Number 110 July 2017

ERCIM NEWS

www.ercim.eu

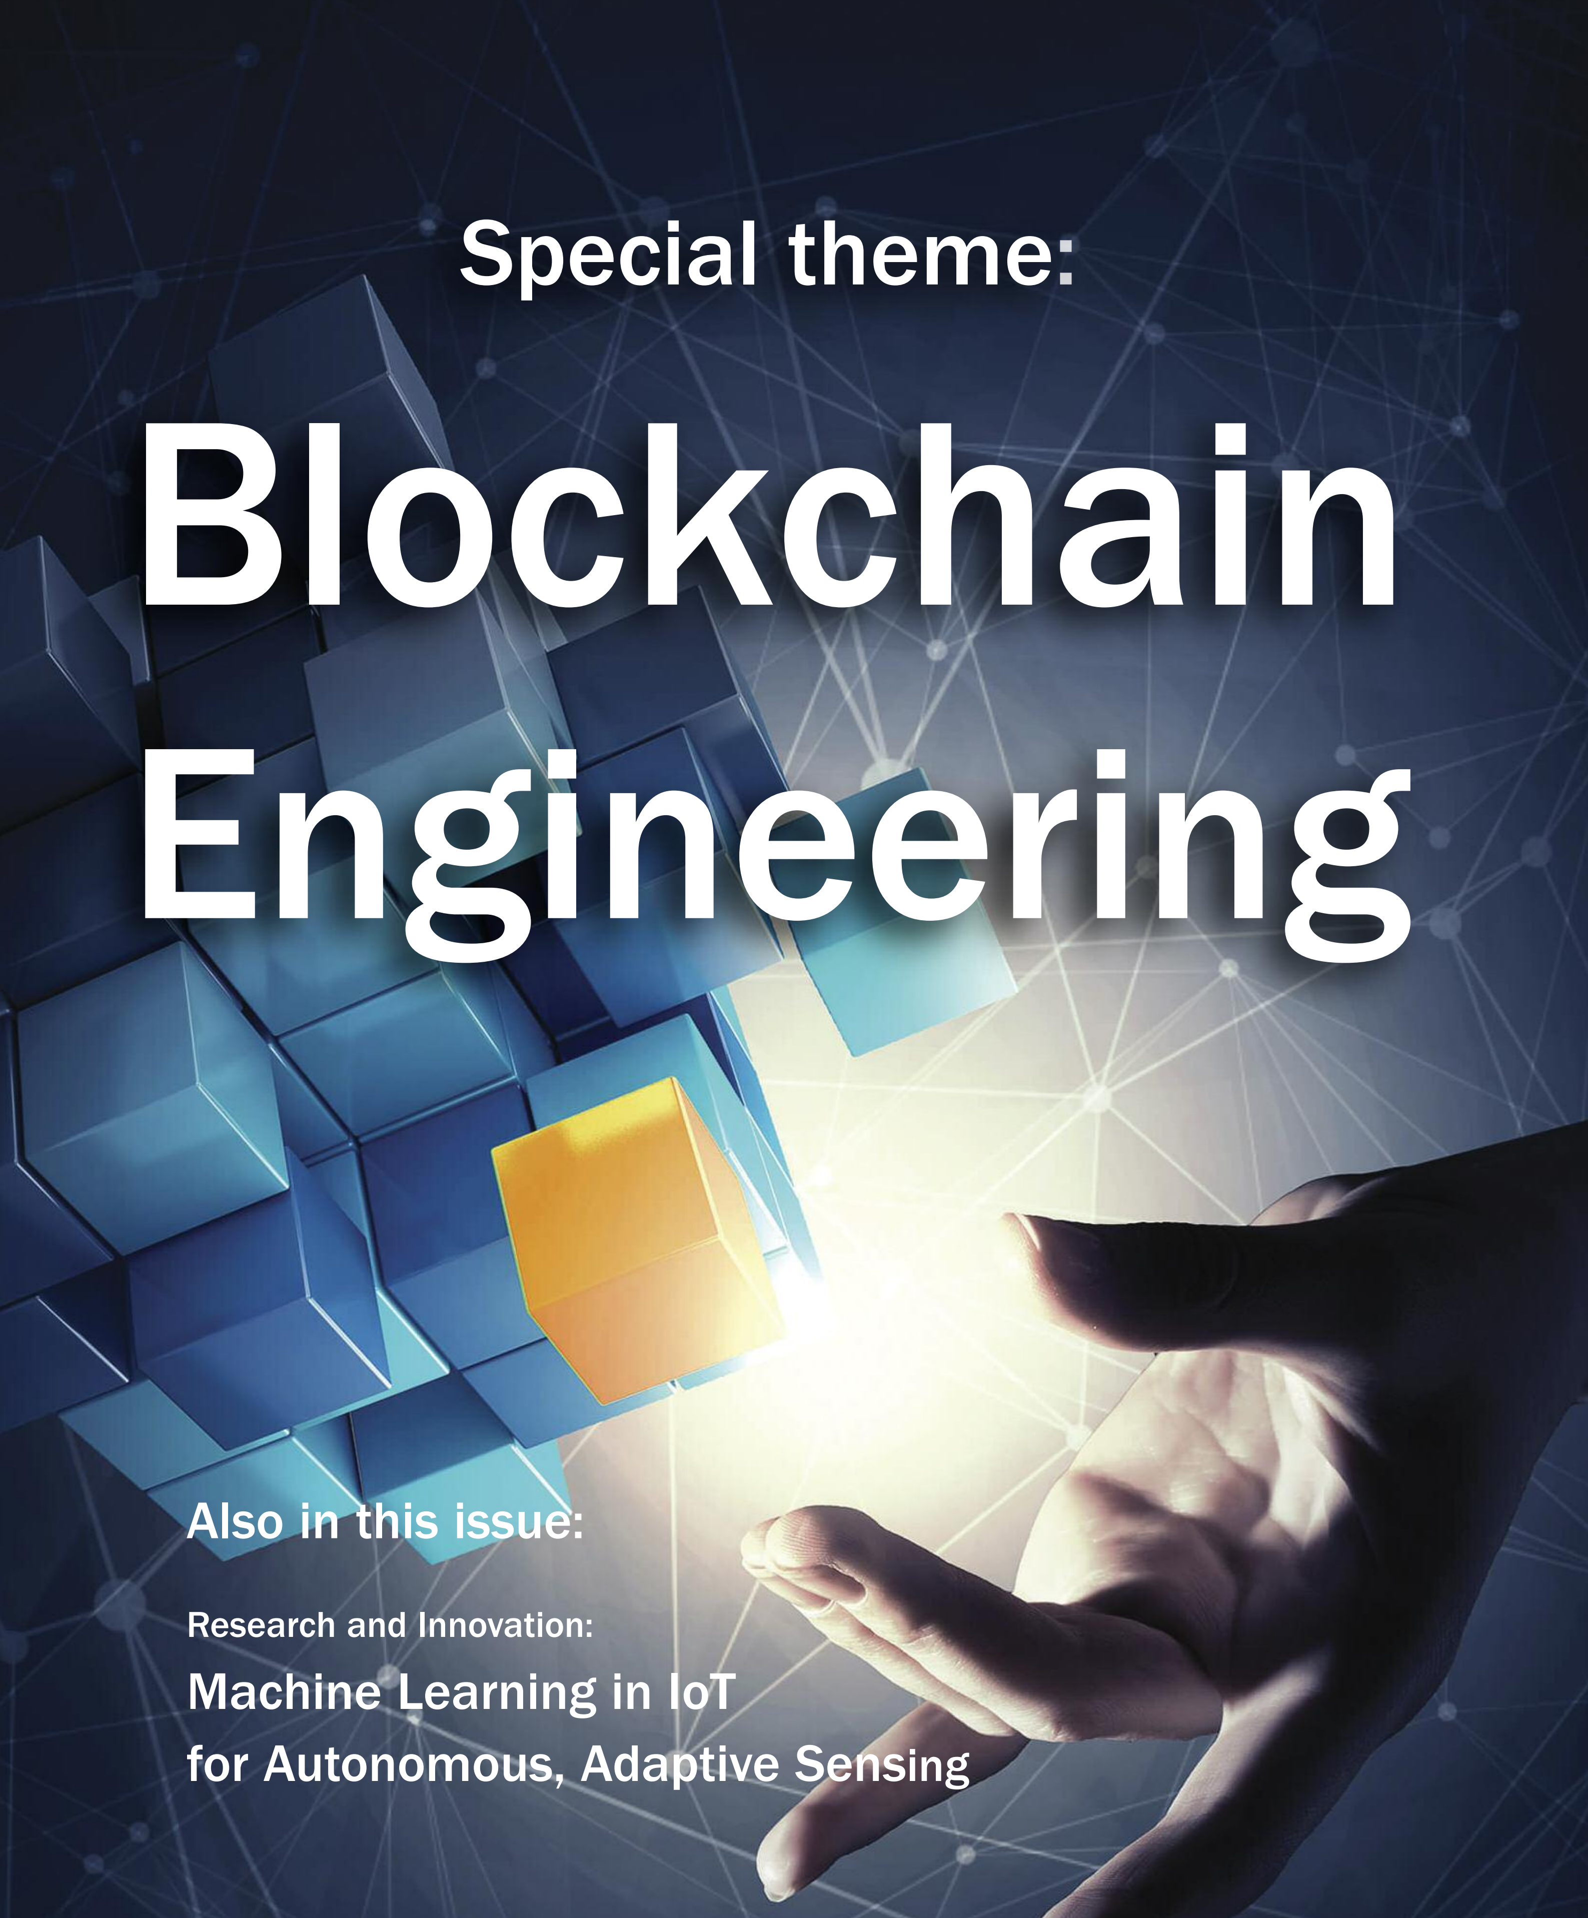

Special theme:

Blockchain Engineering

Also in this issue:

Research and Innovation:

Machine Learning in IoT for Autonomous, Adaptive Sensing

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 6,000 printed copies and is also available online.

ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
Tel: +33 4 9238 5010, E-mail: contact@ercim.eu
Director: Philipp Hoschka, ISSN 0926-4981

Contributions
Contributions should be submitted to the local editor of your country

Copyright notice
All authors, as identified in each article, retain copyright of their work. ERCIM News is licensed under a Creative Commons Attribution 4.0 International License (CC-BY).

