

From Onions to Shallots: Rewarding Tor Relays with TEARS

Rob Jansen¹, Andrew Miller², Paul Syverson¹, and Bryan Ford³

¹ U.S. Naval Research Laboratory, Washington, DC
{rob.g.jansen, paul.syverson}@nrl.navy.mil

² University of Maryland, College Park, MD
amiller@cs.umd.edu

³ Yale University, New Haven, CT
bryan.ford@yale.edu

Abstract. The Tor anonymity network depends on volunteers to operate relays, and might offer higher bandwidth with lower response latencies if more users could be incentivized to contribute relay bandwidth. We introduce TEARS, a system rewarding useful service with traffic priority. TEARS audits relays and rewards them with anonymous coins called *Shallots*, proportionally to bandwidth contributed. Shallots may be redeemed anonymously for *PriorityPasses*, which in turn may be presented to relays to request traffic priority. The *PriorityPass* construction enables relays to prevent double spending locally without leaking information. Unlike previous incentive proposals, TEARS incorporates *transparent and distributed banking* using protocols from distributed digital cryptocurrency systems like Bitcoin. Shallots are publicly-verifiable, minimizing reliance on and trust in banking authorities, making them auditable while naturally distributing bank functionality and associated overhead. Further, these distributed banking protocols resist denial-of-service attacks and can recover from catastrophic failures. TEARS may either be deployed in the existing Tor network or operate alongside it.

1 Introduction

Tor is the most popular deployed anonymous communication system, currently transferring over 8 GiB/s in aggregate [32]. The bandwidth Tor requires is donated by altruistic volunteers without any direct return on their investment. As a result, Tor has primarily grown through the use of social and political means. However, utilizing volunteer resources without providing incentives to contribute may not be an adequate long-term growth strategy. How to recruit bandwidth providers while maintaining anonymity is a well studied problem [1, 7, 15–17, 25, 26]. For several reasons, however, a number of practical changes to Tor would be required before adopting any incentive-based resource model; none of these have yet been implemented.

In this paper, we draw upon the distributed Bitcoin architecture to design a **transparent, efficient, and auditable reward system** (TEARS) for Tor that

rewards relays for providing useful service to the network. The TEARS architecture uses existing anonymous Bitcoin protocols to create a type of “altcoin” token system that Tor may utilize internally to provide a means to obtain priority service, and to provide an incentive to contribute bandwidth or other useful services. In TEARS, relays’ bandwidth contributions are audited and earn them Shallots, anonymous but auditable electronic cash. Relays may redeem their Shallots for relay-specific PriorityPasses, which they may “spend” as clients to request traffic priority [10]. Relays may also privately transfer their Shallots to any user wishing to obtain PriorityPasses. Accountability of Shallots is managed by a distributed and semi-trusted bank using a transparent, publicly auditable banking protocol [19].

Previous Tor incentive schemes that propose payments for service [1, 7] trade off the speed at which double spending may be detected for the amount of information leaked through withdrawal and deposit transactions. In TEARS, relay-specific PriorityPasses allow relays to prevent double spending immediately and locally, without leaking information. Further, PriorityPasses in TEARS are non-transferable and become useless after they are presented, reducing overhead while ensuring that the process of requesting priority does not significantly degrade anonymity—the act of redeeming Shallots for PriorityPasses will be unlinkable to any later transaction.

To the best of our knowledge, we are the first to describe how to distribute the bank among the Tor directory servers and utilize it in a way that protects the anonymity of Tor users. The bank in TEARS uses auditable distributed electronic cash protocols [22, 29] to realize transparent accountability of relay rewards while minimizing trust in authorities. Previous schemes assume the existence of a central bank to manage the dissemination of relay rewards [15, 16].

In addition to outlining the design of a new Tor incentive system based on a distributed bank, we offer extended commentary on some system design decisions common to many Tor incentive systems in order to inform new research. This paper also discusses the social challenges involved in the deployment of TEARS and other radical design changes. We hope to illuminate the challenges and research problems in a way that will provoke useful discussion in the community while facilitating future research in this area.

Section 2 first identifies basic requirements for an incentive design. Section 3 outlines the TEARS architecture, then Section 4 discusses the potential impacts of design decisions on the existing network and its operators. In Section 5 we outline previous Tor incentive proposals and their shortcomings. Finally, we summarize open research problems and conclude in Section 6.

2 Requirements

A Tor incentive system that rewards volunteers for providing useful network services presents numerous requirements and challenges, both technical and social. This section identifies the major technical problems we believe such a system should solve, keeping in mind that solutions should operate “out-of-band” to minimize interference with Tor’s low-latency communication design. See Sec-

tion 4 for a discussion of social challenges involved in deploying a Tor incentive system, and Appendix A for a discussion of useful services.

Preservation of Anonymity The system should not introduce new attacks on the anonymity offered by Tor. As such, the mechanisms used in our reward system should not allow users to be linked to their Tor network usage.

Rewards for Providing Service Members that contribute to the system should be rewarded in proportion to the value of their contribution. Ideally, the reward should have a backing value, in that a member can use it to achieve a desirable system service or attribute.

Proofs of Useful Service A system that rewards members for serving the network should be able to prove that those services were in fact rendered. Ideally, the proofs of service would be publicly verifiable so that any member of the distributed network could validate the utility provided by any other member.

Publicly Auditable Accounting The system should provide a way to account for the rewards obtained by each member, and provide facilities for the exchange of rewards. Double spending of rewards should be prevented or detected and handled. Ideally, the process of accounting for rewards should be publicly verifiable in order to minimize trust in centralized entities.

Deployability The system should integrate well with the existing Tor network in order to maximize the utility provided to existing users and leverage the existing infrastructure and community. Ideally, the system components will be modular so they may be deployed incrementally.

3 Architecture

We now discuss how we realize a practical system for rewarding Tor relays.

3.1 TEARS Overview

TEARS embodies incentives in tokens called Shallots. To obtain Shallots, relays simply forward traffic as usual: their bandwidth contributions are audited and used to reward them with Shallots. Users may redeem Shallots for PriorityPasses, which they may present to Tor relays for temporary traffic priority. Relays will support a new circuit scheduler that will enable them to prioritize PriorityPass traffic over normal traffic. Once redeemed, PriorityPasses become useless.

