolerate a fraction 1/2 of Byzantine challengers.

V. EXPLORING THE DESIGN LANDSCAPE

In this section, we investigate variants of our protocol and variations in security guarantees when different primitives such as throughput fairness (§ V-A), digital signatures (§ V-B) and extra side-links (§ V-C) are present or absent.

A. Stronger Primitive: Fairness

We first consider a stronger primitive when using TCP packets as our challenge flow, whose congestion control algorithm is known to provide fairness [20]. Fairness guarantees that when n challengers are sending traffic simultaneously with the same rate to a link with capacity B , each should have an average rate of B/n . Under fairness assumption, we modify the packet-based termination rule (prover generates the first response on receiving enough packets) of the protocol to a time-based rule that prover will wait for a fixed amount of time Δ before sending the first response. On receiving the first response, challengers stop timer and measure their average throughput during the measurement period. In this variant, the final output is the robust sum of throughput reports. In particular, we take the median throughput from n collected reports and multiply it by $(n - f)$ to get the estimated bandwidth. The report collection phase lasts long enough to make sure all honest reports get delivered, and the throughput of those who never submit reports is set to a default value 0. Since honest challengers send the challenge traffic at rate $\theta_0 = \theta_p/(n - f)$ as is specified in the protocol, they are guaranteed to share the backhaul bandwidth in a same rate $\theta'_0 \leq \theta_0$ (it is possible that $\theta'_0 < \theta_0$ when more than $n - f$ challengers generate traffic).

Security analysis. Since honest challengers share the same throughput $\theta'_0 \leq \theta_0$, and the median of throughput is bounded by honest reports, the output bandwidth must be no larger than $(n - f)\theta'_0 \leq (n - f)\theta_0 = \theta_p$, thus the soundness holds. When more than $(n - f)$ challengers send the packets, honest rates will become lower than θ_0 . But since anyone who send

more than $\theta_0\Delta$ will be considered as malicious, to stay covert corrupted challengers can only send up to $\theta_0\Delta$ packets during the measurement phase. In other words, the final bandwidth has accuracy $1-\beta$, which is better than the accuracy $(1-2\beta)/(1-\beta)$ of the current protocol (Theorem 2) due to fairness.

B. Weaker Primitive: Without Signature

In our main PoB protocol, digital signatures are used to generate unforgeable packets and resolve disputes between prover and challengers in terms of the number of packets that are indeed received by the prover. We now discuss a case where digital signature schemes are not available, instead, packets are generated by a pseudo random generator with a seed picked by each challenger. We describe the key changes in the protocol below. The protocol guarantees an accuracy of $(1-3\beta)/(1-\beta)$.

Measurement phase. At the beginning of measurement phase, challengers generate a seed S_i and commit it to the verifier. Same as the main protocol, each challenger generates a sequence of k packets with rate θ_0 . The q -th packet from the challenger C_i contains the sequence number q and q -th random number R_{iq} generated from seed S_i . The responses are generated in the same way as Algorithm 1.

Verification phase. After the measurements, each challenger reveals the seed S_i to the prover, who can check whether it matches the pre-committed seed from the verifier. Challengers verify responses by constructing Merkle tree just as Algorithm 2. The key change in the verification phase is the termination check. The verifier accepts the output as long as at least $(n-2f)$ packets are reported as received since f corrupted challengers may refuse to report the actual number of packets received by the prover and the prover can not prove the inconsistency, though the honest prover is required to terminate after receiving $(n-f)k$ packets. As a consequence, the correction factor becomes $(1-3\beta)/(1-2\beta)$.

Security analysis. The soundness proof of the new protocol is similar to the proof in Theorem 1. The only difference caused by the new termination rule is that an honest prover may get only $(n-2f)k$ reported packets while they actually receive $(n-f)k$ packets (corrupted challengers all report 0). Different from digital signatures, random packets generated by pseudo random generator does not keep unforgeability after the seed is revealed, thus there is no way to resolve disputes. However, during the measurement phase, the packets generated by honest challengers are still unforgeable, therefore even corrupted prover needs to receive at least $(n-3f)k$ packets from the honest challengers to generate responses, which means the median time $\Delta' \geq (n-3f)kb/\theta_p$ and

$$\theta'_p \leq \frac{(1-3\beta)(n-2f)kb}{(1-2\beta)(n-3f)kb/\theta_p} = \theta_p$$

The termination is guaranteed for honest provers due to the relaxed termination rule with the cost of lower accuracy. When the prover is uncorrupted, it still waits for $(n-f)k$ packets to generate the response. Under withholding attacks and the misreporting attacks, the honest prover can only prove the receipt of $(n-2f)k$ packets. Therefore, the protocol will output

$$\theta'_p \geq \frac{(1-3\beta)(n-2f)kb}{(1-2\beta)(n-f)kb/\theta_p} \geq \frac{(1-3\beta)}{(1-\beta)}\theta_p$$

So accuracy is reduced to $\alpha = (1-3\beta)/(1-\beta)$.

C. Relaxed Threat Model

In §IV-A, we assume adversary has access to extra high-bandwidth channels to every end node in the Internet, which enables unavoidable rushing attacks. In real world, though, maintaining such a backdoor network can be prohibitively expensive. A more realistic privilege adversary may have is a high-bandwidth channel to the network core. This means packets sent by corrupted participants can reach network core almost instantly, but will be added to the message queue together with other honest packets and be transmitted to destination following the order in queue.

Shuffle phase. With the relaxed threat model, we design a new class of protocols to improve the accuracy ratio. When there are no extra links connecting the adversary with other nodes directly, corrupted challengers can not rush packets through extra links. But information share attack is still possible, to address which we design a shuffle phase to involve more than one challengers into the packets generating process.

We consider a slightly different system where challengers join a pool with shares and the number of shares of a challenger C_i is determined by its available bandwidth θ_i and the bandwidth unit θ_0 . For example, a challenger with bandwidth 50 Mbps owns 5 shares in a pool with bandwidth unit 10 Mbps. Each share is assigned with a unique ID $\{1, \dots, U\}$ (we assume $U \rightarrow \infty$). The Byzantine fraction of corrupted bandwidth units $\beta < \frac{1}{3}$.

During the shuffle phase, t shares are drew from the pool uniformly randomly to form a challenger group, we call such a protocol t -shuffle scheme. Suppose we have n challenger groups sampled from the pool, denoted as C_1, \dots, C_n . Specifically, the owner of each bandwidth unit in the challenger group signs a sequence of packets using the same method in Algorithm 1, and sends these packets to the next challenger in the group, who also signs the packets to get a new sequence of packets. Then in the measurement phase, the last challenger serves as the sender to forward the sequence of packets signed by the entire group to the prover.