Advertising
For current advertising rates and conditions, see http://ercim-news.ercim.eu/ or contact peter.kunz@ercim.eu

ERCIM News online edition
http://ercim-news.ercim.eu/

Next issue
October 2017, Special theme: Digital Humanities

Subscription
Subscribe to ERCIM News by sending an email to en-subscriptions@ercim.eu or by filling out the form at the ERCIM News website: http://ercim-news.ercim.eu/

Editorial Board:
Central editor:
Peter Kunz, ERCIM office (peter.kunz@ercim.eu)
Local Editors:
Austria: Erwin Schoitsch (erwin.schoitsch@ait.ac.at)
Belgium:Benoît Michel (benoit.michel@uclouvain.be)
Cyprus: Ioannis Krikidis (krikidis.ioannis@ucy.ac.cy)
France: Steve Kremer (steve.kremer@inria.fr)
Germany: Michael Krapp (michael.krapp@scai.fraunhofer.de)
Greece: Eleni Orphanoudakis (eleni@ics.forth.gr),
Artemios Voyiatzis (bogart@isi.gr)
Hungary: Andras Benczur (benczur@info.ilab.sztaki.hu)
Italy: Maurice ter Beek (maurice.terbeek@isti.cnr.it)
Luxembourg: Thomas Tamisier (thomas.tamisier@list.lu)
Norway: Poul Heegaard (poul.heegaard@item.ntnu.no)
Poland: Hung Son Nguyen (son@mimuw.edu.pl)
Portugal: José Borbinha, Technical University of Lisbon (jlb@ist.utl.pt)
Spain: Silvia Abrahão (sabrahao@dsic.upv.es)
Sweden: Kersti Hedman (kersti@sics.se)
Switzerland: Harry Rudin (hrudin@smile.ch)
The Netherlands: Annette Kik (Annette.Kik@cwi.nl)
W3C: Marie-Claire Forgue (mcf@w3.org)