As in Bitcoin, TEARS users may create accounts simply by generating public/private keypairs. Shallots may be transferred between these accounts, enabling third party markets to form; relays provide the initial *supply* of Shallots, and can trade them with regular users who *demand* prioritized Tor service. Additional Shallots may also be minted and distributed to users for free, according to some prescribed (and publicly known) policy. (See Section 4.2 for more information about the monetary policy and how to “bootstrap” the system.) Allowing regular users and not just relays to obtain Shallots (and therefore PriorityPasses) helps obfuscate the source of prioritized traffic, a key goal also motivating other recent incentive designs [15, 16]. Figure 1 illustrates the TEARS architecture. We next discuss the design of Shallots and PriorityPasses in more detail.

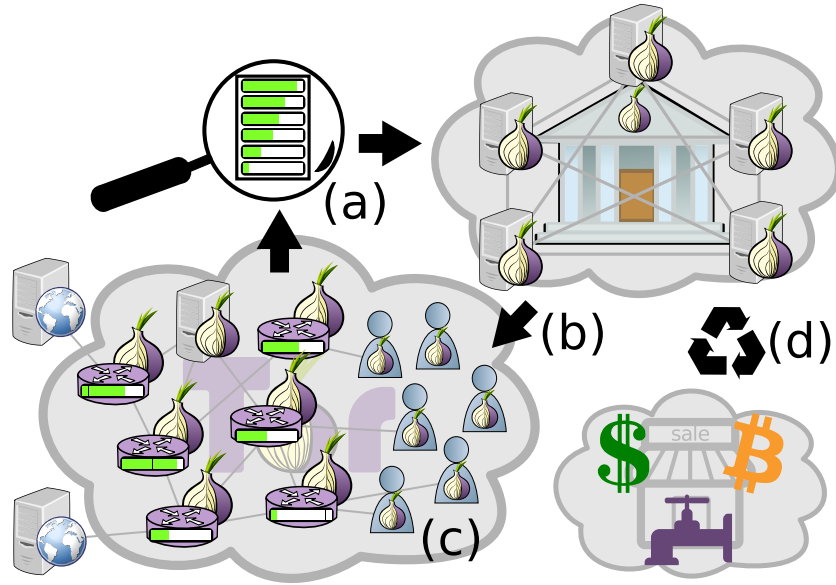


Fig. 1. An overview of our TEARS system. (a) A distributed process *audits* Tor relays’ bandwidth services and informs the bank, who then *mints* new Shallots for each relay and deposits it on their behalf. (b) Tor clients *redeem* Shallots for anonymous and unlinkable PriorityPasses at the bank through an anonymous Tor tunnel. (c) Tor clients *present* PriorityPasses for traffic priority at Tor relays, after which the PriorityPasses become useless. (d) Shallots can also be transferred from one user to another, as often as desired. Shallots may be given out for free via a “faucet”, and an external market may form around the exchange of Shallots for currency such as U.S. Dollars, Bitcoin, or other credentials deemed valuable by Tor users.

3.2 Accountable Banking with Shallots

Although the cash tokens in TEARS are initially minted and disbursed to users who contribute bandwidth, and are eventually redeemed as PriorityPasses in order to request enhanced service, the cash can be transferred among users any number of times in between. Users transfer tokens by communicating with a bank. We envision that the bank system will be operated by (a quorum of) semi-trusted servers, in particular the existing Tor directory servers. Although we refer to these servers as *semi-trusted*, we nonetheless design the system to minimize reliance on them. The key design choices are enumerated below:

Unconditional Privacy The tokens should be *e-cash*. In particular, even if all the bank servers are compromised, users’ transactions should still be unlinkable and private. This suggests that users should communicate with the bank system only through an anonymizing relay such as Tor itself.

Auditability The bank servers should be unable to violate the basic properties of a token currency—such as redeeming “counterfeit” tokens or transmitting

“defective” tokens to honest users—without *getting caught*. This means we must use a *publicly verifiable* (or *auditable*) e-cash protocol [22,29]. These schemes utilize non-interactive zero-knowledge proof systems. The bank maintains a public set of coins (public keys), and a user spends a coin (without revealing which) by proving possession of the private key for *some* coin not yet spent. In these schemes, it is safe to publish all messages between user and bank. In fact, any invalid messages from a bank or user are immediately detectable by anyone reading the transcript, without private information, even if users and banks collude.

We have mentioned that Shallots are initially minted in receipt of proofs of useful bandwidth contribution; other rules for minting Shallots may be supported as well, such as granting an additional fraction of each minted token to the administrators of the Tor Project. Such possibilities are discussed in more detail in Section 4. In all cases, however, an invariant to maintain is that the total quantity of currency in the system at any time is a matter of public record.

Availability Any majority of correctly-functioning servers should be sufficient for banking service to continue uninterrupted, despite Byzantine failures of the others; additionally, no user that is able to connect to a correct server should be denied service. Due to the use of an auditable e-cash protocol as described above, the role of the bank is essentially reduced to receiving transactions from users and assigning to them a sequential ordering. Any efficient variation of Byzantine Paxos [19] can be used so that the correctly-functioning servers arrive at an agreement about the next batch of valid transactions to include; after reaching agreement, the majority of servers then produce a threshold signature [5,27] over these transactions and a sequence number.

Transparency via Gossip or Broadcast As described above, an auditable e-cash protocol is chosen so that every transaction can safely be *published* and any misbehavior detected. However, this is only valuable if members of the public are indeed *watching* and dutifully inspecting the published data. Users and other interested parties should *gossip* about any messages received from bank servers, similar to *Certificate Transparency* [20]. In order to detect a double-spend or otherwise invalid transactions, some member of the gossip network must maintain a copy of the set of outstanding (i.e., minted but not-yet-redeemed) Shallots.

Alternatively, if a public broadcast channel is available, then interactions with the bank may simply be conducted over this broadcast channel. The Bitcoin blockchain is arguably a candidate for this, since arbitrary data may be included in Bitcoin transactions. This provides the further advantage that even a majority of the bank servers would be unable to selectively refuse service to clients; on the other hand, Bitcoin transactions generally require fees.