In this way, as long as one of the challengers in the group is uncorrupted, the information share attack no longer works since corrupted challengers can only share partial information about how to generate packets to the prover. In other words, only those groups in which all bandwidth units belong to corrupted challengers can still mount the information share attack. On the other hand, the multi-signed packets scheme implies that even one corrupted challenger can withhold the packets. We call a challenger group is good if all bandwidth units in the group belong to honest challengers, and a challenger group is bad if all bandwidth units in the group belong to corrupted challengers. We denote the number of good and bad groups as random variables G and B . For each challenger group C_i ,

$$P(C_i \text{ is good}) = (1-\beta)^t, P(C_i \text{ is bad}) = \beta^t$$

Full protocol. After shuffle phase, the last challengers in all groups forward packets to the prover and start timer for transmission process. The time to send each packet is the same as Algorithm 1. On collecting $cnt \geq g = (\theta_p/\theta_0)k$ packets

(g is a liveness parameter), the prover computes the hash of all packets (from different challengers) as the first response, receiving which challengers will stop timer and log time. In the verification phase, prover need to broadcast all received packets to all challengers so that they can verify the hash and submit reports to the verifier. Then similar to Algorithm 2, verifier waits for time reports from at least $n - f$ challengers, and takes the median of RTT Δ' as time measurement. The final bandwidth of output is given by the following formula.

$$\theta'_p \leftarrow \frac{cnt \cdot p \cdot (g - b)}{\Delta' \cdot g} \quad (3)$$

where g is the liveness parameter, b is the accuracy parameter, p is the packet size.

Security analysis. The parameters defined in equation (3) determines the security and accuracy of PoB with shuffle. Formally, to provide soundness we require

$$\Pr(\theta'_p > \theta_p) = \Pr\left(\frac{g - b}{g - B} > 1\right) = \Pr(B > b) \quad (4)$$

to be negligible. Meanwhile, we want to achieve approximate completeness and maximize the accuracy ratio $(g - b)/g$, which implies $\Pr(G < g)$ should also be negligible to ensure termination. Among n challenger groups, let $G = \sum_{i=1}^n \mathbb{I}[\mathcal{C} \text{ is good}]$, $B = \sum_{i=1}^n \mathbb{I}[\mathcal{C} \text{ is bad}]$ (\mathbb{I} is an indicator function), we have

$$E[G] = (1 - \beta)^t n, \quad E[B] = \beta^t n$$

According to Chernoff bounds, for any $0 < \delta_g \leq 1$ and $0 < \delta_b \leq 1$, we have

$$\begin{aligned} P(G < (1 - \delta_g)(1 - \beta)^t n) &= P(G < g) < \exp(-\Theta(\delta_g^2 n)) \\ P(B > (1 + \delta_b)\beta^t n) &= P(B > b) < \exp(-\Theta(\delta_b^2 n)) \end{aligned}$$

Thus by choosing appropriate g, b , we get accuracy

$$\alpha = \frac{g - b}{g} = 1 - \frac{(1 + \delta_b)\beta^t}{(1 - \delta_g)(1 - \beta)^t}$$

In Appendix E, we set error probability $\epsilon = 0.5$ and $\beta = 1/3$. We search for parameters g, b to reach optimal accuracy ratio α given different number of challenger groups $n \in [10, 500]$. The results are shown in Figure 7.

VI. PROTOCOL IMPLEMENTATION

In this section, we present the protocol implementation in a real system. Towards practicality, we discuss the factors that are not addressed in our theoretical modeling (§VI-A) which leads to the modifications in implementation to the basic form of the protocol (§VI-B).

A. Practical considerations

Challenger bandwidth. In §III-C, we assume each challenger can measure its spare bandwidth θ_i precisely. However, this bandwidth may be time-varying and it will be difficult for the challenger to measure every time. We relax this requirement by allowing every challenger simply ensure that it has at least θ_0 bandwidth available for the challenge. Here $\theta_0 = \theta_p/(n - f)$ is the smallest value that satisfies the bandwidth condition in Eq. (2). Each challenger will now send the challenge traffic at rate θ_0 .

Latency. The key requirement of our technique is that the packets from each challenger reach the prover backhaul at the same time. The aggregation condition Eq. (1) ensures this when there is no synchronization error or latency. However, in practice, a packet from the challenger \mathcal{C}_i will take time l_i to reach the prover, where l_i is the one-way latency from challenger \mathcal{C}_i to the prover. The value of l_i can indeed vary for different challengers and to account for such varying latencies, we modify Eq. (1) as

$$t_0 + \frac{b}{\theta_0} + l_0 = t_{11} + \frac{b}{\theta_0} + l_1 = \dots = t_{n1} + \frac{b}{\theta_0} + l_n$$

where t_{i1} is the start time of challenger \mathcal{C}_i to send the first packet. Note that θ_i is replaced by θ_0 as in our implementation; challengers release packets at rate θ_0 .

Likewise, the response packet from the prover will take time l_i to reach challenger \mathcal{C}_i . Accordingly, Δ_i in Algorithm 1 now changes to $\Delta_i = curTime_i - t_i - 2 \cdot l_i$, where $curTime_i$ is the time when challenger \mathcal{C}_i receives the response from the prover. For measuring l_i , before the challenge starts, each challenger sends 20 ICMP ping packets to the prover and takes the average across these 20 packets as RTT . We set the value l_i as $RTT/2$. Note that using ping packets to measure l_i creates a vulnerability: prover can delay the ping response, inflating l_i and thereby θ'_p . To circumvent this, one possible solution is to restrict the challengers to be within certain geographical limit of the prover. This will ensure that the maximum l_i is bounded by a small value of say 10-15 ms and so the error in estimating θ'_p due to incorrect l_i is also small, provided the challenge duration is chosen to be sufficiently long; see error analysis of implemented protocol in Appendix A.

Packet drops. Our premise is that all k packets from a challenger reach the prover. However, simultaneous packet sending by challengers can cause buffer overflow at the prover's last link, leading to packet drops. As we employ UDP protocol, dropped packets will not be retransmitted. As a result, packet drops can hinder challenge termination. In our experiments, we observed that sending $1.1k$ packets from challengers compensates for packet drops. With a presumed packet drop rate of up to 10%, it assures the prover receives enough packets, leading to termination.

Time synchronization. We require that all the challengers are synchronized via Network Time Protocol (NTP) [13]. Note that NTP does not ensure perfect time synchronization, there can still be residual synchronization errors of the order of tens of milliseconds over the Internet [44].