Contents

JOINT ERCIM ACTIONS

SPECIAL THEME

The special theme section "Blockchain Engineering" has been coordinated by Elli Andoulaki (IBM Research – Zurich), Matthias Jarke (RWTH Aachen University & Fraunhofer FIT) and Jean-Jacques Quisquater (Université catholique de Louvain, Belgium, and research affiliate at MIT)

Invited articles

Cryptocurrencies

Contracts and workflows

Public sector applications

Security and privacy

Blockchain labs

RESEARCH AND INNOVATION

This section features news about research activities and innovative developments from European research institutes

EVENTS, IN BRIEF

New Project

Call for Proposals

Call for Participation

Report

In Brief

ERCIM Membership

After having successfully grown to become one of the most recognized ICT Societies in Europe, ERCIM has opened membership to multiple member institutes per country. By joining ERCIM, your research institution or university can directly participate in ERCIM's activities and contribute to the ERCIM members' common objectives playing a leading role in Information and Communication Technology in Europe:

- Building a Europe-wide, open network of centres of excellence in ICT and Applied Mathematics;
- Excelling in research and acting as a bridge for ICT applications;
- Being internationally recognised both as a major representative organisation in its field and as a portal giving access to all relevant ICT research groups in Europe;
- Liaising with other international organisations in its field;
- Promoting cooperation in research, technology transfer, innovation and training.

About ERCIM

ERCIM – the European Research Consortium for Informatics and Mathematics – aims to foster collaborative work within the European research community and to increase cooperation with European industry. Founded in 1989, ERCIM currently includes 15 leading research establishments from 14 European countries. ERCIM is able to undertake consultancy, development and educational projects on any subject related to its field of activity.

ERCIM members are centres of excellence across Europe. ERCIM is internationally recognized as a major representative organization in its field. ERCIM provides access to all major Information Communication Technology research groups in Europe and has established an extensive program in the fields of science, strategy, human capital and outreach. ERCIM publishes ERCIM News, a quarterly high quality magazine and delivers annually the Cor Baayen Award to outstanding young researchers in computer science or applied mathematics. ERCIM also hosts the European branch of the World Wide Web Consortium (W3C).

> **"Through a long history of successful research collaborations in projects and working groups and a highly-selective mobility programme, ERCIM has managed to become the premier network of ICT research institutions in Europe. ERCIM has a consistent presence in EU funded research programmes conducting and promoting high-end research with European and global impact. It has a strong position in advising at the research policy level and contributes significantly to the shaping of EC framework programmes. ERCIM provides a unique pool of research resources within Europe fostering both the career development of young researchers and the synergies among established groups. Membership is a privilege."**
>
> ***Dimitris Plexousakis, ICS-FORTH, ERCIM AISBL Board***

Benefits of Membership

As members of ERCIM AISBL, institutions benefit from:

- International recognition as a leading centre for ICT R&D, as member of the ERCIM European-wide network of centres of excellence;
- More influence on European and national government R&D strategy in ICT. ERCIM members team up to speak with a common voice and produce strategic reports to shape the European research agenda;
- Privileged access to standardisation bodies, such as the W3C which is hosted by ERCIM, and to other bodies with which ERCIM has also established strategic cooperation. These include ETSI, the European Mathematical Society and Informatics Europe;
- Invitations to join projects of strategic importance;
- Establishing personal contacts with executives of leading European research institutes during the bi-annual ERCIM meetings;
- Invitations to join committees and boards developing ICT strategy nationally and internationally;
- Excellent networking possibilities with more than 10,000 research colleagues across Europe. ERCIM's mobility activities, such as the fellowship programme, leverage scientific cooperation and excellence;
- Professional development of staff including international recognition;
- Publicity through the ERCIM website and ERCIM News, the widely read quarterly magazine.

