The Byzantine Empire and Its Generals: An Ancient Empire Back to Life in Computer Security

Pedro Reviriego, Universidad Politécnica de Madrid
Elena Merino-Gómez, Universidad de Valladolid
Fabrizio Lombardi, Northeastern University

Forty years after its initial publication, we revisit the seed contribution of Byzantine fault tolerance, focusing on its application for the security of systems implemented in software. We describe new environments in which it is being used.

Digital Object Identifier 10.1109/MC.2023.3235095
Date of current version: 8 March 2023

Security and dependability are key requirements for most of today's computing systems, and their importance is poised to grow as we increasingly rely on their pervasive use in almost every aspect of our lives. At the same time, the complexity of computing systems unabatedly continues to grow with many different organizations providing interdependent components that, in turn, coordinate to implement services. In this scenario, making sure that such a system will work reliably when some of the components or nodes fail or are compromised by an attacker becomes critical. For example, several attacks, like Spectre or Meltdown, have been investigated; they exploit the advanced mechanisms of modern processors to extract information. Similarly, in the recent SolarWinds attack, a software tool was compromised, and then automatic updates

0018-9162/23©2023IEEE

EDITOR **PHIL LAPLANTE**
The Pennsylvania State University;
plaplante@psu.edu

were exploited to disseminate the infected version. Failures can also disrupt the operation of computing systems in many complex ways. For example, radiation-induced soft errors can flip any bit stored in a memory or register, leading to silent data corruption that can manifest in erratic system behavior.

From a design perspective, in many cases, the same mechanisms can be used to mitigate both attacks and failures. In fact, many of the models commonly used for secure and dependable system design cover both scenarios. This is the case with the Byzantine Generals problem formulated more than 40 years ago; this has led to the concept of Byzantine fault tolerance (BFT), which has found widespread adoption in many technical domains.[1]

In this article, we revisit BFT four decades after its introduction, focusing on software implementations and briefly discussing how it is now being used in new systems, domains, and applications. We also look back to the Byzantine empire to understand how it survived for one millennium and how its history relates to the Byzantine generals problem. This discussion links computing with history and shows that the choice made by the authors for the generals is, in an unintended way, backed by facts.

An analogy with a group of generals who have to act consistently in taking the decision to attack or retreat has been used in Lamport et al.[1] to provide a model for secure and dependable system design. The generals can be loyal or traitors, and there can also be communication failures or restrictions among generals. This models a computing system in which some nodes may have been compromised and in which failures could also either disable nodes or prevent them from communicating.

As has been the case with other famous problems in computing, the analogy can facilitate the understanding of the problem and the algorithms used to solve it. Indeed, formulating the problem with an appealing analogy was one of the objectives of the authors of the article [see http://lamport.azurewebsites.net/pubs/pubs.html#byz (46. The Byzantine Generals Problem)]. Apparently, they chose the generals to be Byzantine to avoid offending any nationality, and thus, the model became the Byzantine Generals problem.[1] From then on, systems and algorithms that can solve this problem and work consistently in that scenario are known as Byzantine fault tolerant (BFT). Hence, BFT has become a key concept in dependable and secure system design.

After presenting the initial model, different scenarios, for example by considering that messages exchanged by the generals can be forged by traitors or, conversely, that they are signed and thus cannot be forged, are analyzed in Lamport et al.[1] This has led to a fundamental result; for the group of loyal generals to act consistently, there can be at most m traitors in a group of $3m + 1$ generals when messages can be manipulated by traitors. For example, when there is a single traitor, there have to be at least three loyal generals for them to act consistently. This illustrates the high cost of building systems that can tolerate failures or attacks; not only are $3m + 1$ generals needed, but they must also exchange a sequence of messages recursively to reach a consensus on the action to take.