Recovery from Catastrophe We have described above a publicly-auditable e-cash system that resists any misbehavior by a minority of the bank servers, and furthermore guarantees that misbehavior by even a majority of the servers is at least detectable. Suppose such misbehavior occurs—for example, that all of the bank servers collude to print their own counterfeit Shallots and then to sell them in external markets—and that the misbehavior is detected as intended.

What should become of the system? Even in this catastrophic scenario, the system (including most Shallots) could be largely salvaged, as long as new (and, hopefully, more trustworthy) bank servers can be appointed. Since the history of Shallot transactions is public knowledge, the new bank can simply resume processing from a “fork” before the counterfeit event occurred.

Redeeming Shallots for PriorityPasses In our e-cash description above, we said that users “spend” Shallots by proving they possess a private key and that it has not yet been spent. The transaction published by the user must also indicate *how* she wishes the Shallots to be spent; two options are available. First, if a user wishes to transfer the Shallots to another user, then the transaction should contain a commitment to the recipient’s public key. The second option is to redeem Shallots for relay-specific PriorityPasses. This similarly involves providing a commitment to the relay’s public key. However, the two transaction types are distinguishable, and, unlike Shallots, PriorityPasses cannot be transferred again. Thus the total quantity of Shallots in the system consists of the total quantity minted minus the total quantity converted to PriorityPasses.

3.3 Using PriorityPasses

Once a user has obtained a PriorityPass, she can present it to the intended relay by opening the commitment in a single private message (which, as in BRAIDS [15], may be contained within an onion-wrapped message). Thus, without needing to interact with the bank, the relay can confirm the PriorityPass is dedicated to itself and has not previously been presented.

To be suitable for use in Tor, PriorityPasses must be **unlinkable** and **private** so that clients’ payments to relays do not deanonymize their traffic. A major problem with most electronic cash in the context of Tor is that it must be deposited after being spent in order to realize payment and detect double spending: the timing of the withdraw and deposit operations *leaks information* about the source of Tor traffic, even if the cash itself is unlinkable (see Appendix B for more details). Therefore, *no bank communication* should occur after relays receive PriorityPasses. To achieve this, PriorityPasses are **locally verifiable**, have **non-transferable value** (are useless after they are spent), and are **bound to a relay identity**: because PriorityPasses are relay-specific, each relay can immediately and locally *prevent double spending*. Finally, PriorityPasses are **non-transferable**: the relay identity bound to the PriorityPass may not be altered and PriorityPasses should not be transferred among users.

While the cryptographic construction of PriorityPasses is out of scope for this paper, they may be produced using blind signatures so that the bank will not be able to determine the relay identity bound to each PriorityPass.

3.4 Prioritizing Traffic and Auditing Bandwidth

To provide service enhancements after receiving PriorityPasses, relays will support a new circuit scheduler based on Proportionally Differentiated Services [9, 10] that is capable of proportionally prioritizing traffic belonging to different payment classes. This approach was also used in both BRAIDS [15] and LIRA [16],

and details about how this type of scheduler would replace Tor’s existing circuit scheduler was outlined by Jansen [13]. We provide a summary in Appendix C.

TEARS also relies on a secure method for auditing relays’ bandwidth contributions. While some work has been done in this area [11, 30, 31], we do not believe any existing design to be suitable for TEARS. While designing a bandwidth auditing system is out of scope for this paper, methods for accounting for bandwidth in general, as well as the limitations of existing measurement techniques, are discussed in Appendix D.

4 Discussion and Open Problems

Section 2 noted some of the technical requirements for a viable Tor incentive scheme, but there are deployment and social challenges as well. We outline these here, and defer a more extended discussion to Appendix E.

4.1 Incentives to Participate

Why should relays follow the protocol and actually give clients priority in exchange for PriorityPasses?

We’ve mentioned earlier that once PriorityPasses are spent, they are thereafter nontransferable and out of circulation, and therefore useless. It is, however, possible for a relay to prove to the world (i.e., in a publicly verifiable way) how many unique PriorityPasses it has received and processed. Relays may be interested in publishing such information for the sake of informal bragging rights, even absent any formal recognition or reward; alternatively, relays that demonstrate receipt of PriorityPasses might be entitled to additional minted Shallots. In either case, it’s important that relays do not *immediately* publish PriorityPasses they receive, since this would enable timing analysis and threaten the anonymity of users. At most, relays should publish the PriorityPasses they receive periodically (for example, once a month). If receipt of PriorityPasses entitle the relay to additional Shallot rewards, this process should also be scheduled periodically to avoid creating any perverse incentive for relays to compromise user privacy.

4.2 Market Economics

TEARS allows quantities of the currency to be transferred from one user to another; however, no explicit market-making mechanism is provided. We believe that as long as the Shallots are transferable, then third-party services will arise to satisfy any demand for trading Shallots for other currencies like U.S. dollars and Bitcoin. This has been the case with Bitcoin, for example, although the ecosystem of currency exchanges serving the Bitcoin community has involved a high degree of failure and fraud [24]. We leave the issue of establishing more trustworthy digital currency exchanges as an open problem.

As previously mentioned, Shallots are initially minted and disbursed to relays as a reward for contributing bandwidth. While this is one way for new Shallots to enter the system, other supplementary minting policies are possible as well. For example, the administrators of the Tor Project may be designated to receive a fraction of each minted Shallot (i.e., as a form of tax), receive a small fraction of a Shallot whenever transferring Shallots among users (i.e., as a transaction fee),

or even to print new currency at their discretion. An expiry might be assigned to each minted Shallot, such that it must be spent within some time interval or else is removed. These decisions constitute the *monetary policy* of TEARS, and we leave choosing appropriate parameters as part of deployment strategy.⁴

Allowing the administrators of the non-profit steering organization to receive additional currency has several benefits; in particular, the Tor project might operate a “faucet”⁵ that gives small amounts of Shallots to individual users who might not have money or access to currency exchanges. Due to the bank system’s public auditability, even if administrators can mint new currency at their discretion, the total amount minted would be publicly visible.