How to Become a Member

- Prospective members must be outstanding research institutions (including universities) within their country;
- Applicants should address a request to the ERCIM Office. The application should inlcude:
    - Name and address of the institution;
    - Short description of the institution's activities;
    - Staff (full time equivalent) relevant to ERCIM's fields of activity;
    - Number of European projects in which the institution is currently involved;
    - Name of the representative and a deputy.
- Membership applications will be reviewed by an internal board and may include an on-site visit;
- The decision on admission of new members is made by the General Assembly of the Association, in accordance with the procedure defined in the Bylaws (http://kwz.me/U7), and notified in writing by the Secretary to the applicant;
- Admission becomes effective upon payment of the appropriate membership fee in each year of membership;
- Membership is renewable as long as the criteria for excellence in research and an active participation in the ERCIM community, cooperating for excellence, are met.

Please contact the ERCIM Office: contact@ercim.eu

First ERCIM Workshop on Blockchain Technology

As part of the 2017 ERCIM spring meetings in Paris, ERCIM held a half-day workshop on blockchain technology on May 23 2017. Co-chaired by Georges Gonthier (Inria) and Wolfgang Prinz (Fraunhofer FIT), the workshop provided a high-level overview of blockchain technology and its opportunities for computer science research to the senior-level workshop attendees. The attendees included executives of ERCIM member institutes as well as a number of researchers.

Wolfgang Prinz (Vice-Chair of Fraunhofer FIT Institute) started out the morning by giving a comprehensive introduction to blockchain technology, its application areas, and related computer science research questions. In particular, he outlined the various areas of computer science research that blockchain technology is touching and using, which include:

- P2P networks
- Distributed systems (in particular scalability)
- Cryptography (with a focus on crypto-agility)
- Consensus-building and validation
- Software lifecycle of smart contracts.

The presentation also provided a classification of the design space which different blockchain technologies are using (unpermissioned versus permissioned, logic-oriented versus transaction oriented). Finally, Wolfgang outlined a number of potential areas of collaboration between ERCIM members, including the creation of an ERCIM blockchain infrastructure.

In the second talk of the day, Georges Gonthier (Inria SPECFUN Unit) talked about the application of formal methods to smart contracts. He outlined the pitfalls of languages currently used for programming smart contracts and their consequences, including the bug in the Ethereum blockchain network that led to the highly visible loss of 53 million dollars (which were later recovered). He argued for the use of formal proof and analysis of smart contracts to prevent this type of issue in the future. The second part of the presentation focussed on new challenges and ideas in the area of name services.

In the final talk of the morning, Arnaud Le Hors (IBM, Member of the Hyperledger Technical Steering Committee) presented the Hyperledger open source project and its quickly growing success in terms of participants and applications. In particular, Arnaud reported that Hyperledger is the fastest growing project in the history of the Linux Foundation, with 300% growth in the first year. He further described the workings of the project, including working groups that are open and free for anyone to participate in, as well as regular hackathons, hackfests and meetups. Then, Arnaud provided a detailed description of the Hyperledger 1.0 "fabric" architecture, covering the ordering service, single and multi channel networks, chaincode and endorsement policies. Arnaud concluded his talk explaining how to get started using Hyperledger, and how to get involved in the community.

Please contact:
Philipp Hoschka, ERCIM Manager
philipp.hoschka@ercim.eu

ERCIM "Alain Bensoussan" Fellowship Programme

ERCIM offers fellowships for PhD holders from all over the world. Topics cover most disciplines in Computer Science, Information Technology, and Applied Mathematics. Fellowships are of 12 months duration, spent in one ERCIM member institute. Fellowships are proposed according to the needs of the member institutes and the available funding.

