Underscoring archival authenticity with blockchain technology

Archives have well-established practices which have been developed over years of working with analogue records. Now they face huge challenges due to the inexorable development of digital technologies. Not only is the heterogeneous nature of the records, their instability and the rapid pace of technological development a threat to the records' survival, but the ease with which digital records can be altered has put archives in a technology arms race with those parties who would seek to falsify our digital inheritance and undermine democracy.

In order to tackle these challenges, the ARCHANGEL project is breaking new ground by using blockchain to record checksums (cryptographic hashes) and other metadata derived from either scanned physical records or born-digital records to allow verification of their integrity over decade- or century-long time spans. This data is permanently preserved through peer-to-peer distribution and consensus checking without the need for a trusted third party, thereby enabling archives to prove the authenticity of the records in their custody.

Keywords

Distributed ledger technology; DLT; blockchain; trusted archives; document provenance

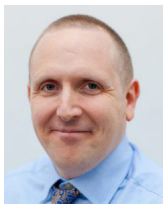

MARK BELL

Researcher
The National Archives, GB

ALEX GREEN

Senior Digital Archivist
The National Archives, GB

JOHN SHERIDAN

Digital Director
The National Archives, GB

JOHN COLLOMOSSE

Professor of Computer Vision
University of Surrey

DANIEL COOPER

Research Software Developer
University of Surrey

TU BUI

Postgraduate Research Student
University of Surrey

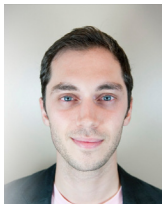

OLIVIER THEREAUX

Head of Technology
Open Data Institute

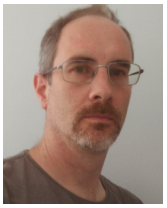

JEZ HIGGINS

Consultant Software Engineer
Open Data Institute

Introduction

Records have been preserved for thousands of years, and modern methods of preserving paper archives are well understood. Over the past two decades, society has experienced rapid technological change which has resulted in vast quantities of information being captured and stored on media other than paper. Although practices around digital preservation have developed over the past 25 years,[1] many of them attempt to replicate archival practices designed for paper collections. To deal with the threats that new technologies pose to the longevity of our digital heritage, archives must be far more willing to experiment with new technologies. The ARCHANGEL project is exploring the possibilities offered by distributed ledger technology (DLT, commonly known as blockchain) and how it could address the challenges around trust, integrity and authenticity that preserving born-digital material introduces. We will begin by describing aspects of paper preservation and the analogous digital preservation processes.

'ARCHANGEL … is exploring the possibilities offered by DLT (commonly known as blockchain)'

Paper and digital preservation practices

The repositories at The National Archives contain around 200 km of shelving holding millions of original paper documents. These documents are kept in very precise atmospheric conditions, with tightly controlled temperature and humidity. Similarly, digital documents are held in conditions conducive to long-term preservation, minimizing the degradation of tapes and spinning disks. Where digital documents differ from paper is that they are not really the original, they have, at some point in their lifetime, been copied from one medium to another. Further back-up copies are made by the archive to de-risk the preservation process. In medieval times, scribes would make copies of documents and only a careful reading of copy and original could verify a faithful copy. In the digital world, we use methods originating in cryptography to automatically verify that not a single byte is out of place in a copied file. When a file is received by the archive, a cryptographic hash, or checksum, of the file is generated and stored in a database. Regular recomputation of file hashes are made and compared with the database to proactively identify corrupted files. If a corruption is found, the file can be replaced with a back-up copy.

'Where digital documents differ from paper is that they are not really the original'

When we come to present the user of the archive with a digital file, there are two options available. They could download the file in its original format, or download the file in an alternative format. One reason for the second option being presented is because as time goes by software is replaced with modern versions and the file formats change with them. For example, WordStar was a very popular word processor in the 1980s but there is no longer a version that runs on modern computers, although emulators, created by enthusiasts, are available. A WordStar file may be opened in a modern version of Microsoft Word after first installing a conversion add-in. This keeps the format alive and usable for now, but can we guarantee that these files will render on a standard computer in 20 years or 50 years? Even then, a modern word processor is not necessarily faithfully rendering the original. In the interests of long-term preservation, and for the convenience of users, the archive may create copies of these WordStar files and then convert them to an open format which is more likely to be still readable decades from now. Similar actions may also be taken with formats such as high definition video, converting them to a compressed format in order to reduce the time to download, again for user convenience.