Computation overhead. The use of cryptographic primitives like *Hash* and *MerkleRoot* (Algorithm 1) inevitably incurs

computation overhead, which will delay the prover from sending responses to challengers and thereby add to an error in measurements. We detail empirical computation times of these primitives in §VII-A as a function of the number of challengers and challenge duration for completeness.

B. Implementation

We implement challengers and the prover as UDP socket applications in C++ and each challenger conducts measurements by sending UDP packets to the prover. Details are described below.

Digital signatures. As per Algorithm 1, each packet needs challenger’s signature. We employ an OpenSSL based implementation of Ed25519 [32], [3], an Edwards-curve digital signature scheme, for its low computational overhead. Our findings reveal that 64 Bytes challenge packets (size of Ed25519 signature) compromise measurement accuracy, particularly over WiFi. For efficiency, modern WiFi uses packet aggregation [49], combining multiple network layer packets into a single MAC layer packet, up to 1MB. Smaller 64 bytes UDP challenge packets hinder WiFi MAC aggregation, decreasing the throughput θ_0 for the challenger. To counter this, we bundle multiple signatures into a single larger packet. We use 1514 bytes challenge packets (1472 byte UDP payload with a 42-byte header), carrying 23 distinct 64 byte signatures. The payload contains only signatures to prevent malicious challengers from omitting plaintext content.

Hashing and verification. As described in Algorithm 1, upon receiving the required number of total packets, the prover generates a hash for each challenger C_i , i.e., $h_{1_i} \leftarrow Hash(M[i])$ where $M[i]$ is the set of all the signatures received from challenger C_i . The prover then generates a *MerkleRoot* of all the hashes from all the challengers. For generating the hash we use *sha256* hash function via the implementation [8] and for generating the Merkle root, we use a C++ open source implementation [9]. The prover sends h_{1_i} and *MerkleRoot* as response to the challenger C_i . The response packet is a UDP packet with a payload of 64 bytes as it contains two 256-bit hashes. In the verification phase (Algorithm 2), the prover sends bitmap B_i and Merkle proof P_i to challenger C_i , who then verifies the Merkle proof and sends RTT Δ_i and number of its packets received by the prover, to the verifier.

Precomputing the signatures. Signature generation incurs computation time too and our benchmarking of OpenSSL implementation [3] of Ed25519 indicates that generating one signature of 64 bytes takes about 50-60 microseconds (μs) on a resource-constrained Linux system consisting of 1 GB of RAM and 1 CPU core. For each packet, a challenger has to generate 23 signatures which will incur a maximum time of $23 * 60 \approx 1.4$ ms. As the signature generation time is more than the packet transmission time of about 1.2 ms even at $\theta_0 = 10$ Mbps, in our implementation challengers precompute all the signatures before the challenge begins. This can be done after the challenger receives the challenge request and while measuring the ping latency l_i .

Benchmarking the technique. Making use of multiple challengers with the additional requirement of security introduces more sources of errors. Particularly, l_i is not a constant and has some jitter. NTP synchronization can result in error of tens of

milliseconds over the Internet. Computation overhead of hash and Merkle tree generation adds delay. Given these sources of error, we evaluate the accuracy as a function of challenger duration and the number of challengers.

VII. EXPERIMENTAL EVALUATION

In this section, we explore the accuracy and robustness of the available backhaul bandwidth measurements through the employment of multiple challengers. To begin with, we compare the measurement results with those of existing tools under the assumption that all participants are operating honestly (§VII-A). Then, we stress test the protocol under Byzantine attacks to assess the security robustness (§VII-B). Lastly, we present a full-stack PoB system and assess the performance of in a real-world setting (§VII-C).

Experimental setup. Our setup consists of a diverse set of challengers in terms of computation capability and geographical location. We carry out experiments with a maximum of ten challengers. The prover has backhaul bandwidth (θ_p) of 250 Mbps enforced by Linux rate limiter `tc`. The details of the prover and different challengers are listed in Table III. Challengers 1-3 are connected to the Internet via WiFi links, while other challengers have a wired link.

	Compute Parameters		Location	RTT (ms)
	RAM (GB)	CPU		
Prover	1	1	AWS X	
Ch. 1-3	12-16	4-8	Y	25
Ch. 4-5	1	1	AWS Z	198
Ch. 6-10	1	1	AWS X	1

TABLE III: Experimental setup details for the prover and challengers. Location of the nodes are in different continents and are anonymized.

A. Performance Evaluation with All Honest Participants

To study how accuracy varies with challenge duration and the number of challengers, we conduct experiments by adjusting the number of selected challengers from 4 to 10 and challenge duration from 25 ms to 200 ms for the prover.

Challenge duration is the time required to transmit the required amount of packets i.e., $(n - f)k$ packets through the prover backhaul. Individual challengers will take longer to complete the challenge due to their latency, l_i , and the fact that they send $1.1k$ packets to account for packet drops. We rate-limit the prover backhaul to 250 Mbps using the Linux utility `tc` [7] and set the bandwidth of each challenger (θ_0) as θ_p/n , where n is the number of challengers.

Fig. 3 shows the backhaul measured by our technique for varying number of challengers and challenge durations as 25, 50, 100 and 200 ms. For each challenge duration, ten experiments are carried out. We plot the average and standard deviation for ten experiments in Fig. 3.

As can be seen from Fig. 3, with the number of challengers set to 4, the measured backhaul is only about 167 Mbps for 25 ms challenge duration, but when the challenge duration is increased to 200 ms, the measured backhaul increases to about 241 Mbps with an error of about 4%. On the other hand, when the number of challengers is increased to 6 or more, the




Fig. 3: Backhaul measured by our technique for different challenge durations. Error bars \approx std. deviation.

measured backhaul has an error of less than 5%, even for 25 ms challenge. However, the standard deviation for 25 ms and 50 ms experiments is higher. The measurement accuracy increases and the standard deviation decreases, if the challenge duration is increased to 100 ms or more. For 100 ms, we observe an error of less than 5% for six or more challengers.

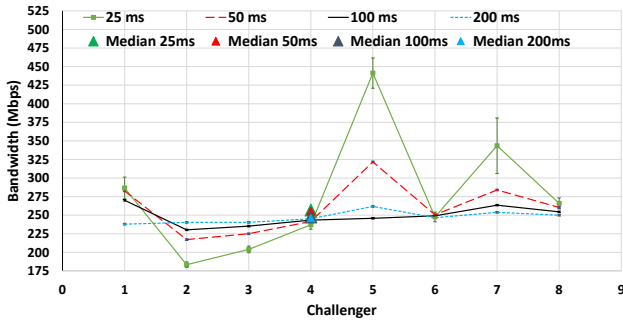


Fig. 4: Backhaul measured by each challenger ($n = 8$) for different challenge durations. Error bars \approx std. deviation.