Application deadlines for the next round: 30 April and 30 September 2017

More information: http://fellowship.ercim.eu/

HORIZON 2020 Project Management

A European project can be a richly rewarding tool for pushing your research or innovation activities to the state-of-the-art and beyond. Through ERCIM, our member institutes have participated in more than 80 projects funded by the European Commission in the ICT domain, by carrying out joint research activities while the ERCIM Office successfully manages the complexity of the project administration, finances and outreach.

The ERCIM Office has recognized expertise in a full range of services, including identification of funding opportunities, recruitment of project partners, proposal writing and project negotiation, contractual and consortium management, communications and systems support, organization of attractive events, from team meetings to large-scale workshops and conferences, support for the dissemination of results.

How does it work in practice?
Contact the ERCIM Office to present your project idea and a panel of experts will review your idea and provide recommendations. If the ERCIM Office expresses its interest to participate, it will assist the project consortium as described above, either as project coordinator or project partner.

Please contact:
Philippe Rohou, ERCIM Project Group Manager
philippe.rohou@ercim.eu

Introduction to the Special Theme

Blockchain Engineering

by Elli Andoulaki (IBM Research – Zurich), Matthias Jarke (RWTH Aachen University & Fraunhofer FIT) and Jean-Jacques Quisquater (Université catholique de Louvain, Belgium, and research affiliate at MIT)

In the last decade, the world of data management has been revolutionised by the influence of universally available distributed and mobile information technology. The jump from desktop and laptop to the smartphone has been a major driver, and the current explosive growth of the internet of things is another. Big data analytics is no longer only a buzzword in computer science, but transcends all levels of business, politics, and society.

In contrast to the explosion of the query processing and data mining side of this development, its equally important impact on transaction management has received much less attention. The problems of misleading information inputs (fake news, chatbots), broken or fraudulent transactions are discussed in public, but scalable solutions around these distributed transaction challenges, most prominently the blockchain technology, has only recent begun to capture more attention, fostered by speculation about crypto-currencies such as Bitcoin.

Conceptually, blockchains can be understood as distributed ledgers, aiming like traditional ledgers at transparent and falsification-proof documentation, while assuming a model where distribution of trust is required. That is, in blockchain systems, operational trust is distributed to two or more mutually distrusting entities. Technically, scalability, anonymity, security and durability are ensured by distributed storage combined with suitable cryptographic primitives and protocols, but many problems remain to be investigated.

In the last couple of years, European industries (e.g., the B3i Blockchain insurance industry initiative) as well as the European Union (e.g., EU blockchain observatory, Blockchain for Industrial Transformation, blockchain architecture call) have started or announced a significant number engineering and policy initiatives. In this special issue of the ERCIM News, we provide an overview of some of the active European research in the field of blockchain engineering.

In the first paper, Jean-Jacques Quisquater – a pioneer of blockchain research since the late 1990s – provides an overview of the concepts, history, and current challenges. In innovative businesses and research, the engineering of blockchain-based solutions is subject to quite a number of commercial and open source initiatives. As a current major open source example, Andoulaki et al. (IBM Research, Zurich)) provide a glimpse on the just released Fabric for permissioned blockchains within the international Hyperledger initiative. This is followed by three sections on different blockchain application engineering domains (finance, public sector, contract and workflow management) and a special focus on security and privacy issues in the context of blockchains. The special theme ends with a couple of blockchain labs and strategic initiatives.

Among the blockchain applications in finance, bitcoin is surely the best-known. Complementing an overview of Bitcoin applications (Judmayer, Zamyatin, Nicholas, SBA Research Vienna), a team from the INRIA and partners (Augot, Chabanne, and George) is specifically studying the question of Identity Management on the Bitcoin blockchain, while a Norwegian-Australian collaboration (Carr, Boyd, Boyen (NTNU Trondheim) and Haines (QUT, Brisbane)) aims to release some restrictions of the current technologies. To strengthen the theoretical foundations, a different space-oriented proof technique for crypto-currencies, called

SpaceMint, is presented in a paper from Inria (Fuchsbauer), whereas the integration of cryptocurrencies such as bitcoin in efficient real-time payment processes is one of the practical challenges (Bocek, Rafati, and Mori, University of Zurich).

