Poster: Thinking Inside the BLAC Box

Smarter Protocols for Faster Anonymous Blacklisting

Ryan Henry
PhD Student

Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1

rhenry@cs.uwaterloo.ca

Ian Goldberg
Associate Professor

Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1

iang@cs.uwaterloo.ca

The Internet can be a dangerous place to visit. As the proportion of our daily activities that occur online continues to increase, so too does our exposure to online privacy risks imposed on us by fraudsters and identity thieves, by intrusive advertising companies, by oppressive governments, and by countless unknown others. Anonymous communications systems like Tor[1] mitigate some of these threats by helping users to access services over the public Internet while concealing their identities and usage patterns from prying eyes. A global user base leverages the anonymity afforded by Tor and other anonymous communication systems to circumvent online censorship, to research taboo and unpopular subjects, and to speak their minds without fear of retaliation. Not only is this a win for privacy and free speech on the Internet, but it is also a potential boon for many online communities that might benefit from added diversity in their respective user populations. Compelling examples of such online communities include collaborative encyclopedias like Wikipedia[2] and community-driven review sites like Yelp[3].

Yet reality is rarely so simple. The providers of such online services must ultimately weigh the expected benefits (both to themselves and to their user communities) of more inclusivity against the risks posed by abusive users, especially those who would hide behind the veil of anonymity to skirt accountability for their actions. A number of popular services—notably including Wikipedia, Yelp, Slashdot[4], Craigslist[5], and most major IRC networks [10]—presently block contributions from anonymous users, despite the implied loss of diversity and the broader implications for free speech and the open exchange of knowledge and ideas.

[1] https://www.torproject.org/
[2] https://en.wikipedia.org/
[3] https://www.yelp.com/
[4] https://slashdot.org/
[5] http://www.craigslist.org/

In response, the cryptographic and privacy research communities have proposed several *anonymous blacklisting* designs, which seek to provide mechanisms through which service providers (SPs) may hold anonymous users accountable for their individual actions *without threatening their anonymity*. SPs can thereby protect their user communities from abuse by the occasional "naughty" anonymous user without inflicting collateral damage on all the "nice" users. An early proposal called Nymble [7] solved the anonymous blacklisting problem elegantly and efficiently; however, Nymble and its progeny [6, 8, 9] rely on powerful trusted third parties (TTPs) that can deanonymize (or link) users' connections undetectably and at will. Subsequent designs [3, 4, 12] have introduced clever cryptography to replace the TTPs, thus solving the trust problem at a cost of much computation and communication overhead both for the users and for the SPs.

One such TTP-free anonymous blacklisting design is Tsang et al.'s *Blacklistable Anonymous Credentials* (BLAC) [11]. In BLAC, a semi-trusted *group manager* (GM) registers each new user into the system by issuing it an anonymous credential $C(x)$ that encodes as an attribute a secret key x unique to that user. (The GM is semi-trusted in the sense that, although the SPs must trust the GM to issue credentials judiciously, lest they succumb to *Sybil attacks*, the users need not trust the GM to maintain their anonymity.) The user holding $C(x)$ authenticates to an SP by producing a *ticket* $\Gamma = (g, g^x)$ together with zero-knowledge proofs that (i) the exponent x used to compute Γ is the same as the secret key x in $C(x)$, and (ii) no ticket on the SP's *blacklist* of tickets from past abusive sessions uses that same x. Both proofs are instantiable using standard techniques for proving statements about the equality and inequality of discrete logarithms. The SP grants the user access (and stores Γ for future

reference) if and only if it accepts both proofs. If it later deems the user's actions during the session to have been abusive, the SP can add Γ to its blacklist to curtail further abuse by that user. The notion of "abuse" in this model is entirely subjective: each SP must define, identify, and penalize abusive behaviour in a way that is appropriate within the context of its user community and the services it provides.

BLAC's all-or-nothing approach to revocation may be overly punitive in some settings. The anonymous blacklisting literature includes two variants of BLAC that seek to address this shortcoming. The first variant does so with a *d-strikes-out revocation policy* [12], wherein each user may authenticate anonymously until it has accumulated d or more tickets on the blacklist (after which, future authentications will fail). The second variant supports *reputation-based blacklisting* [2], wherein SPs can assign scores (both positive and negative) to the anonymous actions of users, and each user may subsequently authenticate only if the aggregate score associated with all of its scored tickets exceeds some minimum threshold value.