We look at how measurement accuracy is affected as a function of the challenge duration for the case of $n = 8$. Fig. 4 shows the backhaul measured by each of eight challengers, as an average across ten experiments for different challenge durations. As can be seen from Fig. 4, the backhaul measured by individual challenger shows higher error when challenge duration is 25 ms or 50 ms. For example, the backhaul measured by challenger 5 is about 440 Mbps and 320 Mbps for 25 ms and 50 ms duration. However, the error decreases when challenge duration is increased to 100 ms or more. Note that standard deviation across 10 experiments for each challenger also decreases as the challenger duration is increased.

Some ISPs use token bucket filter to rate limit the backhaul [58], whereby few packets can be released in a burst at much higher rate than the average backhaul speed. Since in our technique we use large number of packets (much more than the burst size of typical token bucket filter), we will still be able to measure the average backhaul which is our goal.

Sources of error. As shown in Fig. 4, some challengers measured the prover’s backhaul as higher than the actual value of 250 Mbps. This is due to errors affecting measurement

accuracy (see §VI) such as time synchronization and jitters in latency. We observe that due to these errors, there is a time difference of 20-30 ms between the first packet from the first and the last challenger reaching the prover backhaul. To compensate for the maximum packet drop rate of 10% (see §VI), each challenger sends 10% more data. So, the challengers that start late might receive the response from the prover before they finish sending their share of challenge packets, if sufficient challenge packets have been received by the prover from the challengers that start early. Such late starting challenger’s backhaul estimate may be higher than the actual value. However, the median evaluation at the final step, which is primarily designed for security, also provides robustness against such outliers. Consequently, our measurement accuracy increases when there are more challengers.

The computation overhead of hash and Merkle tree generation also adds to the measurement error. We observe that the computation overhead for the case of 4 challengers for 25 ms challenge duration is about 500 μ s, while for 10 challengers for 200 ms challenge duration is about 3ms.

Amount of data. For packet drop rate of 10%, the total amount of data required for different challenge duration to measure a 250 Mbps backhaul link is given in Table IV.

Expt. Duration (ms)	25	50	100	200
Data (MB)	0.86	1.71	3.44	6.88

TABLE IV: Amount of challenge data required.

As seen from Fig. 4, for challenges with 100-ms duration we get a good accuracy for each challenger. Thus, our results show that our technique can measure 250 Mbps backhaul in 100 ms with about 3.5 MB of data and an error of less than 5%, when 6 or more uncorrupted challengers are involved.

Comparison with a single challenger. With a single challenger that has a bandwidth of 250 Mbps, we could measure prover backhaul of 250 Mbps with less than 2% error with challenge duration being only 10 ms and the amount of data required is about 345 KB. Multichallenger technique requires larger challenge duration due to the aforementioned errors. As the duration of the challenge is longer, the amount of data used correspondingly increases. But the primary benefit of multichallenger technique is that each challenger requires much smaller bandwidth. With ten challengers, each challenger requires a bandwidth of only 25 Mbps to measure prover backhaul of 250 Mbps.

Accuracy for larger prover backhauls. The experimental results of accuracy for larger prover backhauls (500 Mbps to 1000 Mbps) with 10 challengers are tabulated in Table V. We observe that the measurement error grows as prover backhaul increases; however even for prover backhaul of 1000 Mbps, the measurement error is less than 8%.

Backhaul (Mbps)	500	750	1000
Measured BW (Mbps)	474.7	705.4	921.4
Error (%)	5.1	5.9	7.9
Overhead (ms)	4.6	7.3	10.2

TABLE V: Measured bandwidth and computation overhead for larger prover backhauls.

One reason for higher error as prover backhaul increases is the increasing computation time to construct Merkle tree. The last row of Table V shows the computation overhead for various prover backhauls. The computation overhead for a prover with 1000 Mbps is about 10 ms, 10% of the challenge duration. The experiments show that larger prover backhaul leads to higher computation overhead. So, for backhauls over 1000 Mbps, the challenge duration should be extended.

Cross traffic on the backhaul link is another factor potentially affecting measurements and possibly leading to protocol non-termination. To mitigate this issue, we propose a protocol patch. Its effectiveness, evaluated under cross-traffic conditions, is detailed in Appendix B.

Comparison with existing techniques. We compare the performance of our technique with pathchar [22], MagicTrain [59] and speedtest [6]. Pathchar uses round trip delays, while MagicTrain uses packet dispersion and speedtest uses bulk TCP transfer to measure bandwidth. Table VI illustrates the results for backhaul of 500 Mbps.

Technique	PoB	pathchar	MagicTrain	speedtest
Measured BW (Mbps)	474.7	101.3	152.1	492.3
Amount of data (MB)	6.88	1.2	0.1	501

TABLE VI: Measurement accuracy of existing techniques.

As evident from Table VI, although pathchar and MagicTrain use less data, they suffer from poor accuracy. Speedtest, while more accurate, consumes about fifty times more data. Recent tools [57], [56] improve speedtest’s data usage, but we couldn’t execute their codes. Moreover, these techniques lack Byzantine tolerance and necessitate their measurement servers to be proximate to the end device.

The poor performance of pathchar and MagicTrain is due to their reliance on precise timing measurements. We empirically demonstrate the reasons in Appendix C. Similar factors account for the inadequate performance of MagicTrain.

B. Security Evaluation

We carry out experiments to study how measurement results are effected in the presence of malicious challengers. We choose $n = 10$ and $f = 2$. We carry out measurements for two different prover backhauls of 100 Mbps and 250 Mbps, with a challenge duration of 100 ms.




Fig. 5: Backhaul measured by each challenger in case of withholding and rushing attack. Prover backhaul is 250 Mbps.

Fig. 5 shows the average bandwidth measured by each challenger across ten experiment runs for the case when prover backhaul is 250 Mbps. As can be seen from Fig. 5, the measured backhaul in the case when all challengers are honest (honest median in Fig. 5) is 251.5 Mbps, while the measured backhaul in the case of withholding attack (Withhold Median in Fig. 5) is 241.4 Mbps. The expected measured backhaul in case of withholding attack is 250 Mbps. So the measurement accuracy in case of withholding attack is within 4%. Note that this is the accuracy of the measurement technique. Our PoB protocol will apply a correction factor $\alpha = (n - 2f)/(n - f)$ (Algorithm 2) and output the guaranteed bandwidth for the prover as $241.4\alpha \approx 181$ Mbps which is about 28% less than the prover backhaul of 250 Mbps.