By presenting the problem in a general manner, considering different scenarios with signed or oral messages and with failures or restrictions in the communications among the generals, the article instantiated a framework for the analysis and design of fault-tolerant systems that has been used in a myriad of applications and designs. Initially, the concept was used for safety-critical applications such as space systems; avionics; military equipment; or industrial and nuclear control systems. However, its adoption has extended to almost every domain in computing. For example, BFT is a key element in many blockchain-based systems, and in particular for cryptocurrencies, to ensure that a group of completely independent nodes can maintain a consistent state. This has motivated a large body of research throughout the years in this area to ensure that consensus can be achieved reliably in systems that involve large numbers of nodes and transactions. These efforts have led to the development of new consensus algorithms, such as, for example, proof of work and proof of stake.[2]

The game with the nationality of the generals seems to continue. For example, three of the main forks in Ethereum (see Figure 1) are named *Byzantium*, *Constantinople*, and *Istanbul*, which does not seem to be a coincidence, and it is likely a play with words and a tribute to the Byzantine generals problem.

Before discussing other areas in which BFT is currently being used, let us go back in time and look at the empire that gives the name to the problem and concept. For more than one millennium, the Byzantine empire was able to survive despite having to face powerful enemies from the East and West. Therefore, in a way, the empire can itself be seen as a complex resilient system from the outset. Most of its rulers had a solid military background; lineage was not a sufficient condition—sometimes not even necessary. For emperors and coemperors to be considered worthy to wear the imperial purple, they had to have demonstrated their ability as generals, which partially explains the strategic strength the empire enjoyed for centuries. Right from the start, the binomial formed by Justinian and his general

FIGURE 1. The Ethereum logo.

Belisarius (Figure 2) paved the way that would forever associate military power with imperial power, often in a single person, a circumstance that would be essential in defensive strategies.

In addition to the skill of many of its rulers, defensive resources, such as the chain of the Golden Horn (Figure 3) or the one known as *Greek fire* (Figure 4), were fundamental for the Byzantine resistance to attacks throughout its history. It is significant that even in the case of war tools from the past, they continue to pose enigmas in the present. The chemical composition of Greek fire, a kind of liquid fire capable of spreading over water and easily reaching enemy ships, has not yet been formulated. The precise mechanism that allowed the closing of the Golden Horn by means of a heavy chain that was pulled up to the surface is even today the subject of speculation. In addition to the scientific challenge of unraveling how these tools worked, their efficient performance can be inspiring for engineers who are currently developing cyberdefense systems.

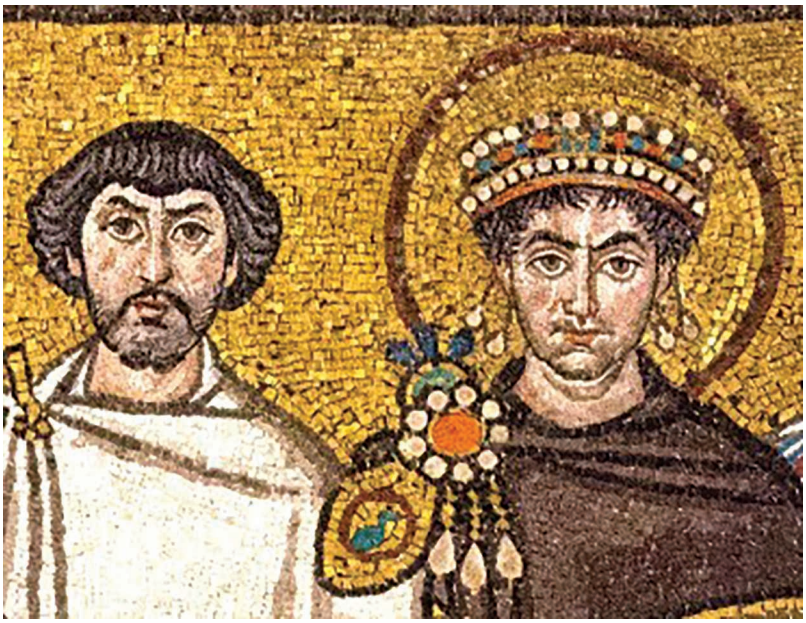

FIGURE 2. A portrait possibly of General Belisarius and Emperor Justinian (from a mosaic in San Vitale, Ravenna.)

FIGURE 3. The iron chain prevents the fleet of Thomas the Slav from entering the Golden Horn. (Source: Biblioteca Nacional de España.)