'a converted file is different from the original and so the system of checksums breaks down'

Changing formats in this way introduces a problem: cryptographically, a converted file is different from the original and so the system of checksums breaks down. Software providers have used checksums for years to allow customers to verify that they are downloading a genuine copy, and the archive can use them to verify born-digital files in the same way. By

changing the format, however, we are offering a cryptographically different file to the one which was originally deposited. How can we assure that nothing nefarious has happened to the file during the conversion process? One answer to this question is to provide both the original and converted version so that they can be compared with each other. Surely, this defeats the purpose of the converted copy which, as has been stated, may be in a more convenient format. In addition, if there is no software available to render the original, this becomes an impossible task. A possible solution to this quandary is to record the provenance of the converted file, and provide a conversion program (which does not rely on rendering the file) to the user. We can record the hashes of the original, the newly formatted file and the converter, making public a transparent and reproducible process that can be independently verified.

This is one solution, but is there a way of demonstrating that two files in different formats are still the same without comparing them side by side? Our project is attempting to tackle this challenge with video files, where it would be a painstaking task to compare two long videos frame by frame, let alone thousands. Researchers at the University of Surrey Computer Vision Centre are using deep neural networks to generate content-aware hashes of video. The technology is still experimental but it is able to create a hash which is invariant to changes in format, but changes more drastically if the file is manipulated in other ways, for example by removing frames. This gives us a method for automatically comparing two video files in different formats and providing assurance that nothing happened to the content during the conversion process. An obvious use case is the detection of malicious or fraudulent actions against archived objects, but a more mundane yet far more likely use case is the detection of corruption during the conversion process; again, questions of scale and resources make it impossible to perform this task manually.

By publishing the reproducible audit trail and a content aware hash of the files this should be enough to demonstrate that the content of the files has not been affected by the custodianship of the archive. But is that right? By publishing this assurance on its own website, is the archive really guaranteeing that it has not amended the original content of the file? The straightforward answer is that the archive is a trusted organization which has been looking after the nation's records of government for hundreds of years. However, we live in turbulent times of fake news and conspiracy theories, a 'post truth' era where public trust in officialdom is at a low ebb.[2] It is also almost impossible to trust digital content. Manipulating a paper document or printed photo without trace requires the work of someone with great expertise; changing a digital photo needs just five minutes with Photoshop.[3] Thousands of text documents can be edited from the command line in minutes.

'changing a digital photo needs just five minutes with Photoshop'

A greater danger has arisen in recent years in the form of AI content-generating technology. Researchers have already developed technology to faithfully reproduce a person's voice and read out any transcript; deep fakes which have been used to insert the faces of celebrities into pornographic videos with no expertise needed;[4] and, most disturbingly, the OpenAI organization demonstrated a new AI model, called GPT2, which writes text that perfectly mimics any person's writing style, given a short sample. This last technology is considered so dangerous that the developers refuse to make it available until they fully understand its implications.[5]

The blockchain solution

This is where we believe DLT can be part of the fight against this onslaught of computer-generated fake content. It was originally developed to enable financial transactions between parties with no basis for trust between them. When describing DLT, it is easy to get bogged down in the technology, and particularly cryptography, but it is clearer to think of it as a database with two important features: it is an immutable, append only, ledger; and it is distributed, which means every party in the network has a copy. This means that when the

aforementioned provenance and content hash information is added to the blockchain, it cannot be removed or edited. It provides a technological underscoring of trust, unnecessary until trust is eroded, by which point it is too late to implement retrospectively. The archive is above all a store of evidence, but if the content is not trusted even while the institution is, then its purpose is uncertain. What we store in the DLT is the evidence of the evidence; it comprises cryptographically or otherwise generated digital fingerprints of the nation's history.

For the archive, we want to be able to demonstrate that we have not amended the records we have been entrusted to look after in the most transparent, immutable way possible. In the wake of OpenAI's GPT2 demonstration and some of the other similar technologies where it is possible to put words into the mouths of any public figure, including journalists and politicians, DLT may offer a layer of protection from fakery. Articles, interviews, speeches and suchlike could be hashed, digitally signed, and written in real time to a blockchain providing a time-stamped, unchangeable record of evidence of what was said. Computer-generated content produced by malicious actors will lack this chain of evidence back to the purported source of their claims.

'What we store in the DLT is the evidence of the evidence … cryptographically or otherwise generated digital fingerprints of the nation's history'

'DLT may offer a layer of protection from fakery'

How it works

The principle behind blockchain is that it would take someone with command of at least 51% of the computing power to maliciously rewrite the main chain of records. This is because each block in the chain contains a cryptographic link to the previous one. Therefore, to rewrite a transaction in a block would require all subsequent blocks to be recomputed. Since the chain is always moving forward, only an entity with more computational resource than the rest of the participants could effectively catch up. Obviously, the more participants involved, the harder it is for anyone to take control in this way. This is especially the case with the ARCHANGEL network which is using a permissioned blockchain, meaning that all participants are invited to join as trusted organizations in their own right. This is also means that the (pseudo) anonymity of a typical blockchain network is not there, and therefore suspicious behaviour is identifiable to particular participants.