The ability for users (other than relays) to obtain Shallots via the faucet mechanism as well as third-party markets is important as it provides a non-trivial amount of added security: without these mechanisms, only service providers would be able to accumulate Shallots and would therefore be more easily identified as a potential source of traffic on circuits where PriorityPasses are presented for priority. However, we note that we want the system to be accessible to as many users as possible, and so Shallots and PriorityPasses should always be an optional part of the system and should never be required to receive basic service.

4.3 Community Effects

Adding explicit incentive mechanisms to a system like Tor that relies on voluntary behavior can have perverse effects. This has been widely observed, and quoted here from Benkler [2]:

Compare the level of use and success of pay-per-cycle distributed computing sites like Gomez Performance Networks or Capacity Calibration Networks, as compared to the socially engaged platforms like SETI@Home or Folding@Home. It is just too simplistic to think that if you add money, the really good participants will come and do the work as well as, or better than, the parallel social processes. . . . Adding money alters the overall relationship. It makes some people ‘professionals,’ and renders other participants, ‘suckers.’ It is not impossible to mix paid and unpaid participants, as we see in free and open source software and even to a very limited extent in Wikipedia. It is just hard, and requires a cultural form that is definitely not ‘now at long last we can tell who’s worth something and pay them, while everyone else is just worthless.’

If a new incentive scheme is incorporated into Tor, there is a danger of such a crowding out of intrinsically motivated volunteers, thus reducing rather than

⁴ Although Bitcoin, the first and currently most popular cryptocurrency, mints new coins only as a reward for network participation (and at a fixed rate, which is scheduled to gradually diminish over hundreds of years), other competing “altcoins” such as Freicoins (<http://freico.in/>) implement alternate monetary policies such as *demurrage* and allocating a portion of newly minted coins to a designated administrator.

⁵ See https://en.bitcoin.it/wiki/Bitcoin_faucet for a list of faucets that offer small quantities of Bitcoin for free to users.

increasing network size, diversity, and capacity. When faced with challenges to users, the network, or their relays, intrinsically motivated operators are more likely to resist on principle than someone for whom this is a mere economically rational question. However, we have not proposed straightforwardly liquid compensation: contribution to Tor simply allows an operator to have prioritized traffic in relay queues, an enhanced version of the basic service you would already get for free (much like the members entrance at a zoo or museum).

4.4 Recommended Deployment

A deployment strategy for radical Tor design changes is challenging, especially for the TEARS system because of the many problems and risks outline above. Both for our own research questions and for other potential innovations, we believe it would be useful for Tor to offer *experimental releases*. Appendix E.2 presents our rationale for this recommendation.

Experimental releases would typically introduce features that already have been vetted through academic publication, have done the necessary network experimentation to satisfy the community, and might typically have associated Tor proposals. They should also have been through some degree of scrutiny from the Tor developers. Nonetheless, these should clearly be labeled as experimental, similar to the labeling of alpha releases. Participating relays should be capable of handling both experimental traffic and production Tor network traffic, and we assume that Tor should always support both clients and relays that choose to ignore any feature in any experimental release.

To deploy TEARS using an experimental release and network, we expect that alternative consensus files would be created while new infrastructure may be required to provide an operational incentive network (e.g., new directory servers that also run the TEARS bank). Users that would like to try out TEARS features could use the new experimental consensus to choose relays supporting TEARS as part of the experimental release; other users would ignore the experimental release and continue using Tor as usual. We believe this deployment path will help reduce risk to the existing production network while still providing useful information about the potential deployment of new designs. We can also deploy all the experimental mechanisms and software for purposes of testing and bug fixing but initially turn the faucet mentioned in Section 4.2 on high. The inflationary effect will be to eliminate any value from prioritization since so much traffic will be eligible for PriorityPasses that prioritization will not provide significant performance benefit. Once TEARS has been deployed and tested it will then be possible to turn down the faucet at whatever rate and to whatever level is desired to evaluate the effect on prioritization and prioritization-motivated contribution to the network. In this way we can separate the deployment and testing of the experimental software from the experimentation with incentives.

5 Related Work

Tor-specific incentives systems have been studied many times before, but none of them discuss how to use a decentralized transaction processing service to handle the management of accounts.

Ngan *et al.* propose that the fastest 7/8 of relays get flagged as **gold star** relays in the consensus, and that **gold star** relays prioritize the traffic that is initiated at other **gold star** relays [26]. Similarly, Tortoise applies a universal rate limit to all clients, but exempts those that run relays marked with the **fast** and **stable** consensus flags [25]. Both of these approaches severely harm the anonymity of relays using Tor’s client functionality: the set of potential initiators for traffic receiving priority (or not being throttled) can be narrowed from the full set of clients to the set of relays with the correct flags.

Monetary payments are offered in exchange for routing services in PAR [1] and XPAY [7], where digital coins are inserted into the routing protocol. Clients obtain coins from a centralized bank, and relays must send coins back to the bank soon after receiving them in order to detect double spending. The timing of the withdrawal by the client and the deposit by the relay may leak information about the client’s traffic. Coins may be held by the relay longer to disrupt this timing information, but it also reduces the speed at which double spending can be detected.

The double spending problem was eliminated in BRAIDS [15] with the introduction of relay-specific tickets: relays themselves are responsible for verifying that their tickets have not been previously spent without contacting the centralized bank. BRAIDS does not require clients to provide a ticket in order to use the system, and freely distributes a small amount of tickets to all users. Although this increases the difficulty in distinguishing clients from relays based on their payments, relays also accumulate tickets for providing service and will therefore still be able to receive traffic priority more often than clients. Finally, BRAIDS is somewhat inefficient in that it depends on a central bank that must continuously exchange tickets for all users.

The efficiency of a central bank was improved in LIRA [16], where a cryptographic lottery was used to decide when clients could receive priority. In LIRA, clients simply guess a random number and receive priority with tunable probability. The central bank only has to manage accounts for relays, who receive the ability to obtain guaranteed winners to the lottery in exchange for providing service. Unfortunately, probabilistic guessing reduces flexibility for clients wanting to receive continuous priority over time, and creates a potential for cheating the system because it enables clients to continuously create circuits until a correct guess is found. Finally, LIRA relies on both a central bank and a secure bandwidth measurement system to inform the bank about the number of guaranteed winners each relay is allowed to obtain; as we will discuss, measuring relay bandwidth securely is an open research problem.

