

# Building and verifying a **quasi-certification entity** over Distributed Hash Tables

Join work with X. Bonnaire, R. Cortes & O. Marin UPMC (France), UFSM (Chile), NYUS (China)

Fabrice.Kordon@lip6.fr





## **B** Certification, why?

- Motivations
  - Digital filling for tax purpose
    - Certify that somebody did it before a given deadline
  - Certified emails

Use emails for legal purposes
 Online game refereeing
 e-voting, etc.



#### **Certification, why?** 6

- Motivations
  - Digital filling for tax purpose
    - Certify that somebody did it before a given deadline
  - Certified emails
  - Use emails for legal purposes Online game refereeing
  - e-voting, etc.

### **Existing Solutions**

- Centralized
  - Public Key Infrastructures (traditional PKI)
  - Scaling problem/prone to faults/implementation (atomic multicast)
- Decentralized
  - Certification on top of Distributed Hash Tables (DHT)
  - Rapidly brings Byzantine consensus (for a 100.0% guarantee)



#### **Certification, why?** 6



- Digital filling for tax purpose
  - Certify that somebody did it before a given deadline
- Certified emails
- Use emails for legal purposes Online game refereeing
- e-voting, etc.

(quasi)-certify that a given action has been performed at a certain time Distributed context (DHT)

### **Existing Solutions**

- Centralized
  - Public Key Infrastructures (traditional PKI)
  - Scaling problem/prone to faults/implementation (atomic multicast)
- Decentralized
  - Certification on top of Distributed Hash Tables (DHT)
  - Rapidly brings Byzantine consensus (for a 100.0% guarantee)

### Objective





Totally decentralized + built)in

### Retrieve data (key + value)

 $\Rightarrow$  put (v,k)  $\Rightarrow$  get(k)  $\rightarrow$  v

![](_page_6_Picture_4.jpeg)

### Retrieve data (key + value)

 $\Rightarrow$  put (v,k)  $\Rightarrow$  get(k)  $\rightarrow$  v

![](_page_7_Picture_4.jpeg)

### Leafset L close nodes (+ root)

### Retrieve data (key + value)

 $\Rightarrow$  put (v,k)  $\Rightarrow$  get(k)  $\rightarrow$  v

![](_page_8_Picture_4.jpeg)

### Classical values for L

8, 16, 32 (best)

### Leafset L close nodes (+ root)

![](_page_9_Picture_0.jpeg)

 $\bigcirc$  A  $\rightarrow$  an actor performing a service

![](_page_9_Picture_2.jpeg)

![](_page_9_Picture_4.jpeg)

## **Quasi-certification** — entities 6 $\bigcirc$ A $\rightarrow$ an actor performing a service $\bigcirc$ S $\rightarrow$ leafset hash(service) offering the service

![](_page_10_Figure_1.jpeg)

![](_page_10_Picture_3.jpeg)

## **Quasi-certification** — entities 6 $\bigcirc$ A $\rightarrow$ an actor performing a service $\bigcirc$ S $\rightarrow$ leafset hash(service) offering the service

![](_page_11_Figure_1.jpeg)

![](_page_11_Picture_3.jpeg)

# $\bigcup_{i=1}^{n} \mathbf{G}_{i}$ Quasi-certification — entities $\bigcup_{i=1}^{n} \mathbf{A} \rightarrow \text{an actor performing a service}$

- $\bigcirc$  S  $\rightarrow$  leafset hash(service) offering the service

![](_page_12_Figure_3.jpeg)

# the service hash(A/service)

### 3 - transaction ack

С

## **Quasi-certification** — entities 6 $\bigcirc$ A $\rightarrow$ an actor performing a service

- $\bigcirc$  S  $\rightarrow$  leafset hash(service) offering the service
- $\bigcirc$  C  $\rightarrow$  certification authority leafset hash(A/service)

![](_page_13_Figure_3.jpeg)

# - transaction ack 3 С - certificate generation Certificate Log Entry

Д

![](_page_14_Picture_3.jpeg)

![](_page_15_Figure_1.jpeg)

![](_page_15_Picture_4.jpeg)

![](_page_16_Figure_1.jpeg)

![](_page_16_Picture_4.jpeg)

![](_page_17_Figure_1.jpeg)

![](_page_18_Figure_1.jpeg)

![](_page_19_Figure_1.jpeg)

