Brief Announcement: Byzantine Agreement with Homonyms

Carole Delporte-Gallet*
LIAFA, Université Paris Diderot
cd@liafa.jussieu.fr

Hugues Fauconnier
LIAFA, Université Paris Diderot
hf@liafa.jussieu.fr

Rachid Guerraoui
Ecole Polytechnique Fédérale de Lausanne
Rachid.Guerraoui@epfl.ch

Anne-Marie Kermarrec
INRIA Rennes
Anne-Marie.Kermarrec@inria.fr

ABSTRACT

In this work, we address Byzantine agreement in a message passing system with homonyms, *i.e.* a system with a number l of authenticated identities that is independent of the total number of processes n, in the presence of $t < n$ Byzantine processes.

We prove the following results: (i) agreement is possible if (and only if) $l > 3t$ in a synchronous model; (ii) agreement is impossible, independently of the number of failures, in an eventually synchronous model; (iii) eventual agreement is possible, if (and only if) $l > 3t$, in an asynchronous model.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*distributed networks*; C.2.4 [**Computer-Communication Networks**]: Distributed Systems; C.4 [**Performance of Systems**]: [Fault tolerance]

General Terms

Algorithms, Theory, Reliability

Keywords

Consensus, Message-Passing, Byzantine Agreement, Authentication.

1. INTRODUCTION

We study Byzantine agreement in a message passing system with a limited number of *authenticated identities*, i.e., a system with *homonyms*. Basically, there is a set of l different identities used to identify processes. This set can be much smaller than the set of processes and hence several processes may have the same identity. Such processes can be considered as homonyms. Identities are authenticated in the sense that if a process p receives a message from a process q with identity id, p knows that the message does not come from a process with name $id' \neq id$. The message could however come from a homonym process $q' \neq q$, also with identity id.

In this context, we ask whether Byzantine agreement can be solved if l is independent of the total number of processes n. Beyond intellectual curiosity, the motivation of addressing this question is twofold.

Firstly, the assumption, underlying many distributed systems, e.g. [12, 14], that all processes have unique (unforgeable) identities might, we believe, be considered too strong. More specifically, authenticated unique identities are typically achieved through collision-free hash functions. Yet, such functions are potentially breakable, e.g.[15], namely they may lead to the same output with different inputs. Even the ones that have not been broken yet, such as SHA-256, might be in a near future.

Secondly, in many cases, processes (users) do care about their privacy and would rather not reveal their identity. In a fully anonymous system where processes do not have identities [2], reaching agreement is simply impossible in the face of Byzantine players. With a limited number of identities, one can preserve some level of anonymity and hide, to some extent, the association between every user and every identity. A limited number of identities, may render some agreement possible, as we will show in the sequel.

2. RESULTS

Assuming a system of n processes, t potential Byzantine failures and l identities, we prove the three following main results:

1. In a synchronous model [9], Byzantine agreement is possible if (and only if) $l > 3t$. This result is to be contrasted with the classical impossibility of Byzantine agreement if $n \leq 3t$: in a sense, we show that what really matters is the ratio between the number of Byzantine processes and the number of authenticated identities rather than the one between the number of Byzantine processes and the total number of processes.

   We prove the existence of our algorithm by simulation: we show how to transform any synchronous Byzantine agreement protocol, where all processes have uniquely authenticated identities (i.e., $n = l$), into a Byzantine agreement protocol using only l ($n \geq l > 3t$) au-

*Work of Carole Delporte-Gallet and Hugues Fauconnier was supported by grant ANR-08-VERSO-SHAMAN and the INRIA project GANG.

SPAA'10, June 13–15, 2010, Thira, Santorini, Greece.
ACM 978-1-4503-0079-7/10/06.

thenticated identities. In short, we do so by adding a communication round between any two communication rounds of the original agreement protocol. The added rounds are used by the processes with the same identities to agree on the same message to send in the next round of the original protocol.

2. In an eventually synchronous model [5] agreement is impossible. More specifically, we show, using a partitioning argument, that processes can decide on different values if $n \geq 2l$. The proof does not require any failure, which might be surprising at first glance.