Johnson *et al.* argue that Tor should allow its services to be purchased [17] to take advantage of a large pool of new users that already pay for tools that circumvent censorship without the strong privacy guarantees that Tor offers. This new pool of clients could help fund new service providers and could increase Tor’s anonymity by diversifying its traffic sources.

6 Conclusions and Future Work

In this paper, we presented the design of the first Tor incentive system that is based on a distributed bank with fully transparent and auditable protocols for accountability. We believe our system minimizes the trust placed in the banking authorities while reducing or eliminating problems related to the double spending of tokens, the information leaked by transactions, and the processing limitations of a central bank.

Future Work There are several open problems that future work should consider. TEARS relies on a bandwidth auditing service to securely measure relays' contributions. We believe that more work is needed to determine the extent to which previous bandwidth measurement proposals [30,31] may help fulfill our requirement. Another challenging problem is how to prove that a relay was honest in giving priority when presented with a PriorityPass, and how to prove the relay did not change the scheduling parameters set by the network.

More work is needed before TEARS can be deployed. A system analysis should be done to ensure TEARS provides the properties we claimed it does, and an efficiency analysis should be done to determine the scalability as both the number of relays and the number of banking authorities increases. In addition, experiments with TEARS should be done using tools like Shadow [14] to determine how performance changes as the number of PriorityPasses that are available to users varies. Finally, usability and behavioral analyses would be extremely useful in helping to understand how an incentive system would affect the altruism of the existing community of volunteers.

Acknowledgments We thank Roger Dingledine, Mathew Green, Virgil Griffith, and Aaron Johnson for useful discussions regarding this work. We thank the TorProject for providing the Tor logo graphics used in Figure 1 under a Creative Commons attribution license,⁶ and acknowledge that they in no way endorse Figure 1 specifically or the contents of this paper generally. This work was supported by the Office of Naval Research. This material is based upon work supported by the Defense Advanced Research Agency (DARPA) and SPAWAR Systems Center Pacific, Contract No. N66001-11-C-4018.

References

1. Androulaki, E., Raykova, M., Srivatsan, S., Stavrou, A., Bellovin, S.M.: PAR: Payment for anonymous routing. In: Privacy Enhancing Technologies Symposium (PETS) (2008)
2. Benkler, Y.: Benkler on Calacanis's wallet. <http://www.rough.type.com/?p=479> (July 28 2006)
3. Biryukov, A., Pustogarov, I., Weinmann, R.: Trawling for Tor hidden services: Detection, measurement, deanonymization. In: Symposium on Security and Privacy (SP) (2013)
4. Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W.: An Architecture for Differentiated Services. RFC 2475 (1998)

⁶ <http://creativecommons.org/licenses/by/3.0/us/>

5. Boldyreva, A.: Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In: Public key cryptography (PKC) (2002)
6. Chaabane, A., Manils, P., Kaafar, M.: Digging into anonymous traffic: A deep analysis of the tor anonymizing network. In: Conference on Network and System Security (NSS) (2010)
7. Chen, Y., Sion, R., Carbunar, B.: XPay: Practical anonymous payments for Tor routing and other networked services. In: Workshop on Privacy in the Electronic Society (WPES) (2009)
8. Danezis, G., Syverson, P.: Bridging and fingerprinting: Epistemic attacks on route selection. In: Privacy Enhancing Technologies Symposium (PETS) (2008)
9. Dovrolis, C., Ramanathan, P.: A case for relative differentiated services and the proportional differentiation model. *Network* 13(5), 26–34 (1999)
10. Dovrolis, C., Stiliadis, D., Ramanathan, P.: Proportional differentiated services: Delay differentiation and packet scheduling. *Transactions on Networking* 10(1), 12–26 (2002)
11. Ghosh, M., Richardson, M., Ford, B., Jansen, R.: A TorPath to TorCoin: Proof-of-bandwidth altcoins for compensating relays. In: Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs) (2014)
12. Jakobsson, M.: Ripping coins for a fair exchange. In: Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT) (1995)
13. Jansen, R.: Privacy Preserving Performance Enhancements for Anonymous Communication Systems. University of Minnesota PhD Thesis. (October 2012)
14. Jansen, R., Hopper, N.: Shadow: Running Tor in a box for accurate and efficient experimentation. In: Network and Distributed System Security Symposium (NDSS) (2012)
15. Jansen, R., Hopper, N., Kim, Y.: Recruiting new Tor relays with BRAIDS. In: Conference on Computer and Communications Security (CCS) (2010)
16. Jansen, R., Johnson, A., Syverson, P.: LIRA: Lightweight Incentivized Routing for Anonymity. In: Network and Distributed System Security Symposium (NDSS) (2013)
17. Johnson, A., Jansen, R., Syverson, P.: Onions for sale: Putting privacy on the market. In: Financial Cryptography and Data Security Conference (FC) (2013)
18. Johnson, A., Wacek, C., Jansen, R., Sherr, M., Syverson, P.: Users get routed: Traffic correlation on tor by realistic adversaries. In: Conference on Computer and Communications Security (CCS) (2013)
19. Lamport, L.: Fast Paxos. *Distributed Computing* 19(2), 79–103 (2006)
20. Laurie, B., Langley, A., Kasper, E.: Certificate transparency. Network Working Group Internet-Draft, v12, work in progress. <http://tools.ietf.org/html/draft-laurie-pki-sunlight-12> (2013)
21. McCoy, D., Bauer, K., Grunwald, D., Kohno, T., Sicker, D.: Shining light in dark places: Understanding the Tor network. In: Privacy Enhancing Technologies Symposium (PETS) (2008)
22. Miers, I., Garman, C., Green, M., Rubin, A.D.: Zerocoin: Anonymous distributed e-cash from bitcoin. In: Symposium on Security and Privacy (SP) (2013)
23. Miller, A., Juels, A., Shi, E., Parno, B., Katz, J.: Permacoin: Repurposing bitcoin work for long-term data preservation. In: Symposium on Security and Privacy (SP) (2014)
24. Moore, T., Christin, N.: Beware the middleman: Empirical analysis of bitcoin-exchange risk. In: Financial Cryptography and Data Security Conference (FC) (2013)