& M. Curie - CC2016

- Université

F. Kordon

![](_page_20_Figure_1.jpeg)

#### The verification process 6

![](_page_21_Picture_1.jpeg)

#### Proven to be undecidable [FLP 85]

AND

Abstract. The consensus problem involves an asynchronous system of processes, some of which may be unreliable. The problem is for the reliable processes to agree on a binary value. In this paper, it is shown that every protocol for this problem has the possibility of nontermination, even with only one faulty process. By way of contrast, solutions are known for the synchronous case, the "Byzantine Generals"

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocolsprotocol architecture; C.2.4 [Computer-Communication Networks]: Distributed Systems-distributed applications; distributed databases; network operating systems; C.4 [Performance of Systems]: Reliability, Availability, and Serviceability; F.1.2 [Computation by Abstract Devices]: Modes of Computationparallelism; H.2.4 [Database Management]: Systems-distributed systems; transaction processing

& M. Curie - CC2016 Kordon - Université P. щ

#### Impossibility of Distributed Consensus with One Faulty Process

#### MICHAEL J. FISCHER

Yale University, New Haven, Connecticut

#### NANCY A. LYNCH

Massachusetts Institute of Technology, Cambridge, Massachusetts

#### MICHAEL S. PATERSON

University of Warwick, Coventry, England

![](_page_21_Picture_26.jpeg)

General Terms: Algorithms, Reliability, Theory

Additional Key Words and Phrases: Agreement problem, asynchronous system, Byzantine Generals problem, commit problem, consensus problem, distributed computing, fault tolerance, impossibility

#### 1. Introduction

The problem of reaching agreement among remote processes is one of the most fundamental problems in distributed computing and is at the core of many

Editing of this paper was performed by guest editor S. L. Graham. The Editor-in-Chief of JACM did not participate in the processing of the paper.

This work was supported in part by the Office of Naval Research under Contract N00014-82-K-0154, by the Office of Army Research under Contract DAAG29-79-C-0155, and by the National Science Foundation under Grants MCS-7924370 and MCS-8116678.

This work was originally presented at the 2nd ACM Symposium on Principles of Database Systems,

Authors' present addresses: M. J. Fischer, Department of Computer Science, Yale University, P.O. Box 2158, Yale Station, New Haven, CT 06520; N. A. Lynch, Laboratory for Computer Science, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139; M. S. Paterson, Department of Computer Science, University of Warwick, Coventry CV4 7AL, England

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. © 1985 ACM 0004-5411/85/0400-0374 \$00.75

Journal of the Association for Computing Machinery, Vol. 32, No. 2, April 1985, pp. 374-382.

## **b** The verification process

![](_page_22_Picture_1.jpeg)

### Proof (by any method?)

Proven to be undecidable [FLP 85]

### So what?

Being pragmatic Going for «quasi»

![](_page_22_Picture_6.jpeg)

#### Impossibility of Distributed Consensus with One Faulty Process

#### MICHAEL J. FISCHER

Yale University, New Haven, Connecticut

#### NANCY A. LYNCH

Massachusetts Institute of Technology, Cambridge, Massachusetts

#### AND

#### MICHAEL S. PATERSON

University of Warwick, Coventry, England

Abstract. The consensus problem involves an asynchronous system of processes, some of which may be unreliable. The problem is for the reliable processes to agree on a binary value. In this paper, it is shown that every protocol for this problem has the possibility of nontermination, even with only one faulty process. By way of contrast, solutions are known for the synchronous case, the "Byzantine Generals"

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocolsprotocol architecture; C.2.4 [Computer-Communication Networks]: Distributed Systems-distributed applications; distributed databases; network operating systems; C.4 [Performance of Systems]: Reliability, Availability, and Serviceability; F.1.2 [Computation by Abstract Devices]: Modes of Computationparallelism; H.2.4 [Database Management]: Systems-distributed systems; transaction processing

General Terms: Algorithms, Reliability, Theory

Additional Key Words and Phrases: Agreement problem, asynchronous system, Byzantine Generals problem, commit problem, consensus problem, distributed computing, fault tolerance, impossibility

#### 1. Introduction

The problem of reaching agreement among remote processes is one of the most fundamental problems in distributed computing and is at the core of many

Editing of this paper was performed by guest editor S. L. Graham. The Editor-in-Chief of JACM did not participate in the processing of the paper.