The seemingly paradoxical combination of transparency and privacy offered by blockchains make them suitable for many applications beyond crypto-currencies. Generalising from crypto-currencies to general asset exchange, a team around FORTH-ICS (Askoxylakis, Alexandris and Demetriou) discuss this aspect in a circular economy, and an Italian team around CNR-IIT looks at healthcare applications (Lo Duca, Bacciu, and Marchetti) whereas Christian Welzel (Fraunhofer FOKUS, Berlin) weighs the threats and opportunities of blockchains from the viewpoint of the public sector in general. This section ends with a discussion (Nowostawski, NTNU Trondheim) how blockchains can be used to make institutions more autonomous.

The concepts of smart contracts and associated workflows is ubiquitous in almost all blockchain application domains. Linking back to the financial application domain, Fridgen, Urbach and Sablowsky (Fraunhofer FIT Bayreuth) present the blockchain-based workflow management system at a German bank. Three other papers investigate the important proof of work (Biryukov, University of Luxembourg), consistency enforcement (Osterland and Rose, Fraunhofer FIT Sankt Augustin), and smart contract security (Stifter, Judmayer, and Weippl, SBA Research Vienna). Another important challenge is the merger of multiple workflows or blockchains (Mellissen, Storro B.V.).

Despite full transparency of the transactions, blockchains also need to protect the privacy of their users and of the person-related data within them. Indeed, blockchains can even support the implementation of the new European Data Privacy Regulation with respect to transparency of person data usage (Roth), and a flexible transparency approach can be employed to control the degree of user privacy as well (Christofi and Gouget, Trusted Labs Versailles). At the corporate and the individual level alike, data sovereignty has recently become an important goal in European policy making and system engineering; a suitable Identity Framework can be combined with blockchains to get closer to this challenging goal (Joosten). Also at the corporate level, the paper by Di Francesco Maesa, Ricci, and Mori (CNR-IIT and Pisa University) demonstrates the usefulness of blockchain technologies for data access control in large systems.

A number of blockchain research labs, national initiatives, have sprung up recently in several European countries. The special issue ends with a description of two examples from the Netherlands (CWI Amsterdam) and Germany (Fraunhofer).

This special theme shows that there are still many challenges to overcome from the perspective of engineering as well as business models and public policy regulations. Nevertheless, a growing number of applications already indicates the enormous potential of blockchain technologies.

Please contact:
Elli Andoulaki
IBM Research – Zurich, Switzerland
LLI@zurich.ibm.com

Matthias Jarke
Information Systems Group, RWTH Aachen University & Fraunhofer FIT, Germany
jarke@dbis.rwth-aachen.de

Jean-Jacques Quisquater
Crypto Group, Université catholique de Louvain, Belgium, and research affiliate at MIT
jjq@uclouvain.be

Blockchains for Everybody: Individuals, Companies, States and Democracy

by Jean-Jacques Quisquater

Trust, transparency and traceability (or nontraceability) are important in online transactions, which may involve banks, notaries, public administrations, trusted-third-parties, witnesses and others. Even long before the internet, people in ancient civilizations used tools to create a permanent trace, such as a public (or private) ledger: Assyrian people used tablets and Incas used khipus, for instance.

In the 19th century, people dealt with the problems of synchronisation of clocks and being able to know the correct time in different locations, which was necessary to schedule trains. Telegraphy largely solved these issues - but only after lengthy negotiations (in France, it was not until 1891 that the time was unified). Synchronisation of clocks in practical situations was a research subject for Albert Einstein and others, with the eventual winner being the theory of special relativity, which is applied today in GPS.

Timestamping was an important subject for the authentication of actions. But it often needed trust in a particular authority, such as a notary, which left open the possibility of errors or cheating. Coordinated timestamping was also required for patents, music, contracts, auctions and other purposes.