25. Moore, W.B., Wacek, C., Sherr, M.: Exploring the potential benefits of expanded rate limiting in tor: Slow and steady wins the race with tortoise. In: Annual Computer Security Applications Conference (ACSAC) (2011)
26. Ngan, T.W.J., Dingleline, R., Wallach, D.S.: Building incentives into Tor. In: Financial Cryptography and Data Security Conference (FC) (2010)
27. Pedersen, T.P.: A threshold cryptosystem without a trusted party. In: Advances in Cryptology (EUROCRYPT) (1991)
28. Reiter, M., Wang, X., Wright, M.: Building reliable mix networks with fair exchange. In: Conference on Applied Cryptography and Network Security (ACNS) (2005)
29. Sander, T., Ta-Shma, A.: Auditable, anonymous electronic cash. In: Advances in Cryptology (CRYPTO) (1999)
30. Snader, R., Borisov, N.: Eigenspeed: secure peer-to-peer bandwidth evaluation. In: International Workshop on Peer-To-Peer Systems (IPTPS) (2009)
31. Snader, R., Borisov, N.: Improving security and performance in the tor network through tunable path selection. Transactions on Dependable and Secure Computing 8(5), 728–741 (2011)
32. Tor Network Metrics. <https://metrics.torproject.org/network.html>
33. Tor Proposals. <https://gitweb.torproject.org/torspec.git/tree/HEAD:/proposals>

Appendices

A Useful Service in Tor

The main goal of a Tor incentive system is to convince users that they should contribute useful services to the network by providing rewards. A Tor service may include bridge, guard, middle, and exit relays, directory authorities and mirrors, and hidden service directories. While simply running these services may provide some benefit to the network, composing any of the following properties will generally increase utility:

- *High bandwidth*: more bandwidth to better support Tor’s load.
- *Location diversity*: expanding Tor’s reach to geographic regions containing few or no relays will help users route around untrusted parts of the Internet.
- *Stable*: longer uptimes improves network stability.
- *Exit policies*: liberal exit policies will increase relay utility.
- *Metrics*: share relay or other statistics about network operation.
- Newer Tor *versions* or support for experimental features.

There is no direct existing means to encourage useful service such as location-diverse or high bandwidth relays. As a result, existing Tor relays naturally gravitate towards familiar, cheap, and Tor friendly ISPs, and many of them provide such a low level of bandwidth that their relays are rarely chosen for client circuits. An incentive design could differentially reward contributions based on their usefulness to the network. While this may add complexity to the process for verifying contributions, it would allow Tor to get the most out of its contributors.

An incentive design could incorporate a diversity weight to provide rewards for relays based on location thus improving general security as well as performance of the Tor network. We could also create a super-linear reward scale to

encourage higher bandwidth capacity relays, or a sub-linear scale to encourage greater uniformity of relay bandwidth. Although it is clear that useful service is certainly not limited to relay bandwidth, we focus on it in this paper as it is currently the most demanded Tor resource.

B Payments Leak Information

To request a service or service enhancement, clients in a Tor incentive system generally need to provide some kind of bank-authorized payment to the relays. These payments sometimes use electronic cash schemes, but are usually constructed with blind signatures and zero-knowledge protocols. For the purposes of this section, we will refer to these digital certificates more generally as *certs*.

Certs provide unlinkability between when they are withdrawn from the bank and spent at a merchant, making their use in a system of accounting an attractive option for an anonymity network like Tor. Here we discuss how using certs to request enhanced service may leak information in Tor. Further, if relays' contributions are accounted for via an independent process (as discussed in Appendix D), then the question of "*Should the certs themselves have value to the relay?*" can be thought about by considering the roughly analogous notion of receiving "tips" in addition to "salary". We will also briefly speculate on how this may affect relay behavior.

B.1 Certs that Retain Value After Transfer

With generic certs that have transferable value, the merchant must generally deposit each cert immediately in order to detect double spending. Because the timing of the withdraw and associated deposit can leak information about a client's usage activity to the bank, the client should add some "mixing time" before spending it to improve unlinkability at the cost of flexibility.

Relay-specific certs help mitigate this problem to some extent, because the relays are able to prevent double spending locally. However, if these certs retain value that can be transferred to another relay, then the receiving relay will still need to deposit certs it receives via transfers. The receiving relay should add some amount of "mixing time" before depositing transferred certs, to obfuscate from the bank the withdraw operation associated with each deposited cert. Some flexibility in usage is lost because the certs are good only at a single relay unless they are transferred, and they must be mixed again before each transfer.

To summarize, with general certs:

- clients may use them at any relay;
- relays should deposit them immediately to detect double spending;
- clients should delay usage after withdrawal to reduce linkability;

With relay-specific certs:

- clients withdraw them from the bank for specific relays (the chosen relay is blinded to the bank);
- clients may spend them immediately;
- relays may prevent double spending locally;
- relays should delay deposit after receipt to reduce linkability;

- long delays are imposed if clients exchange unspent certs (e.g., because the relay bound to the original cert went offline).

The “mixing” happens on the client end with generic certs, and the relay end with relay-specific certs. These approaches may be complimentary, such that combining them would improve flexibility: clients keep a stash of general certs usable anywhere in the long term, and supplement that with relay-specific certs usable at a single relay in the short term.

If we allow clients to “tip” relays – by directly transferring value to relays above and beyond the system’s normal reward mechanism based on measured bandwidth – then we create an incentive for relays to provide service differentiation (see Appendix C). If this incentive exists, a selfish relay will always prefer to forward “paid” traffic over “unpaid” traffic. Relays will try to make priority traffic as attractive to clients as possible, and thus prefer scheduling parameters that result in bigger performance boosts. With enough “tipping” clients, the “non-tipping” clients would eventually stop using relays that provide horrible service to them. Of course this will always be true to some extent due to the functionality of the scheduler, but relays would have a reason to adjust their scheduler in a way that forces the situation earlier.

B.2 Certs that are Valueless After Transfer

With relay-specific certs that have no transferable value, the relay can simply detect double spending locally without ever depositing anything back to the bank, thereby addressing the linkability problem. However, some flexibility is lost in this design because the certs are bound to specific relays and cannot be transferred to other relays. Further, relays would need to earn rewards via a separate channel (such as a collective bandwidth measurement mechanism as proposed in this paper).