The adoption of the term *Byzantine* to describe the problem of the Byzantine Generals is a cautious decision that contrasts with the apparently inappropriate use to denominate the Eastern Roman Empire. Leslie Lamport assigned the Byzantine nationality so as not to offend any reader in the certainty that an extinct empire was a safe bet. However, the name "Byzantine" turned out to be one of the most controversial of the historical empires. Curiously, the Byzantines themselves would surely have felt bothered with the name attributed to them to avoid calling them "Romans," as they considered themselves. The name "Byzantine" to name the Eastern Roman Empire is subsequent to the disappearance of the empire itself. The origin of the name and its connotations are still part of a scientific discussion today.[3]

Lamport also confesses that he took the idea of the generals from the problem in distributed computing that is sometimes called the *Chinese Generals problem*, "in which two generals have to come to a common agreement on whether to attack or retreat, but can communicate only by sending messengers who might never arrive."[1] The idea of going back to the past arises intuitively at the moment in which the figure of the messenger appears. The times when messages could be transmitted only through intermediaries

seem to contrast with the current situation in which the immediacy of sending and receiving can give the impression that the endpoints of the communication are enough. However, it is obvious, although imperceptible, that the messenger, which connects the sender and receiver, continues to be present, merged, or frequently confused with the channel.

The immediacy of the transmission of information, so common nowadays, can lead to a feeling of false security in the minds of communicators. It is possible to mistakenly perceive that the instantaneity and the apparent absence of intermediaries guarantee the veracity of the message. The speed of communication does not seem to offer enough time to intentionally alter the message. In a similar way, immediacy seems to establish a direct thread with the addressee, without intermediaries. However, messages and the channels through which they travel can be as insecure today as they were 500 years ago. The messages today must go through a myriad of hardware and software components and systems before reaching their destination. Not only can the senders or recipients of the message be malicious, but all those complex elements can also be manipulated to interfere with and disrupt communication, opening a vast attack surface. In fact, the security of communications has been a critical issue since the beginning of its existence and is one of the oldest problems in the history of communication.

The Byzantine Empire offers heroic examples of safe message delivery. One of the most thrilling episodes in the transmission of a message occurred only a few days before the Fall of Constantinople (see Figure 5). Constantine XI Palaiologos, the last Byzantine emperor, urgently needed to know if more reinforcements from Venice would arrive in Constantinople. Without them, the city was doomed. Twelve men trusted by Constantine embarked in a small brigantine, with a false flag, disguised as Turks, toward the Aegean Sea, to see if the necessary help was approaching. The 12 messengers found that there was no help on the horizon. They were to return to Constantinople to deliver their hopeless message. One of the messengers proposed to the others to continue toward Christian lands to save their own lives. Returning to Constantinople to deliver the message to the emperor meant certain death. However, loyalty prevailed, and they returned to the city to transfer the information even at the cost of their lives.[4] Although the precious message arrived uncorrupted on 23 May 1453, it was already too late for everything,

FIGURE 4. An illustration of Greek fire, from J. Scylitzes (flourished), *History of Byzantium*. (Source: Biblioteca Nacional de España.)