In case of a rushing attack, the measured bandwidth is inflated to $250/\alpha \approx 333$ Mbps. The measured backhaul (rushing median in Fig. 5) 331.5 Mbps matches our theoretical prediction. The output of the protocol in this case will be $331.5\alpha \approx 249$ Mbps, close to the actual prover backhaul.

Note that our security guarantees require us to curtail bandwidth inflation. Indeed, we can observe that even under a rushing attack, the guaranteed bandwidth does not exceed the actual bandwidth. This is enabled by multiplying by a shrinkage factor to compensate for adversarial challengers trying to help the prover to claim an inflated bandwidth. However, this comes at the cost of lower guaranteed bandwidth even when all challengers are reporting honestly. After repeating the experiment for a backhaul of 100 Mbps, the results stay similar, validating our theoretical predictions.

C. A Full-stack PoB System

We have developed a full-stack implementation of our PoB framework. This refined code base facilitates the deployment of provers and challengers, and is available for reference [5]. Crucially, we have achieved integration with Ethereum [1] blockchain. Briefly, the system enables a *payer* to initiate a PoB challenge using blockchains. This request entails a payment in test tokens of the underlying blockchains. For each specific challenge, we randomly select a group of challengers, who are subsequently rewarded with the test tokens provided by the payer at the end of the challenge. The details of the system architecture, workflow and smart contract interfaces are presented in Appendix D.

Our system is alive and allows anyone who has a Ethereum wallet to join as a challenger by simply deploying our code. At the time of writing, we have approximately 25-30 active challengers situated in diverse geographic regions including the US, Europe, and Asia.

Measurements from real setting. For a better understanding of the system’s practical performance, Table VII presents the measurement results for three different backhauls of 500, 700 and 1000 Mbps.

Backhaul (Mbps)	500	700	1000
Measured BW [stdev] (Mbps)	478.8 [11.9]	671.5 [13.6]	901.0 [16.0]
Average Error	4.2%	4.1%	9.9%

TABLE VII: Measurement results in a real setting.

Each measured backhaul in Table VII is average across ten challenges carried out with different set of challengers and provers at different times. The average measurement error is about 4% for backhaul of 500 and 700 Mbps and about 10% for backhaul of 1000 Mbps, similar to our testbed experiments.

VIII. CONCLUSION AND DISCUSSION

Summary. Trustfree telemetry is a central problem in decentralized networks. Our Proof of Backhaul protocol addresses a core requirement by providing a secure and accurate backhaul bandwidth measurement service for wireless access points while also allowing open participation. The protocol is operated by a group of challengers, whose latency and bandwidth can be ordinary, with the goal of measuring a prover hotspot who may have a high-bandwidth backhaul link. We have established a trust model for the PoB problem, designed precise specifications of the PoB protocol, and tested a high-performance, low-overhead implementation.

Improving accuracy in practice. Our protocol guarantees soundness and completeness of backhaul measurements with a reasonable accuracy in the presence of Byzantine parties. The accuracy ratio $(1 - 2\beta)/(1 - \beta)$ is determined by the Byzantine fraction due to a correction made for an unavoidable rushing attack – corrupted challengers can always rush their packets through an external high-bandwidth link to lower RTT and inflate backhaul bandwidth to be measured. However, such backdoor links may incur substantial costs in practice, necessitating a more relaxed threat model and a family of extended protocols. Without rushing links, we equip PoB protocols with a shuffle phase where a pair of challengers are asked to jointly sign packets. This mechanism improves accuracy by making information sharing attack harder in a probabilistic manner, with a cost of higher communication overhead for verification. Implementing an efficient system for such a shuffle protocol in practice is an active area of research.

Cross traffic. In our proposed method for handling cross-traffic in Appendix B, we run experiments for increasing values of bandwidth below the claimed link capacity. This requires fresh data to be sent for each value and increases the amount of data needed. To reduce the data consumption, a naive approach could be that the prover replies to the challengers with the number of packets received in a fixed duration. An alternative approach is to send intermediate responses when appropriate amounts of data are received. Both approaches cannot guarantee a fixed accuracy for different available bandwidths. Designing a protocol which is more data efficient and has such guarantees is an open problem.

ACKNOWLEDGMENT

This work was conducted when all authors were working at Witness Chain (<https://witnesschain.com/>).

REFERENCES

[1] “Ethereum,” <https://ethereum.org/en/>, [Online; accessed 26-June-2023].
 [2] “Fast internet speed test.” <https://fast.com/>, [Online; accessed 26-June-2023].
 [3] “Openssl ed25519 implementation,” <https://www.openssl.org/docs/man1.1.1/man7/Ed25519.html>, [Online; accessed 26-June-2023].