If we do not allow “tips” or re-use of certs by the relays, then we expect the relays to remain agnostic about the operation of the scheduler. They will be happy as long as they are able to provide bandwidth, because that is how they receive rewards. This will allow Tor to better control the scheduling parameters in a way that makes service for “non-tipping” clients bearable. The relays have an incentive to follow the protocol, because otherwise clients may be able to detect malicious behavior, and might even form a reputation system to gossip opinions about relays exhibiting bad or suspiciously deviant behavior. While *proving* that relays didn’t change their scheduling parameters would be challenging, at least they don’t have a direct incentive to be dishonest – provided the collective bandwidth measurement mechanism is secure and relays cannot game it to collect rewards disproportionate to the utility they offer.

Given the discussion in this section, we believe the benefits of relay-specific certs without transferable value outweigh the loss in flexibility, and therefore designed PriorityPasses accordingly.

C Rewards for Providing Service

Relays will receive rewards for their useful contribution to the network. We believe a desirable reward for providing service is preferential treatment of traffic

by bandwidth providers. We intend that rewards be used to enhance performance by differentiating the traffic priority of rewarded flows from unrewarded flows. Note, however, that this approach does not *guarantee* service differentiation, because the ability to provide performance enhancements depends on network dynamics such as relay capacities and client load distribution.

C.1 Differentiated Services

The differentiated services architecture [4] can be used to improve the control over traffic priority in Tor, as previously outlined by Jansen [13]. More specifically, the proportional differentiation model [9] allows for predictable (i.e., consistent as load increases) and controllable (i.e., adjustable differentiation) performance between N traffic classes. The model allows for the configuration of a differentiation parameter p_i for each class i , and enforces the proportional priority of a traffic quality metric q between all pairs of classes i and j for measurement timescale σ as:

$$\forall i \in [N], \forall j \in [N] : \frac{q_i(t, t + \sigma)}{q_j(t, t + \sigma)} = \frac{p_i}{p_j} \quad (1)$$

where $p_1 < \dots < p_N$ and p_i/p_j defines the quality proportion between classes i and j . The model is well defined when there is enough traffic in each class to allow a work-conserving scheduler to meet the desired proportions.

Dovrolis *et al.* design a scheduler under the proportional differentiation model using a queuing delay metric [10], which in our case corresponds to Tor cell waiting times. For class i , the quality metric q_i combines the queuing delay $D_i(t)$ of the longest waiting cell with the long-term average delay $\delta_i(t)$ of all previously scheduled cells at time t :

$$q_i(t) = D_i(t) \cdot f + \delta_i(t) \cdot (1 - f) \quad (2)$$

where f is an adjustable fraction that tunes the scheduler’s ability to react to short term spikes in delay. When a scheduling decision is to be made at time t , the longest waiting cell from the class with the maximum priority $P(t) = q(t)/p(t)$ is chosen and scheduled.

C.2 DiffServ Parameters

Any number of classes can be configured as well as the prices and proportions between them. We suggest three service classes: basic (free), standard, and premium. Clients can select their priority class by supplying the correct number of PriorityPasses independently to each relay in their circuit, and each PriorityPass can provide the selected priority for a configurable amount of data. Although heuristic measurements could suffice, how to *prove* that a presented PriorityPass actually resulted in a service enhancement reward is an open research problem.

C.3 Distinguishing Traffic by Priority

We acknowledge that traffic priority will fundamentally allow users that are categorized into different service classes to be somewhat distinguishable from one another, which may result in a partitioning of Tor’s anonymity set. However, as Johnson *et al.* argue [17], users with lower security requirements may be willing

to trade off reduced security for improved flexibility and performance: users already use Tor for downloading large files even though their activities look very different than most [6,21]. Further, a faster network that is more flexible (clients may specify their desired performance level) may attract a significant number of new users and result in a net increase in anonymity. We note that it is important to make clear to users the risks and trade-offs they assume by participating in such a network.

D Bandwidth Accounting

A Tor incentive system should provide a means to prove that the service was actually rendered, i.e., to account for relay bandwidth contributions in a way that is *accurate* for all relay positions and in both upstream and downstream directions, is *secure against cheating*, and *cannot link clients* to the relays they choose, the circuits they use, or the traffic they send.

In this section, we will use the generic term *bwtoken* to represent bandwidth contribution levels: obtaining bwtokens should be roughly correlated with contributing bandwidth. There are at least two methods for using bwtokens to account for bandwidth:

1. relays receive bwtokens directly from clients in exchange for utilizing their bandwidth resources to forward traffic; and
2. relays receive bwtokens indirectly through a distributed process that decides the number of bwtokens each relay should receive, e.g., based on bandwidth measurement audits.

D.1 Direct Accounting

In order to use the first method above to accurately account for all contributed bandwidth, clients must be able to continuously send bwtokens in exchange for service in a way that does not break client anonymity and cannot be exploited by service providers to earn more bwtokens than they deserve. Also, relays should be able to use bwtokens to prove their contributions to other system members.

The BRAIDS incentive scheme [15] shows how to use coin ripping [12] and fair exchange [28] to ensure that relays don't receive "tickets" unless they actually perform the forwarding service. However, BRAIDS suffers from somewhat high performance costs because it distributes tickets to every client in the system (though tickets are not required to receive service). If all clients were required to use tickets so that we could accurately account for relay bandwidth, the overhead would also increase.

In the TorPath scheme [11], participating clients work with the three relays on a Tor circuit to mint coins in proportion to the useful end-to-end bandwidth they observe. To prevent clients from choosing colluding relays to mint coins dishonestly, TorPath uses verifiable shuffles to assign relays to clients using secret but publicly verifiable randomness. This approach measures bandwidth end-to-end and limits the system's vulnerability to unfair or malicious ticket distribution, at the cost of limiting clients' freedom in choosing relays.

We would like clients to receive enhanced service whether or not they are required to directly handle bwtokens for bandwidth accounting purposes. Note that the problem of handling enhanced service payments without leaking information is discussed in Appendix B and is independent of the problem of accurately accounting for bandwidth contributions: although many Tor incentive schemes use a single token to handle both problems at once, this need not be the case.