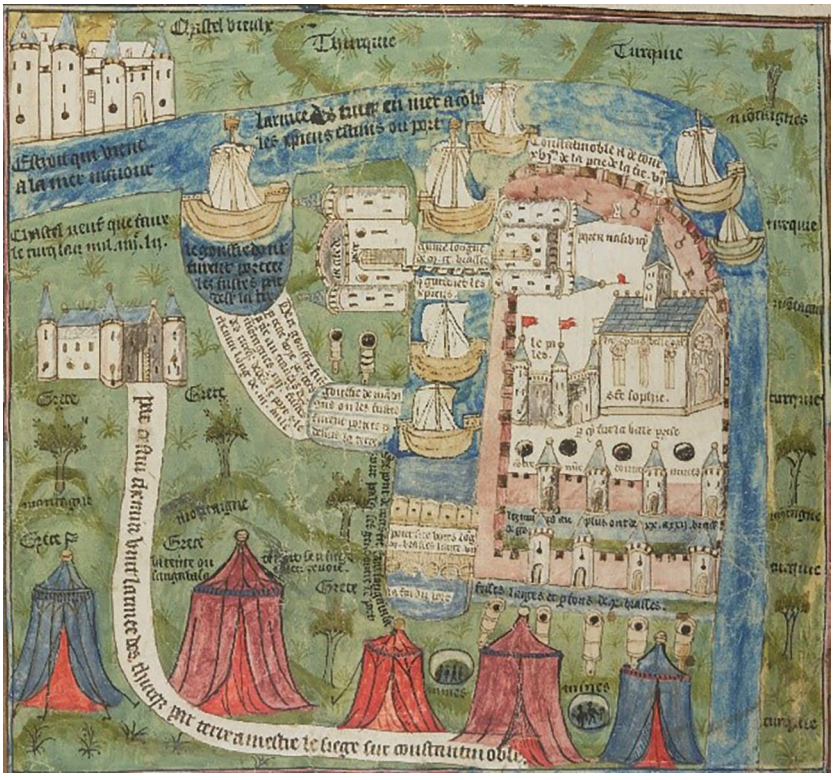

FIGURE 5. A large miniature depicting a view of besieged Constantinople from Jacques Tedaldi, *Recueil de textes historiques – Récit de la prise de Constantinople* (1453). (Source: Bibliothèque nationale de France.)

and just six days later, the city would fall, and with it, the last stronghold of the empire.

Although the error-free transmission of this last message could do nothing to prevent the final Fall of Constantinople, throughout the history of the Byzantine Empire, there are other kinds of episodes that exemplify how the system was able to survive despite the spread of erroneous messages. One of the versions of the events that occurred in the famous Battle of Manzikert in 1071, between Byzantines and Seljuk Turks, illustrates how false messages about the defeat of Emperor Romanus IV Diogenes were spread. It is possible that the jealous Byzantine general Andronikos Doukas, belonging to the family that had ruled Byzantium in the previous generation, took advantage of the confusion in the transmission of a message to abandon the emperor in battle.

Romanus IV Diogenes gave the signal to bring the pursuit against the Turks to a halt,[5] fearing an ambush. The message was misinterpreted in the rear guard. The order to return to the camp was interpreted as a withdrawal, and it was deduced that the emperor had fallen in his advance against the enemies. Some argue that the rumor was actually started by Andronikos Doukas, who did not forgive Romanus IV for having interrupted the succession of the house of Doukas to the throne of Byzantium. The contaminated message of the fall of Emperor Romanus IV Diogenes caused the abandonment of his people and determined that, indeed, he was finally captured by the Seljuk Sultan Alp Arslan. Although the defeat of Manzikert, due in part to the spread of an adulterated message, is considered as one of the greatest disasters in the history of the empire, the Byzantines continued to persist.

Seven decades later, the transmission of a corrupted message once again put another emperor in trouble. During the maneuvers to recapture Antioch for the Byzantine Empire, Emperor John II Komnenos was betrayed by two apparent allies: Raymond of Antioch and Joscelin II, Count of Edessa. The latter sent secret messengers to spread the false message to the citizens of Antioch that Emperor John II Komnenos wanted to harm them. The rumor that Antioch had been sold to the Byzantine Greeks and that the citizens should leave their homes forced one of the greatest emperors in the history of Byzantium to leave Antioch in 1142.[6] However, the empire still had more than three centuries to live.

After 40 years, security and dependability have become critical design requirements, and BFT has been used in a myriad of systems, domains, and applications.

The problem of the Byzantine Generals is still valid 40 years later, and its name, even if it was adopted to avoid potentially more controversial terms, demonstrates its relevance in the face of the many examples that can be extracted from the long history of Byzantium. It would also be interesting to consider some of the weak points of Byzantine history to name possible security flaws. For example, according to Doukas, a contemporary historian of the Fall of Constantinople, the door next to the circus, known as the *kerkoporta*, was left ajar, and 50 Janissaries slipped through the unattended door.[7] The chapter of the forgotten door could have been decisive for the final blow to the city on 29 May 1453. This is very similar to backdoors that are created to gain unauthorized access to computing systems today.