[4] “PM-WANI Central Registry,” <https://pmwani.gov.in/wani>, [Online; accessed 26-June-2023].
 [5] “proof-of-backhaul,” <https://github.com/multichallengerpob/proof-of-backhaul>, [Online; accessed 26-June-2023].
 [6] “Speedtest,” <https://www.speedtest.net>, [Online; accessed 26-June-2023].
 [7] “Tc - traffic control, linux manual,” <https://man7.org/linux/man-pages/man8/tc.8.html>, 2001, [Online; accessed 18-September-2022].
 [8] “C++ sha256 implementation,” <http://www.zedwood.com/article/cpp-sha256-function>, 2012, [Online; accessed 26-June-2022].
 [9] <https://github.com/IAIK/merkle-tree>, 2015, [Online; accessed 26-June-2023].
 [10] “Celona: Private wireless for the enterprise,” <https://www.celona.io/>, 2019, [Online; accessed 26-June-2023].
 [11] “Ericsson private 5g: Private network for your industry,” <https://www.ericsson.com/en/private-networks/ericsson-private-5g>, 2019, [Online; accessed 26-June-2023].
 [12] “Take your business into the digital age with private wireless,” <https://www.nokia.com/networks/private-wireless-network>, 2019, [Online; accessed 26-June-2023].
 [13] “Ntp: The network time protocol,” <http://www.ntp.org/>, 2020, [Online; accessed 26-June-2023].
 [14] “Multi-server testing,” <https://www.ookla.com/articles/how-ookla-ensures-accurate-reliable-data-2021>, 2021, [Online; accessed 26-June-2023].
 [15] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker, “Accountable internet protocol (aip),” in *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, 2008, pp. 339–350.
 [16] S. Banerjee and A. K. Agrawal, “Estimating available capacity of a network connection,” in *Proceedings IEEE International Conference on Networks 2000 (ICON 2000). Networking Trends and Challenges in the New Millennium*. IEEE, 2000.
 [17] R. L. Carter and M. E. Crovella, “Dynamic server selection using bandwidth probing in wide-area networks,” Boston University Computer Science Department, Tech. Rep., 1996.
 [18] —, “Measuring bottleneck link speed in packet-switched networks,” *Performance evaluation*, vol. 27, pp. 297–318, 1996.
 [19] I. Castro, A. Panda, B. Raghavan, S. Shenker, and S. Gorinsky, “Route Bazaar: Automatic Interdomain Contract Negotiation,” in *15th Workshop on Hot Topics in Operating Systems (HotOS XV)*. Kartause Ittingen, Switzerland: USENIX Association, May 2015. [Online]. Available: <https://www.usenix.org/conference/hotos15/workshop-program/presentation/castro>
 [20] D.-M. Chiu and R. Jain, “Analysis of the increase/decrease algorithms for congestion avoidance in computer networks. j-comp-net-isdn, 17 (1): 1–14,” 1989.
 [21] J. Douceur, J. Howell, and J. J. Douceur, “Scalable byzantine-fault-quantifying clock synchronization,” 2003.
 [22] A. B. Downey, “Using pathchar to estimate internet link characteristics,” *ACM SIGCOMM Computer Communication Review*, vol. 29, no. 4, pp. 241–250, 1999.
 [23] FCC, “Title 47, Chapter I, Subchapter D, Part 96, Citizens Broadband Radio Service,” Federal Communications Commission, Regulatory Information, Oct. 2020. [Online]. Available: <https://www.govinfo.gov/content/pkg/CFR-2020-title47-vol5/pdf/CFR-2020-title47-vol5-part96.pdf>
 [24] M. Ghosh, M. Richardson, B. Ford, and R. Jansen, “A torpath to torcoin: Proof-of-bandwidth altcoins for compensating relays,” NAVAL RESEARCH LAB WASHINGTON DC, Tech. Rep., 2014.
 [25] A. Haleem, A. Allen, A. Thompson, M. Nijdam, and R. Garg, “Helium: A Decentralized Wireless Network,” Helium Systems, Inc., White Paper, Nov. 2018. [Online]. Available: <http://whitepaper.helium.com>
 [26] K. Harfoush, A. Bestavros, and J. Byers, “Measuring bottleneck bandwidth of targeted path segments,” in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies*, 2003.
 [27] N. Hu and P. Steenkiste, “Evaluation and characterization of available

- bandwidth probing techniques,” *Journal on Selected Areas in Communications*, vol. 21, no. 6, pp. 879–894, 2003.
- [28] V. Jacobson, “Traceroute,” <https://linux.die.net/man/8/traceroute6>.
- [29] —, “Pathchar,” <ftp://ftp.ee.lbl.gov/pathchar/>, 1999.
- [30] M. Jain and C. Dovrolis, “End-to-end available bandwidth: Measurement methodology, dynamics, and relation with tcp throughput,” in *ACM SIGCOMM Computer Communication Review*, 2002.
- [31] —, “Pathload: A measurement tool for end-to-end available bandwidth,” in *In Proceedings of Passive and Active Measurements (PAM) Workshop*. Citeseer, 2002.
- [32] S. Josefsson and I. Liusvaara, “Edwards-curve digital signature algorithm (eddsa),” Tech. Rep., 2017.
- [33] G. O. Karame, B. Danev, C. Bannwart, and S. Capkun, “On the security of end-to-end measurements based on packet-pair dispersions,” *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 1, pp. 149–162, 2012.
- [34] S. Keshav, “A control-theoretic approach to flow control,” in *Proceedings of the conference on Communications architecture and protocols*, 1991.
- [35] H. Kim and N. Feamster, “Improving network management with software defined networking,” *IEEE Communications Magazine*, vol. 51, no. 2, pp. 114–119, 2013.
- [36] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, “Software-defined networking: A comprehensive survey,” *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2014.
- [37] K. Lai and M. Baker, “Measuring link bandwidths using a deterministic model of packet delay,” in *Proceedings of the conference on applications, technologies, architectures, and protocols for computer communication*, 2000, pp. 283–294.
- [38] —, “Nettimer: A tool for measuring bottleneck link bandwidth,” in *3rd USENIX Symposium on Internet Technologies and Systems (USITS 01)*, 2001.
- [39] “Magma: A modern mobile core network solution,” magma Core Foundation. [Online]. Available: <https://magma-core.org>
- [40] B. A. Mah, “pchar: A tool for measuring internet path characteristics,” <http://www.employees.org/~bmah/Software/pchar/>, 2000.
- [41] B. Melander, M. Bjorkman, and P. Gunningberg, “A new end-to-end probing and analysis method for estimating bandwidth bottlenecks,” in *Globecom’00-IEEE. Global Telecommunications Conference*. IEEE, 2000.
- [42] —, “Regression-based available bandwidth measurements,” in *International Symposium on Performance Evaluation of Computer and Telecommunications Systems*, 2002.
- [43] R. C. Merkle, “A digital signature based on a conventional encryption function,” in *Conference on the theory and application of cryptographic techniques*. Springer, 1987, pp. 369–378.
- [44] D. L. Mills, “On the accuracy and stability of clocks synchronized by the network time protocol in the internet system,” in *ACM SIGCOMM Computer Communication Review*, 1989.
- [45] “ORAN: Transforming the Radio Access Networks Towards Open, Intelligent, Virtualized and Fully Interoperable RAN,” o-RAN Alliance e.V. [Online]. Available: <https://www.o-ran.org>
- [46] C. Pappas, R. M. Reischuk, and A. Perrig, “Fair: Forwarding accountability for internet reputation,” in *2015 IEEE 23rd International Conference on Network Protocols (ICNP)*. IEEE, 2015, pp. 189–200.
- [47] A. Pasztor and D. Veitch, “Active probing using packet quartets,” in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*, 2002, pp. 293–305.
- [48] T. Peng, C. Leckie, and K. Ramamohanarao, “Protection from distributed denial of service attacks using history-based ip filtering,” in *IEEE International Conference on Communications, 2003. ICC’03.*, vol. 1. IEEE, 2003, pp. 482–486.
- [49] E. Perahia and R. Stacey, *Next generation wireless LANs: 802.11 n and 802.11 ac*. Cambridge university press, 2013.
- [50] R. Prasad, C. Dovrolis, M. Murray, and K. Claffy, “Bandwidth estimation: metrics, measurement techniques, and tools,” *IEEE network*, vol. 17, no. 6, pp. 27–35, 2003.
- [51] V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell, “pathchirp: Efficient available bandwidth estimation for network paths,” in *Passive and active measurement workshop*, 2003.
- [52] K. M. Salehin and R. Rojas-Cessa, “Packet-pair sizing for controlling packet dispersion on wired heterogeneous networks,” in *2013 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 2013, pp. 1031–1035.
- [53] R. Snader and N. Borisov, “Eigenspeed: secure peer-to-peer bandwidth evaluation,” in *IPTPS*, 2009, p. 9.
- [54] J. Strauss, D. Katabi, and F. Kaashoek, “A measurement study of available bandwidth estimation tools,” in *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. ACM, 2003.
- [55] J. Tremback and J. Kilpatrick, “Althea: An incentivized mesh network protocol,” Althea Network, Inc., White Paper, May 2017. [Online]. Available: <https://github.com/althea-net/althea-whitepaper/blob/master/whitepaper.pdf>
- [56] X. Yang, H. Lin, Z. Li, F. Qian, X. Li, Z. He, X. Wu, X. Wang, Y. Liu, Z. Liao *et al.*, “Mobile access bandwidth in practice: measurement, analysis, and implications,” in *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022, pp. 114–128.
- [57] X. Yang, X. Wang, Z. Li, Y. Liu, F. Qian, L. Gong, R. Miao, and T. Xu, “Fast and light bandwidth testing for internet users,” in *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, 2021, pp. 1011–1026.
- [58] E. Zhang and L. Xu, “Capacity and token rate estimation for networks with token bucket shapers,” *Computer Networks*, vol. 88, pp. 1–11, 2015.
- [59] P. Zhou, R. K. Chang, X. Gu, M. Fei, and J. Zhou, “Magic train: design of measurement methods against bandwidth inflation attacks,” *IEEE Transactions on Dependable and Secure Computing*, vol. 15, no. 1, pp. 98–111, 2015.