Judged solely on the basis of privacy and functionality, the BLAC approach to anonymous blacklisting is very attractive indeed; judged also on the basis of scalability, however, it becomes much less so. In all three BLAC variants, the bottleneck operation is the second zero-knowledge proof (in which the user demonstrates that its own tickets on an SP's blacklist do not meet that SP's revocation criteria). The 'size' of this proof scales as the total number of tickets on the blacklist, which can introduce considerable delays and consume considerable bandwidth and computation capacity for large SPs that cater to millions of users. Prior work [2, 11, 12] essentially regards the zero-knowledge proofs as "black boxes", to be instantiated using the standard techniques from the literature. Unfortunately, those standard techniques become prohibitively expensive even for moderate-sized blacklists (say, those containing a few hundred tickets). This fact has contributed to the common conception [1, 3, 5] that—despite being both novel and elegant—the BLAC approach to anonymous blacklisting is impractical for large SPs.

In this work, we improve on BLAC and its derivatives by peering *inside* their zero-knowledge black boxes and optimizing the underlying protocols; that is, we innovate by *"thinking inside the BLAC box"*. We find, in particular, that existing *batch proof and verification* techniques can reduce substantially the communication and computation overhead in vanilla BLAC's bottleneck zero-knowledge proof. We then extend our optimized protocol in a novel way to deal also with the bottleneck proofs in the two above-mentioned BLAC variants. At the heart of our new constructions is a novel system for batch zero-knowledge proofs of partial knowledge for discrete logarithms over *non-monotone* access structures. Our new protocols appear to be the first in the literature for batch zero-knowledge proofs over non-monotone access structures and we suspect that our techniques will find applications in speeding up other cryptographic protocols that require proofs of similar statements in other contexts.

We demonstrate both analytically and experimentally that anonymous blacklisting with our new and improved black boxes is practical even for today's largest SPs using only modest commodity hardware.

References

[1] M. H. Au and A. Kapadia, "PERM: Practical reputation-based blacklisting without TTPs," in *Proceedings of CCS 2012*, Raleigh, NC, USA, October 2012, pp. 929–940.

[2] M. H. Au, A. Kapadia, and W. Susilo, "BLACR: TTP-free blacklistable anonymous credentials with reputation," in *Proceedings of NDSS 2012*, San Diego, CA, USA, February 2012.

[3] M. H. Au, P. P. Tsang, and A. Kapadia, "PEREA: Practical TTP-free revocation of repeatedly misbehaving anonymous users," *ACM Transactions on Information and System Security*, vol. 14, no. 4, p. 29, December 2011.

[4] E. Brickell and J. Li, "Enhanced Privacy ID: A direct anonymous attestation scheme with enhanced revocation capabilities," *IEEE Transactions on Dependable and Secure Computing (TDSC)*, vol. 9, no. 3, pp. 345–360, May 2012.

[5] R. Henry and I. Goldberg, "Formalizing anonymous blacklisting systems," in *Proceedings of IEEE S&P 2011*, Berkeley, CA, USA, May 2011, pp. 81–95.

[6] R. Henry, K. Henry, and I. Goldberg, "Making a Nymbler Nymble using VERBS," in *Proceedings of PETS 2010*, ser. LNCS, vol. 6205, Berlin, Germany, July 2010, pp. 111–129.

[7] P. C. Johnson, A. Kapadia, P. P. Tsang, and S. W. Smith, "Nymble: Anonymous IP-address blocking," in *Proceedings of PETS 2007*, ser. LNCS, vol. 4776, Ottawa, ON, Canada, June 2007, pp. 113–133.

[8] Z. Lin and N. Hopper, "Jack: Scalable accumulator-based nymble system," in *Proceedings of WPES 2010*, Chicago, IL, USA, October 2010, pp. 53–62.

[9] P. Lofgren and N. Hopper, "BNymble: More anonymous blacklisting at almost no cost (a short paper)," in *Proceedings of FC 2011*, ser. LNCS, vol. 7035, Gros Islet, St. Lucia, February 2011, pp. 268–275.

[10] The Tor Project, "List of IRC/chat networks that block or support Tor," https://trac.torproject.org/projects/tor/wiki/doc/BlockingIrc (Retrieved: 2013-04-08).

[11] P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith, "Blacklistable Anonymous Credentials: Blocking misbehaving users without TTPs," in *Proceedings of CCS 2007*, Alexandria, VA, USA, October 2007, pp. 72–81.

[12] ——, "BLAC: Revoking repeatedly misbehaving anonymous users without relying on TTPs," *ACM Transactions on Information and System Security (TISSEC)*, vol. 13, no. 4, pp. 39:1–39:33, December 2010.