However, different from the Byzantine Generals problem, the link with the past was not made, thus losing the opportunity to use the term "kerkoporta" to denote security backdoors, and thus, keeping the term for the collective memory of humanity. However, there may be some hope as recently "kerkoporta" has been used to name a ransomware, so in the long run, the term may be adopted, increasing the links between the old Byzantine empire and computer security.

After 40 years, security and dependability have become critical design requirements, and BFT has been used in a myriad of systems, domains, and applications. In fact, in recent years, new scenarios for BFT have emerged. For example, due to storage and processing limitations or privacy concerns, machine learning is increasingly being implemented in multiple nodes. Typically, each node stores or generates a part of the dataset, and all nodes cooperate to implement training or inference. For example, distributing the dataset among several computing nodes in a data center can provide large speedups for training, while in Internet of Things applications, the nodes commonly operate in a decentralized manner with limited capability to exchange data.[8]

The use of several nodes creates the need for the system to operate reliably when some of the nodes fail or are compromised. A good example is federated learning, which is emerging as a technology that can enable learning from many users or sensors while preserving privacy. Basically, training is done locally without sharing the data, and the results from many devices are aggregated to obtain a model based on data from all of them. The implementation of federated learning poses challenges to developing efficient algorithms to coordinate training but also to ensure that it is robust when some of the devices fail or act maliciously. Therefore, there is a strong need to implement BFT at scale in federated learning. Different schemes have been proposed; they try, for example, to detect the

updates from malicious nodes by comparing them with those of the rest of the nodes or to reduce their impact on the aggregated result.

Distributed nodes or sensors are used not only for training but also for inference, and then again, there is a need to make sure that the system can withstand the failure or misbehavior of some of them. This can be achieved by carefully analyzing the information coming from each sensor to estimate their reliability and use them accordingly for the inference process.[9] Therefore, the trend to use distributed systems to implement both training and inference makes BFT a key element for future machine learning systems.

Networking is another area in which BFT is becoming increasingly important. For example, in software-defined networks, controllers are critical, and thus, they are typically replicated to tolerate failures. As security is also a major issue in networks, providing BFT for the control plane when some of the controllers may have been compromised is also desirable in all cases and needed for networks used for critical applications. Several schemes have been proposed to reduce the overhead of implementing BFT by first identifying the disagreement among a subset of controllers and only then activating all the controllers needed to implement BFT or to detect malicious controllers.[10]

Similarly, BFT is also fundamental in wireless sensor networks that are by nature decentralized systems and for which attackers can use sophisticated mechanisms or direct jamming to disrupt communications. The ability to send broadcast messages in real time is also critical in some systems, and thus, BFT has to be implemented.[11] In summary, networks are by nature distributed, and thus, they can suffer failures and compromised nodes, thus making BFT imperative when reliable operation is needed.

Distributed optimization, similarly to distributed machine learning, relies on different nodes to optimize a function; this can be done locally and independently at each node or in a coordinated way.[12] In all cases, there can again be faulty or malicious nodes, and thus, there is a need to implement BFT. New mechanisms to support the coordination of multiple agents to perform a given computation with BFT are being proposed[13] by using replication; such a scheme can be used as a general solution when the cost introduced by replication is acceptable.

The use of quantum technologies has also been proposed to reach an agreement between generals.[14] Soon, with computing systems moving toward more complex, distributed, and in many cases, decentralized systems, the importance of BFT is poised to keep growing. Therefore, these first 40 years seem to be only the beginning of a new Byzantine era, but this time in computer science rather than as an empire.