This work was supported in part by the Office of Naval Research under Contract N00014-82-K-0154, by the Office of Army Research under Contract DAAG29-79-C-0155, and by the National Science Foundation under Grants MCS-7924370 and MCS-8116678. This work was originally presented at the 2nd ACM Symposium on Principles of Database Systems,

Authors' present addresses: M. J. Fischer, Department of Computer Science, Yale University, P.O. Box 2158, Yale Station, New Haven, CT 06520; N. A. Lynch, Laboratory for Computer Science, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139; M. S. Paterson, Department of Computer Science, University of Warwick, Coventry CV4 7AL, England

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. © 1985 ACM 0004-5411/85/0400-0374 \$00.75

Journal of the Association for Computing Machinery, Vol. 32, No. 2, April 1985, pp. 374-382.

#### The verification process 6

![](_page_23_Picture_1.jpeg)

### Proof (by any method?)

Proven to be undecidable [FLP 85]

### So what?

Being pragmatic Going for «quasi»

![](_page_23_Picture_6.jpeg)

### **Two steps**

- 1. modeling the protocol in a perfect world (no error) **Use of Petri nets**
- 2. Probabilistic analysis to evaluate the failure rate Use of a classical fault model, building a formula + numeric evaluation

#### Impossibility of Distributed Consensus with One Faulty Process

#### MICHAEL J. FISCHER

Yale University, New Haven, Connecticut

#### NANCY A. LYNCH

Massachusetts Institute of Technology, Cambridge, Massachusetts

#### AND

#### MICHAEL S. PATERSON

University of Warwick, Coventry, England

Abstract. The consensus problem involves an asynchronous system of processes, some of which may be unreliable. The problem is for the reliable processes to agree on a binary value. In this paper, it is shown that every protocol for this problem has the possibility of nontermination, even with only one faulty process. By way of contrast, solutions are known for the synchronous case, the "Byzantine Generals"

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocolsprotocol architecture; C.2.4 [Computer-Communication Networks]: Distributed Systems-distributed applications; distributed databases; network operating systems; C.4 [Performance of Systems]: Reliability, Availability, and Serviceability; F.1.2 [Computation by Abstract Devices]: Modes of Computationparallelism; H.2.4 [Database Management]: Systems-distributed systems; transaction processing

General Terms: Algorithms, Reliability, Theory

Additional Key Words and Phrases: Agreement problem, asynchronous system, Byzantine Generals problem, commit problem, consensus problem, distributed computing, fault tolerance, impossibility

#### 1. Introduction

The problem of reaching agreement among remote processes is one of the most fundamental problems in distributed computing and is at the core of many

Editing of this paper was performed by guest editor S. L. Graham. The Editor-in-Chief of JACM did not participate in the processing of the paper.

This work was supported in part by the Office of Naval Research under Contract N00014-82-K-0154, by the Office of Army Research under Contract DAAG29-79-C-0155, and by the National Science Foundation under Grants MCS-7924370 and MCS-8116678.

This work was originally presented at the 2nd ACM Symposium on Principles of Database Systems,

Authors' present addresses: M. J. Fischer, Department of Computer Science, Yale University, P.O. Box 2158, Yale Station, New Haven, CT 06520; N. A. Lynch, Laboratory for Computer Science, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139; M. S. Paterson, Department of Computer Science, University of Warwick, Coventry CV4 7AL, England

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. © 1985 ACM 0004-5411/85/0400-0374 \$00.75

Journal of the Association for Computing Machinery, Vol. 32, No. 2, April 1985, pp. 374-382.

Hypotheses

H1: perfect world H<sub>2</sub>: service reduced to 1 interaction H<sub>3</sub>: L+1 answer requested instead of L/2+1 Symmetric net with Bags?

![](_page_24_Picture_4.jpeg)

![](_page_24_Picture_5.jpeg)

### **Types and variables**

type tsid is 0...L; type tsidxtsid is <tsid, tsid>;

var i in tsid;

![](_page_25_Picture_5.jpeg)

![](_page_25_Picture_6.jpeg)

![](_page_26_Figure_1.jpeg)

![](_page_27_Figure_1.jpeg)

![](_page_28_Figure_1.jpeg)

![](_page_29_Figure_1.jpeg)

CC2016

.

& M. Curie

٩.