Blockchain questions

A project workshop held in March 2018 explored a variety of questions around the technology and its use as envisaged by the project. Attendees from the fields of records management, archives, law, academia, and digital preservation system vendors gave us an excellent insight into the different concerns and potential uses of blockchain. Three main areas emerged:

- collaboration
- longevity of the technology
- viability of technology within archives.

'the ARCHANGEL blockchain relies on the need for a number of participants from as many different institutions as possible'

Collaboration

Collaboration is essential to the value offered by blockchain technology.
The success of the ARCHANGEL blockchain relies on the need for a number of participants from as many different institutions as possible. Without a minimum number of participants, we think at least seven, the trust that the technology engenders is in danger of being lost. In a live, future environment, we would hope to involve participants beyond the archive sector, such as news organizations and other transparency groups[6] who as well as providing external oversight also have a stake in the assurance of transactions. Only by maintaining

a minimum number of participants throughout its life can a blockchain survive to provide evidence for users. This distributed approach to assuring trust is another example of the long running practice of memory institutions being reliant on each other. Archival capability has long been distributed and shared in terms of know-how, the development and maintenance of tools, and of key archival resources (for example PRONOM[7] and LOCKSS[8]). And for those organizations that hold web archives, the collections themselves overlap and content is shared as archives supply each other with content to fill gaps in their collections.

Because the challenge is great and, as institutions, archives are quite small, this approach is the key to winning the technology arms race between archives and those parties who use the tools to falsify our digital inheritance. If this approach to distributed services works, it is exciting to think about what else could be distributed in future.

Longevity of the technology

DLT is a relatively new technology and is currently riding high on the Gartner hype cycle.[9] The technology is advancing but it is far from mature and it is likely to experience growing pains over the coming years. The first Ethereum 'hard fork' of 2016 is a good example of the challenges of technological change in a distributed system.[10] Ethereum is one of a number of platforms using DLT and was chosen for the project because it is one of the most easily accessible platforms for executing smart contracts – essentially pieces of executable computer code that can automatically generate an activity when certain conditions are met – such as writing a block to the blockchain when it is validated. This hard fork resulted in two versions of the Ethereum platform, and the decision to split this way caused much disagreement in the Ethereum community. For The National Archives, longevity is critical. This means the technology needs to be around for decades, not just years, to be at all useful. Furthermore, any technology that is based on cryptography should not be considered safe over a long period of time, with cryptographic systems typically lasting around 20 years before their security is broken. Quantum computing is also on the horizon, potentially rendering some current encryption algorithms obsolete. Moreover, it is not just the longevity of the technology that is a concern; it is the distributed network itself. As already discussed, a DLT system relies on having many participants. As the technology becomes more popular, there will be many competing vendors and so participants may wish to migrate to a different platform. There is an initial barrier since there is no benefit to migrating alone, but it is feasible a network could split if a large enough group of participants prefer a different platform. Fortunately, there are efforts under way to create standards for interoperability between platforms, which may mitigate against this problem.[11]

'any technology … based on cryptography should not be considered safe over a long period of time'

Viability of technology within archives

Digital preservation approaches are developing and vary across archives, as each has different priorities and different levels of automation. In order for the ARCHANGEL blockchain to be viable, it is essential that the software can be integrated into those existing processes and automated workflows. At every point where a change could be made to the record, the checksum and metadata should be sent to the blockchain. Only if the registering of checksums and other metadata on the blockchain is performed at predetermined points in the process will it provide the provenance that users and archivists can trust.

The ARCHANGEL blockchain application

The ARCHANGEL project has developed a prototype desktop application which a number of archival institutions have volunteered to test. This will enable the project team to understand how it could fit into varying archival processes of institutions around the world who work with different levels of automation. It will also bring to light issues around implementing the software in different technological environments, the feasibility of integrating it with other processes, enabling the benefits of the technology, not previously used in archives, to be better understood.

To investigate these questions and understand how blockchain can be used to establish the authenticity of records, the ARCHANGEL project has constructed a private Ethereum test bed. This will enable participants to take part in a permissioned blockchain. They will use the network to upload metadata about their records, validate the metadata uploaded by other participating organizations and be able to search and view the results on the blockchain.

When designing the application, the project team had to decide what data would be written to the ARCHANGEL blockchain. As data written to the blockchain cannot be deleted or amended, the sensitive nature of some archival records meant that the records themselves, including in some cases their file-names, could not be made available. This means that, along with the checksums, only the file's unique identifier, its size in bytes and the file format in the form of the PUID (PRONOM Unique Identifier) are being written to the blockchain by the application.