Looking forward, we think that using analogies when presenting new algorithms and ideas can be a powerful tool to catch the attention of the readers, enable a formulation of problems and solutions that is more general, and link computer science with other fields like history. The Byzantine Generals problem is an excellent example of how those benefits can be achieved. However, that has not always been the case. In fact, Leslie Lamport used another analogy to describe a consistency algorithm, relating it to an ancient parliament formed by part-time legislators in the Greek island of Paxos, but in this case, it seems that, at least initially, the analogy was not well received [see http://lamport.azurewebsites.net/pubs/pubs.html#lamport-paxos (123. The Part-Time Parliament)]. Therefore, as with any powerful tool, analogies should be used with caution. C

REFERENCES

1. L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, Jul. 1982, doi: 10.1145/357172.357176.
2. A. Bessani, E. Alchieri, J. Sousa, A. Oliveira, and F. Pedone, "From Byzantine replication to blockchain: Consensus is only the beginning," in *Proc. 50th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2020, pp. 424–436, doi: 10.1109/DSN48063.2020.00057.
3. A. Kaldellis, "From Rome to New Rome, from Empire to Nation-State," in *Two Romes: Rome and Constantinople in Late Antiquity*, L. Grig and G. Kelly, Eds. London, U.K.: Oxford Univ. Press, 2012, pp. 387–404.
4. N. Barbaro, *Giornale Dell'assedio di Constantinopoli*. Vienne, France: Libreria Tendler, 1856, p. 35.
5. M. Attaleiates, *The History*. (Transl.: A. Kaldellis and D. Krallis, Dumbarton Oaks Medieval Library 16). Cambridge, MA, USA: Harvard Univ. Press, 2012, p. 293.
6. W. of Tyre, *A History of the Deeds Done Beyond the Sea*, vol. 2. New York, NY, USA: Columbia Univ. Press, 1943, p. 97.
7. M. Philippides and W. K. Hanak, *The Siege and the Fall of Constantinople in 1453: Historiography, Topography, and Military Studies*. Burlington, VT, USA: Ashgate Publishing, 2011, p. 622.
8. Z. Yang, A. Gang, and W. U. Bajwa, "Adversary-resilient distributed and decentralized statistical inference and machine learning: An overview of recent advances under the byzantine threat model," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 146–159, May 2020, doi: 10.1109/MSP.2020.2973345.
9. J. Choi, Z. Hakimi, J. Sampson, and V. Narayanan, "Byzantine-tolerant inference in distributed deep intelligent system: Challenges and opportunities," *IEEE Trans. Emerg. Sel. Topics Circuits Syst.*, vol. 9, no. 3, pp. 509–519, Sep. 2019, doi: 10.1109/JETCAS.2019.2933807.
10. E. Sakic, N. Đerić, and W. Kellerer, "MORPH: An adaptive framework for efficient and Byzantine

fault-tolerant SDN control plane," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 10, pp. 2158–2174, Oct. 2018, doi: 10.1109/JSAC.2018.2869938.
11. D. Kozhaya, J. Decouchant, and P. Esteves-Verissimo, "RT-Byz-Cast: Byzantine-resilient real-time reliable broadcast," *IEEE Trans. Comput.*, vol. 68, no. 3, pp. 440–454, Mar. 1, 2019, doi: 10.1109/TC.2018.2871443.
12. L. Su and N. H. Vaidya, "Byzantine-Resilient multiagent optimization," *IEEE Trans. Autom. Control*, vol. 66, no. 5, pp. 2227–2233, May 2021, doi: 10.1109/TAC.2020.3008139.
13. R. Guerraoui and A. Maurer, "Byzantine-resilient multi-agent system," *IEEE Trans. Dependable Secure Comput.*, vol. 19, no. 6, pp. 4032–4038, Nov./Dec. 1, 2022, doi: 10.1109/TDSC.2021.3116488.
14. M. Fitzi, N. Gisin, and U. Maurer, "Quantum solution to the Byzantine agreement problem," *Phys. Rev. Lett.*, vol. 87, no. 21, Nov. 2001, Art. no. 217901, doi: 10.1103/PhysRevLett.87.217901.

PEDRO REVIRIEGO is an associate professor at the Universidad Politécnica de Madrid, 28040 Madrid, Spain. Contact him at pedro.reviriego@upm.es.

ELENA MERINO-GÓMEZ is an assistant professor at the Universidad de Valladolid, 47011 Valladolid, Spain. Contact her at elena.merino.gomez@uva.es.

FABRIZIO LOMBARDI is the International Test Conference Endowed Chair Professor at Northeastern University, Boston, MA 02215 USA. Contact him at lombardi@coe.neu.edu.

Digital Object Identifier 10.1109/MC.2023.3241503