APPENDIX

A. Security Analysis of Implemented Protocol

Taking practical factors such as packet drop rate bound r and an overall latency bound L (including time synchronization error, transmission latency and computation overhead) into consideration, we analyze the security properties of our implemented protocol below.

Theorem 3 (Soundness). *In implemented protocol, by adjusting the correction factor to $(n - 2f - rf)/(n - f) - L\theta_0/(kb)$, the prover cannot inflate the bandwidth when $f < n/3$.*

Proof: When at least $n - f$ honest challengers send packets at rate θ_0 , according to the bandwidth condition in Eq.2, the backhaul link will never be empty. To account for packet drops, challengers are required to send $(1 + r)k$ packets, which means at most $(1 + r)fk$ packets can be sent by corrupted challengers without actually consuming the bandwidth (in rushing attack or information sharing attack). Thus, the number of packets sent by uncorrupted challengers is at least $(n - 2f - rf) \cdot k$.

For handling latency, note that the median Δ' of the RTTs is at least $(n - 2fk)b/\theta_p - L$, whereby

$$\begin{aligned} \theta'_p &\leq \frac{(n - f)kb}{(n - 2f - rf)kb/\theta_p - L} \cdot \frac{(n - 2f - rf)kb - L(n - f)\theta_0}{(n - f)kb} \\ &\leq \theta_p. \end{aligned}$$

■

Theorem 4 (Approximate completeness). *When $f < n/3$ and the prover is uncorrupted, the implemented protocol*

will always output bandwidth with accuracy $\alpha' = \theta'_p/\theta_p \geq \frac{(n-2f-rf)kb/(n-f)-L\theta_0}{(1+r)kb+L\theta_0}$.

Proof: In presence of latency and packet drops, the maximum median RTT is $\Delta' = (n-f)(1+r)kb/\theta_p + L$. But the total number of packets are still no less than $(n-f)k$. Consequently, the bandwidth of honest challenger output by the protocol is at least

$$\begin{aligned} \theta'_p &= \frac{(n-f)kb}{(n-f)(1+r)kb/\theta_p + L} \cdot \frac{(n-2f-rf)kb - L(n-f)\theta_0}{(n-f)kb} \\ &\geq \frac{(n-2f-rf)kb - L(n-f)\theta_0}{(n-f)(1+r)kb + L(n-f)\theta_0} \theta_p \\ &= \frac{(n-2f-rf)kb/(n-f) - L\theta_0}{(1+r)kb + L\theta_0} \theta_p. \end{aligned}$$

■

B. Effect of cross traffic

Our PoB protocol terminates when $(n-f)k$ packets are received by the prover. The number of packets k sent by each challenger is determined by the prover backhaul and challenge duration. However, if there is cross-traffic, the available bandwidth at the prover will be reduced and the challenge packets may experience a larger drop rate than 10% that we assume for our experiments. In this situation the experiment may not terminate.

We propose a modification to measure the available bandwidth in the presence of cross traffic, up to a fixed accuracy δ . The protocol repeats the basic PoB protocol, but instead of verifying θ_p , it verifies iteratively $\theta_p^{(1)} = \delta, \theta_p^{(2)} = 2\delta, \dots, \theta_p^{(\ell)} = \ell\delta$ and so on till θ_p . In more detail, we proceed as follows.

- 1) At step i , execute multichallenger PoB protocol with $\theta_p^{(i)} = i\delta$. Note that each challenger must release challenge traffic at rate $\theta_0 = \theta_p^{(i)}/n$ at this step.
- 2) Each challenger sets a timeout of $5\times$ (challenge duration). If the response from the prover is not received during this period, the challenger declares not terminate. If majority of the challengers declare not terminate, we say that the protocol does not terminate.
- 3) If the protocol for the i th step terminates, increment $i \leftarrow i + 1$ and repeat the steps above.
- 4) Else if the protocol for the i th step does not terminate, output the bandwidth obtained in the previous execution of the PoB protocol.

We assume that the amount of cross traffic does not vary throughout all the steps of the experiment. Using the approach above, we carried out experiments to measure available bandwidth in the presence of different amounts of cross-traffic. The backhaul of the prover is set to 250 Mbps and the number of challengers is 10. $\theta_{P,0}$ is set to 40 Mbps and δ to 20 Mbps.

Table VIII summarizes the results. The measured bandwidths are close to available bandwidths, except for 90 Mbps. The likely reason for this discrepancy is that the presence of challenge traffic reduces cross traffic, resulting in an overestimation of the available bandwidth.

Available BW (Mbps)	220	140	90
Measured BW (Mbps)	219.6	144.6	104.1

TABLE VIII: Measured bandwidth in the presence of cross traffic.