D.2 Indirect Accounting

Accounting for bandwidth indirectly requires a measurement infrastructure and service audit system that prevents providers from cheating to receive bwtokens disproportionate to their actual contributions. The LIRA incentive scheme [16] also relies on such a system. A bandwidth audit system measures the expected bandwidth at each relay and produces weights for the Tor consensus that are used during path selection. If accurate and secure, these weights could also dictate how to distribute bwtokens to relays in proportion to their contributions. Tor currently runs a bandwidth audit service, but it has been shown to be easily manipulable [3].

The most recent alternative Tor bandwidth measurement system is called EigenSpeed [30, 31]. In EigenSpeed, relays opportunistically measure their interactions with other relays and send the observation vectors to the authorities. The authorities combine the measurements from all relays using principal component analysis (PCA), and produce a set of authoritative weights that can be distributed via the Tor consensus file. Unfortunately, EigenSpeed has the following problems:

- It does not account for asymmetric bandwidth.
- It is unclear how to measure relays acting in entry and exit positions. Since EigenSpeed measures only the peer-to-peer bandwidth relays offer *each other*, a relay might obtain a high EigenSpeed rating without ever accepting connections from non-relay clients or forwarding traffic to external destinations.
- Because measurements are opportunistic, EigenSpeed tends to under-estimate the contributions of the fastest relays.
- It is unclear how to handle sybil attacks by malicious collectives.

Correctly attributing bandwidth contributions to providers via either of the two methods outlined in this section remains an open research problem, and one that future work should consider.

E Extended Discussion

E.1 Direct Payment for Anonymity

There have been previous approaches deployed involving some form of direct payment for Tor. It is instructive to consider what these do and do not provide. Torservers.net is an international federated consortium of nonprofit organizations that uses donated funds to purchase service especially for large capacity exit relays in various locations (so far in the U.S. and Europe). This system offers direct financial payment for provided service rather than providing either

basic service or service improvement as compensation. It also does not aspire to the market driven exchangeability of incentives that we envision. It is based more on the cost of running a given capacity server in a given market than on direct measurement of amount of service provided.

A commercial pre-Tor onion routing network, the Freedom Network from Zero Knowledge Systems, was operated in the early 2000s. This network allowed clients to purchase a number of pseudonyms that were granted access to the Freedom Network, and like `torservers.net`, network relay operators were paid monetarily for bandwidth provided. This, however, was designed so that all network providers were expected to operate for compensation and essentially all clients were expected to pay proportionally to service obtained or contracted.

E.2 Path to Deployment

It may be useful for cleaner experimental results to deploy a separate pay network along the lines of Freedom, for example, to see the relation of differentiated service parameters to incentives for providing various amounts of service. Our long-term goal, however, is to grow the Tor network for all users regardless of ability to pay. Such experimental results are thus unlikely to be at all representative of intended usage patterns and may thus be of limited usefulness—without even considering the impact of such partitioning on anonymity.

This necessitates a plan for deploying an incentive scheme that dovetails with existing and otherwise planned Tor protocols and Tor infrastructure, which will thus require the explicit approval of principal Tor developers and the tacit approval of a statistically large fraction of Tor relay operators, directory operators, etc. To gain the approval of developers, it is not enough to have vetted research results. Tor proposals [33] will need to be created and accepted.

We will further explore the need for tacit relay operator approval presently. We note here simply that anything deployed should minimize major unpleasant surprises for relay operators. A measure of this might be resistance to adopting new Tor releases that incorporate new incentive mechanisms. Rollout must also be gradual in multiple senses: the more relays running versions of Tor incorporating features of any incentive scheme are able to function smoothly with relays running versions of Tor without those features, the easier it will be to introduce those features into the network. At the same time, the more multiple incentive-scheme feature changes are simultaneously deployed rather than one at a time, the harder it will be to understand, isolate, and debug problems that inevitably arise when features are deployed in the wild. We can also separate rollout of a bank and network information directory system and associated experiments of their interaction with participating relays and simulated clients from subsequent rollout of experimental clients.

Tor follows common practice of making alpha releases available to those who are willing to take some risks to help understand and debug new features before they are incorporated into stable releases for the general user and operator population, for example, when introducing a new handshake or directory information protocol. This amounts to a live experimental network interoperating with the stable network. Alpha releases are, however, intended to incorporate

features and changes that may need some further adjusting and debugging but are already intended for deployment rather than for experimental evaluation.

Approaches to experimental research include simulation and live or testbed network experimentation. Shadow [14] permits full scale simulation of the entire current Tor network on a single server. But this cannot address a primary question of how incentives will affect behavior of client and relay operating users. PlanetLab has been used both for entirely experimental Tor networks and to run experimental relays that participate in the public Tor network. As a general shared network experimentation resource PlanetLab may not be suitable for sustained large scale experimentation involving differentiated service. Also, it is important to manage the sorts of experiments running on the live Tor network lest they significantly harm security or performance for the users or network elements. We must also consider how to make client software adequately available.

As discussed in Section 4.4, we advocate for a new experimental Tor release as a means for creating an experimental network. Research experiments may require relays that are only capable of handling experimental traffic but that can pass this transparently to other parts of the Tor network. As with network experimentation, some of the infrastructure may run on PlanetLab as appropriate, but that need not be the case. After experimental release, we expect appropriate parts of new designs to follow Tor's traditional deployment process: first in alpha releases, then standard releases; older versions may eventually become unsupported.

E.3 Managing Network Information

Tor's current directory system is distributed but centralized. A relatively small set of directory authorities (currently roughly 2.5 orders of magnitude smaller than the network) vote on the status of relays in the network. The results are pushed out to all relays so that, except when unusual problems arise, Tor clients do not need to interact with this small set.

It is an open question how far Tor's current approach will scale. However, it is unclear how far it practically needs to do so given the limitations to diversity of locations for a large adversary to observe in the underlying network [18]. The anonymity benefit of attempting to maintain a single uniform network picture versus allowing partitioning may not be as great as often assumed [8]. Nonetheless, alternative approaches to distributing network information have been explored in the research literature, including P2P and DHT based schemes and PIR schemes. None have yet been vetted to the point of deployment or even to having a Tor proposal. It is an open research question if decentralized coin systems already in use, such as Bitcoin, could be used for storage of Tor network information [23].