Kordon - Université

щ

![](_page_30_Figure_1.jpeg)

.

٩.

Kordon - Université

щ

![](_page_31_Figure_1.jpeg)

![](_page_32_Figure_1.jpeg)

$$F_{abort} = |SstopAbort| > \frac{L}{2} \lor |CstopAbort| > \frac{L}{2} \lor |CstopOK| = L + 1 \land |CstopOK| = L + 1 \land |CstopAbort| > 0 \lor |cst$$

![](_page_33_Figure_4.jpeg)

![](_page_34_Figure_1.jpeg)

### About the complexity of the state space

- Roughly 10<sup>L</sup> states
- 21 state= 13 int + 17 multistep  $\Rightarrow$  memory problem to check for L=32

![](_page_35_Figure_5.jpeg)

![](_page_35_Picture_6.jpeg)

- About the complexity of the state space
  - Roughly 10<sup>L</sup> states
  - $\stackrel{\scriptstyle <}{\scriptstyle \sim}$  1 state= 13 int + 17 multistep  $\Rightarrow$  memory problem to check for L=32

### GreatSPN (use of symmetries)

- L=24 fails after 11h45mn of CPU (costly canonisation function)
  - Implementation limit = size of a hash table

![](_page_36_Figure_7.jpeg)

![](_page_36_Picture_9.jpeg)

- About the complexity of the state space
  - Roughly 10<sup>L</sup> states
  - 2 1 state= 13 int + 17 multistep  $\Rightarrow$  memory problem to check for L=32

### GreatSPN (use of symmetries)

- L=24 fails after 11h45mn of CPU (costly canonisation function)
  - Implementation limit = size of a hash table

### **PNXDD** (decision diagrams + variable ordering)

- Unfolding to P/T nets
- $\downarrow$  L=10 fails after 3h20mn (memory overflow > 16GB)

![](_page_37_Picture_11.jpeg)

- About the complexity of the state space
  - Roughly 10<sup>L</sup> states
  - 21 state= 13 int + 17 multistep  $\Rightarrow$  memory problem to check for L=32

### GreatSPN (use of symmetries)

- L=24 fails after 11h45mn of CPU (costly canonisation function)
  - Implementation limit = size of a hash table
  - **PNXDD** (decision diagrams + variable ordering)
- Unfolding to P/T nets
- L=10 fails after 3h20mn (memory overflow > 16GB)
- **ITS-Tools (hierarchical decision diagrams)**
- Completed for L=32 in less than one minute
  - Handling of symmetries in the system by means of a dedicated encoding

![](_page_38_Picture_14.jpeg)

![](_page_39_Picture_1.jpeg)

### http://cosyverif.org

![](_page_39_Picture_3.jpeg)

![](_page_39_Picture_4.jpeg)

![](_page_39_Picture_6.jpeg)

# Probabilistic analysis (step 2)

### Classical approach of the domain

- Based on p, probability of node failure
  Hypotheses required
  - Diversity routing to avoid coalitions

![](_page_40_Picture_5.jpeg)

## B Probabilistic analysis (step 2)

### Classical approach of the domain

- Based on **p**, probability of node failure
- Hypotheses required Diversity routing to avoid coalitions

### Origin of problems

- Source 1  $\rightarrow$  failure of the protocol
  - No answer to A + no ack to A
  - Interactions between A, S (leafset)
- Source 2  $\rightarrow$  inappropriate certificate
  - Lost of a certificate (inconsistency)
  - Interactions between S (leafset) and C (leafset)

![](_page_41_Picture_12.jpeg)

## B Probabilistic analysis (step 2)

### Classical approach of the domain

- Based on **p**, probability of node failure **Hypotheses required**
- Diversity routing to avoid coalitions

![](_page_42_Picture_4.jpeg)

### Origin of problems

- Source 1  $\rightarrow$  failure of the protocol
  - No answer to A + no ack to A
  - Interactions between A, S (leafset)
- Source 2  $\rightarrow$  inappropriate certificate
  - Lost of a certificate (inconsistency)
  - Interactions between S (leafset) and C (leafset)

![](_page_42_Picture_13.jpeg)

![](_page_43_Figure_1.jpeg)