'there is a real risk that our digital heritage will be lost'

Conclusion

Digital presents many challenges to traditional archival practice. In the area of preservation, the approach of ensuring the original paper record remains virtually untouched while in the archive's custody almost guarantees its loss in the case of digital. Without refreshing the storage media, checking and rechecking for change at byte level and creating copies in different formats, there is a real risk that our digital heritage will be lost. Moreover, while trust in archival institutions remains high, the ease with which digital files can be altered means that archives cannot rely solely on their reputation to guarantee the authenticity of the records in their custody.

It could be argued that blockchain may be a solution for a problem that does not yet exist but there are already services available that provide storage for digital records which cannot be overwritten or deleted for a predetermined period of time.[12] Challenges to archives' trustworthiness will be made, and it will be too late to demonstrate their careful custodianship by implementing a solution years after the records were received. As a store of evidence, an important pillar of democracy, the archive needs to be able to prove the authenticity and integrity of the records they hold in a way that cannot be challenged.

'the archive needs to … prove the authenticity and integrity of the records … in a way that cannot be challenged'

There are more lessons to learn and the ARCHANGEL project will continue to explore the feasibility, challenges and opportunities of a blockchain network with the prototype Ethereum network and the participating institutions from around the world.

Acknowledgements
The ARCHANGEL Project would like to acknowledge the funding received from the ESPRC Grant Ref EP/P03151X/1.

Abbreviations and Acronyms
A list of the abbreviations and acronyms used in this and other *Insights* articles can be accessed here – click on the URL below and then select the 'full list of industry A&As' link: **http://www.uksg.org/publications#aa**

Competing Interests
The authors have declared no competing interests.

References

1. John D. Garrett et al., "Preserving digital information: Report of the task force on archiving of digital information," (PDF). *Commission on Preservation and Access and the Research Libraries Group* (1996), https://archive.org/details/PreservingDigitalInformationTaskForceReport1996 (accessed 15 May 2019).
2. Louella-Mae Eleftheriou-Smith, "Donald Trump says interview with *The Sun* criticising Theresa May is 'fake news' even though it was recorded," *i*, July 13, 2018, https://inews.co.uk/news/politics/donald-trump-the-sun-theresa-may-fake-news/ (accessed 15 May 2019).
3. John Swain, "Trump inauguration crowd photos were edited after he intervened," *The Guardian*, September 6, 2018, https://www.theguardian.com/world/2018/sep/06/donald-trump-inauguration-crowd-size-photos-edited (accessed 15 May 2019).
4. Stephen Maher, "Fake video is a big problem. In 2019, it gets worse," *MSN*, December 28, 2018, https://www.msn.com/en-ca/money/topstories/fake-video-is-a-big-problem-in-2019-it-gets-worse/ar-BBRwxZ4 (accessed 15 May 2019).
5. Alex Hearn, "New AI fake text generator may be too dangerous to release, say creators", *The Guardian*, February 14, 2019, https://www.theguardian.com/technology/2019/feb/14/elon-musk-backed-ai-writes-convincing-news-fiction (accessed 15 May 2019).
6. Transparency groups are organizations which campaign for the release of information to ensure that particular bodies, such as governments or businesses, are not corrupt, for example Transparency International, https://www.transparency.org.uk/ (accessed 20 May 2019).
7. "The technical registry, PRONOM," *The National Archives*, http://www.nationalarchives.gov.uk/PRONOM/Default.aspx (accessed 15 May 2019).
8. The LOCKSS Program, https://www.lockss.org/ (accessed 15 May 2019).
9. "Gartner Hype Cycle," Gartner, https://www.gartner.com/en/research/methodologies/gartner-hype-cycle (accessed 15 May 2019).
10. A hard fork in the context of distributed ledger technology is a phenomenon in which a change in the technology results in a divergence that creates two separate technologies that are incompatible with each other.
11. A working group on the Interoperability of blockchain and distributed ledger technology systems (ISO/TC 307/SG 7) has been set up, https://www.iso.org/committee/6266604.html (accessed 15 May 2019).
12. "Amazon S3 Object Lock Overview," *AWS*, https://docs.aws.amazon.com/AmazonS3/latest/dev/object-lock-overview.html (accessed 15 May 2019).

Corresponding author:
Alex Green
Digital Archivist
The National Archives, GB
E-mail: alex.green@nationalarchives.gov.uk

To cite this article:
Bell M, Green A, Sheridan J, Collomosse J, Cooper D, Bui T, Thereaux O and Higgins J, "Underscoring archival authenticity with blockchain technology", *Insights*, 2019, 32: 21, 1–7; DOI: https://doi.org/10.1629/uksg.470

Submitted on 05 April 2019 Accepted on 10 June 2019 Published on 26 June 2019

Published by UKSG in association with Ubiquity Press.