C. Measurements for Pathchar

As we outlined earlier, the performance of pathchar depends on how robust are minimum delay estimates over the Internet and how long will it take for us to get a robust estimate. Thus, to evaluate the performance of pathchar, we measure the RTT to the prover node using ping for 15 different packet sizes in multiples of 100 Bytes, starting from 100 Bytes and ending at 1500 Bytes. We ran the experiment five times and took 500 measurements for each packet size. Pathchar [29], [22] suggests taking the minimum RTT for each packet size and fitting the linear least squares line to the data.




Fig. 6: Minimum round trip time to the prover versus packet size in Bytes. The line shows the linear least squares fit.

Expt.	1	2	3	4	5
RTT	16.589	16.511	16.635	16.565	16.500

TABLE IX: Linear least squares line's intercept values for each experiment's dataset. RTT is in *ms*.

Figure 6 shows the minimum RTT (ms) versus packet size in bytes and the fitted line for the first experiment. Table IX shows the y-intercepts for five experiments; note that the y-intercept represents latency. We can see that the y-intercepts have a difference of 50-100 ms. Thus, we can say that the jitters experienced over the internet is not negligible; in particular, we cannot estimate the minimum latency below accuracy of 50-100 microseconds. Consequently, it is not feasible to use pathchar to measure 100 Mbps or higher bandwidth.

D. PoB System Design

We construct the PoB framework utilizing Ethereum as the verifier blockchain. Beyond the basic roles of challengers and provers, our system incorporates several auxiliary roles delineated in Table X. Specifically, payers are the entities that

TABLE X: Parties involved in PoB system

Party	Descriptions
Payer	A party who pays for the challenge
Prover	The end-point whose backhaul capacity is being measured
Challengers	A pool of servers which can send challenge traffic to the prover
Challenge coordinator	Centralized services for (i) communication between the parties; (ii) computing challenge meta data; and (iii) interacting with the ledger
Blockchain full nodes	Parties maintain a decentralized ledger for recording all the challenge requests and outcomes

cover the expense of the challenge, and challenge coordinators facilitate the coordination among challengers, provers, and blockchain full nodes. The responsibility of a challenge coordinator includes maintaining an active list of challengers, selecting a subset of them randomly for each challenge, and interacting with the blockchain to record challenge-related data. Importantly, although a poorly performing challenge coordinator can compromise the system’s liveness, it does not affect its accuracy or soundness. In practice, we encourage system liveness through commission fees and by establishing a pool of coordinators.


We have categorized the primary stages of the protocol as follows:

- **Challenge Initiation.** The *payer* triggers an on-chain request via the `STARTCHALLENGE` function. This function submits the public keys of both the payer and the prover, and the claimed capacity. To accommodate the expenses related to the challenge (which fluctuate with the claimed capacity), the payer deposits escrow tokens into the contract.
- **Challenge Setup.** The *challenge coordinator* operates a blockchain watcher to track the contract’s states. Upon the submission of a new request, it employs the `GETPOB` function to extract the request information from the chain and selects a suitable number of challengers from its list of active challengers. The challenge coordinator subsequently informs the prover and the chosen challengers of the impending challenge.
- **Challenge Execution.** Each challenger calculates the RTT to the prover. The prover provides instructions to each challenger regarding the location and the timing for the transmission of challenge data. In accordance with the protocol, the challengers proceed to dispatch the challenge data. The prover, after collecting a sufficient number of packets, computes a response hash from the received data and shares it with each challenger. The challengers record both the initiation time and the time of response hash receipt. The prover then sends signed verification data to each challenger, which allows them to calculate the transmission time and the count of received packets. This data is passed on to the challenge coordinator, who determines the bandwidth and uploads all the data (including a list of challengers and their responses) on-chain using the `ENDCHALLENGE` function. Post data verification, the PoB contract is terminated and the escrowed tokens are distributed among the challengers as a token of appreciation for their contribution to the challenge.

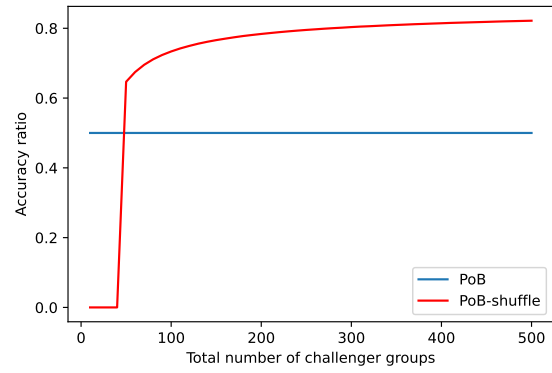
E. Practicality Analysis of PoB-Shuffle

In order to evaluate the practical application of PoB-Shuffle, we establish a Byzantine fraction $\beta = 1/3$ and designate a maximum error probability $\epsilon = 0.5$. We then compare four different variations of the t -shuffle protocols, characterized by varying challenger group sizes of $t = 2, 3, 4, 5$. We employ a binary search of parameters g, b to calculate the optimal accuracy ratio under differing quantities of challenger groups.


Our results indicate that a larger number of challenger groups n enhances the accuracy ratio across all three cases. However, the rate of improvement declines over n until the accuracy eventually converges and the limit can be improved by choosing larger t . Furthermore, a larger t improves measurement accuracy with the same number of groups, albeit at the expense of inherently larger latency. Also, the minimum group number requirements to maintain achieve the specific error probability will be increased. For $t = 3, 4, 5$, there exist no feasible parameters to attain the required error probability when the number of challenger groups is small. Given our choice of a non-negligible error probability, it is necessary to repetitively execute the protocol to obtain more dependable results. In conclusion, PoB-shuffle protocol can achieve better accuracy compared to the constant ratio of PoB when we have sufficiently large number of challenge groups.



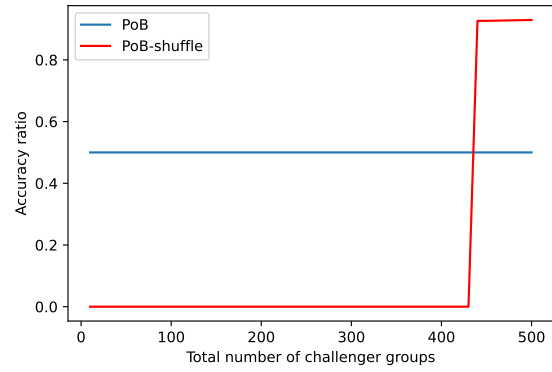
(a) $t = 2$



(b) $t = 3$



(c) $t = 4$



(d) $t = 5$

Fig. 7: The comparison of accuracy ratio between PoB and t -shuffle protocols. Byzantine ratio $\beta = 1/3$, we calculate the best accuracy ratio under different number of challenger groups n . The maximal failure probability of PoB-shuffle protocol is $\epsilon = 0.5$.