- More that L/2 nodes are malicious
- Fre formula:

at most L+1

at most L/2  $\underbrace{\text{malicious nodes}}_{\sum_{i=1}^{L+1} {L+1 \choose i} p^i (1-p)^{L+1-i}} - \underbrace{\sum_{i=1}^{\frac{L}{2}} {L+1 \choose i} p^i (1-p)^{L+1-i}}_{i=1}$ 

![](_page_43_Picture_7.jpeg)

![](_page_43_Picture_8.jpeg)

![](_page_44_Figure_1.jpeg)

![](_page_44_Picture_3.jpeg)

![](_page_45_Figure_1.jpeg)

### Inappropriate certificate generation

🗳 Two parts

Existence of more L/2 malicious nodes

Unable to retrieve at least L/2 + 1 identical answers

The formula (combines problems between S and C)

$$1 - (1 - P_{>\frac{L}{2}})$$

![](_page_45_Picture_9.jpeg)

![](_page_46_Figure_1.jpeg)

### Inappropriate certificate generation

🗳 Two parts

Existence of more L/2 malicious nodes

Unable to retrieve at least L/2 + 1 identical answers

The formula (combines problems between S and C)

$$1 - (1 - P_{>\frac{L}{2}})^2$$
  
$$1 - \left(1 - \sum_{i=1}^{L+1} {L+1 \choose i} p^i (1-p)^{L+1-i} + \sum_{i=1}^{\frac{L}{2}} {L+1 \choose i} p^i (1-p)^{L+1-i} \right)$$

![](_page_46_Picture_9.jpeg)

 $p^i(1-p)^{L+1-i}\Big)^2$ 

![](_page_47_Figure_1.jpeg)

8

16

32

- 🗳 Two parts
  - Existence of more L/2 malicious nodes
  - Unable to retrieve at least L/2 + 1 identic
- The formula (combines problems b)

$$1 - (1 - P_{>\frac{L}{2}})^2$$
  
$$1 - \left(1 - \sum_{i=1}^{L+1} {L+1 \choose i} p^i (1-p)^{L+1-i} + \sum_{i=1}^{\frac{L}{2}} {L+1 \choose i} p^i (1-p)^{L+1-i} \right)$$

![](_page_47_Picture_8.jpeg)

<b>HT</b> - $p = 0.3$	CODDC
0.216	<b>CORPS -</b> $p = 0.05$
0.075	$\frac{4.97 \times 10^{-4}}{3.50 \times 10^{-8}}$
0.014	$\frac{3.30 \times 10^{-8}}{4.24 \times 10^{-13}}$
	1.24 × 10 10

10

 $p^i(1-p)^{L+1-i}\Big)^2$ 

# **B** Conclusion

![](_page_48_Picture_1.jpeg)

Low probability of failure + Good message complexity (not discussed)

# stion + verification) complexity (not discussed)

## **B** Conclusion

### Quasi-certification entity (elaboration + verification)

- Low probability of failure + Good message complexity (not discussed)
   Application to digital tax filling in France
  - ▶ 36.5 M revenues declarations (in 2012)
  - ▶ a wrong tax certificate every 5132 years
  - Iost of a tax certificate every 997 years

# stion + verification) complexity (not discussed)

#### Conclusion 6

![](_page_50_Picture_1.jpeg)

### Quasi-certification entity (elaboration + verification)

- Low probability of failure + Good message complexity (not discussed) Application to digital tax filling in France
  - ▶ 36.5 M revenues declarations (in 2012)
  - a wrong tax certificate every 5132 years
  - Iost of a tax certificate every 997 years

### 3 years of work (details recently fixed)

- Work partially published in TrustCom'2013
- Without formal proof (there was yet glitches details with hypotheses)

#### Conclusion 6

![](_page_51_Picture_1.jpeg)

### Quasi-certification entity (elaboration + verification)

- Low probability of failure + Good message complexity (not discussed) Application to digital tax filling in France
  - ▶ 36.5 M revenues declarations (in 2012)
  - a wrong tax certificate every 5132 years
  - Iost of a tax certificate every 997 years

### 3 years of work (details recently fixed)

- Work partially published in TrustCom'2013
- Without formal proof (there was yet glitches details with hypotheses)

### Realistic problem with applicability to e-government

- Probably numerous applications in the future
- The model is now a metric for the Model Checking Contest
- Potential applicability for Symmetric Nets with Bags
  - Excellent playground for this type of problems

![](_page_52_Picture_0.jpeg)