   Interestingly, this highlights the very fact that, in a setting with limited identities, the demarcation line between the possibility and impossibility of agreement is different that in a model where all processes have identities [6].

3. In an asynchronous model, eventual agreement is possible, if (and only if) $l > 3t$. In *eventual agreement* (also called *stabilizing consensus* [1]), processes can decide several times, possibly on different values, but eventually converge to the same value.

   We first present a binary eventual agreement protocol and then show how to use a transformation close to [2, 10] to obtain a multi-valued agreement.

3. CONCLUDING REMARKS

To the best of our knowledge, this paper is the first to study a distributed system model with a limited number of identities, i.e., homonyms. In a sense, the model unifies both classical non-anonymous [3, 9] and anonymous models [1, 2, 4, 7, 8, 11, 13]. As we argued in the introduction, we believe this model to be interesting for both intellectual and practical considerations.

While studying the spectrum of possibilities for Byzantine agreement in a system with homonyms, we shed a different light on this problem. Whereas the traditional agreement impossibility result holds only in a fully asynchronous model only, we show than in a system with homonyms, the impossibility holds as soon as there is some (even partial) asynchrony. In addition, in a synchronous setting, the relevant crucial ratio turns out to be between the number of identities l and the number of Byzantine processes t, while the classical result considers a ratio between n and t.

It is important to note that we only scratched the surface of what can be computed with homonyms: many challenging issues are open. For instance, we considered a message passing model and it would be interesting to explore the impact of a shared memory. Also, we focused on agreement and many other problems should be considered. Finally, we focused on computability, complexity is yet to be explored.

Acknowledgments

We are grateful to Christian Cachin for his useful comments on our model with homonyms and to the reviewers for their helpful comments.

4. REFERENCES

[1] Dana Angluin, Michael J. Fischer, and Hong Jiang. Stabilizing consensus in mobile networks. In *DCOSS*, volume 4026 of *LNCS*, pages 37–50, 2006.

[2] Hagit Attiya, Alla Gorbach, and Shlomo Moran. Computing in totally anonymous asynchronous shared memory systems. *Inf. Comput.*, 173(2):162–183, 2002.

[3] Hagit Attiya and Jennifer Welch. *Distributed Computing: fundamentals, simulations and advanced topics, 2nd edition*. Wiley, 2004.

[4] Harry Buhrman, Alessandro Panconesi, Riccardo Silvestri, and Paul M. B. Vitányi. On the importance of having an identity or, is consensus really universal? *Distributed Computing*, 18(3):167–176, 2006.

[5] Cynthia Dwork, Nancy A. Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, April 1988.

[6] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985.

[7] Rachid Guerraoui and Eric Ruppert. Anonymous and fault-tolerant shared-memory computing. *Distributed Computing*, 20(3):165–177, 2007.

[8] Rachid Guerraoui and Eric Ruppert. Names trump malice: Tiny mobile agents can tolerate byzantine failures. In *ICALP*, volume 5556 of *LNCS*, pages 484–495. Springer, 2009.

[9] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

[10] Achour Mostéfaoui, Michel Raynal, and Frederic Tronel. From binary consensus to multivalued consensus in asynchronous message-passing systems. *Inf. Process. Lett.*, 73(5-6):207–212, 2000.

[11] Michael Okun and Amnon Barak. Efficient algorithms for anonymous byzantine agreement. *Theory Comput. Syst.*, 42(2):222–238, 2008.

[12] Antony I. T. Rowstron and Peter Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware*, volume 2218 of *LNCS*, pages 329–350, 2001.

[13] Eric Ruppert. The anonymous consensus hierarchy and naming problems. In *OPODIS*, volume 4878 of *LNCS*, pages 386–400. Springer, 2007.

[14] Ion Stoica, Robert Morris, David R. Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *ACM SIGCOMM*, pages 149–160, 2001.

[15] Xiaoyun Wang and Hongbo Yu. How to break md5 and other hash functions. In *EUROCRYPT*, volume 3494 of *LNCS*, pages 19–35, 2005.