In the late 1980s I was working for Philips Research in Belgium. At that time I was the head of the crypto group, which was making great inroads into the security of smart cards. In 1989, my boss asked the team to imagine new applications that might be enabled by the transition from binary flow (Shannon) to multimedia streams (sound, images, videos, etc). The idea was to translate every action (very often analogue) into the digital world. So we began considering how cryptography might be used for watermarking, timestamping and geolocalisation. We then communicated with Belgian notaries and they were very interested in our ideas. Alas, it was too early because the research into cryptographic hash functions was not yet mature enough, and the standardisation process (ISO, IETF) was then being lobbied for by banking sector, which did not understand the challenges (can you imagine today that people did not approve proposals taken into the anniversary's paradox because it was paranoid …). Practical functions were finally proposed by Ron Rivest (MIT): the MD4 and MD5 cryptographic hash functions in 1990 and 1991 respectively. Curiously, with the exception of Raph Merkle, nobody at that time was really interested in working with these functions. However, hash functions were to become the future of digital signature, as well as blockchains and bitcoins.

The first public secure timestamping scheme, based on cryptography, was set by Stuart Haber and Scott Stornetta (1990) [1] and, even at this time, their proposals were very mature: the first one proposed chaining using cryptographic hash functions, the second one distributed the chain with a random positioning of the actors, that is, blockchain of today! They also added blocks using an idea of Raph Merkle's (tree): then the blockchain as we know it today was ready – except for the mining and the solutions for possible forks. Mining was invented several times including the "Chinese Lotto" (1987-1991) [2]. A company, "surety.com", acted as a trusted-third-party for a chain with only one trusted point, and a journal (NYT) as the public ledger, which didn't require the use of internet.

A second early use of cryptographic chaining in the context of secure timestamping with broadcast was described for voting protocols by Josh Benaloh and Michael de Mare (1991) taking into account Haber-Stornetta. It is ironic that people are trying to solve voting problems using bitcoin, for instance, including the internal blockchain, when direct solutions have existed for a long time [3].

In 1996 an important timestamping project was initiated in Belgium: TIMESEC [4]. Its goals included: to improve the network time protocol for internet; to push trusted timestamping using chains; to integrate blocks as we know today, and redundant hash functions; to use several servers in a distributed and decentralized way; to examine the possible uses of cryptographic accumulators. This work took us one step closer to blockchains. A complete working demo was put on online for two years. But it was also too early for a widespread adoption.

In 2001 an important report for the Bank of Japan was written by Masashi Une under the direction of Professor Matsumoto [5]. A comparison of the seven systems of digital timestamping was described and some classification was done by including the solutions by Haber-Stornetta and TIMESEC . The challenge of a really distributed timestamping was clearly set and the solution ended up being the one by Satoshi Nakamoto inside bitcoin [6]! In fact, the introduction, together with other experiments of peer-to-peer networks on internet provided the missing link for the success of timestamping.

New ideas are continually emerging: smart contract is a promising one, with complex internal verifications in order to avoid problems (it is possible to write a "nearly" undetectable virus in powerful Turing languages like Solidity [L1]: see also openzeppelin [L2]). Current challenges are scalability, time to register (latency is too big), how to put together several blockchains (I don't want to have hundreds of blockchains on my smartphone in the future), how to renew a blockchain if a systematic error is found, how to handle the right of forgotten (oblivion). And what about the possible power of quantum computers against the cryptographic primitives (not a complete science-fiction because NIST and NSA are thinking of soon replacing the primitives in use for bitcoin)? What are the relationships – if any - of blockchains to states and governments? How can we handle conflicts, errors, cheated contracts (a new area for lawyers?). When is consensus enough?

There are enough questions and problems to occupy many scientists and fuel numerous new projects, and I'm sure a future issue of ERCIM News is already on the cards to keep us up to date with their results. Industry-proved applications are on the way, which is very good, but we need to be very careful not to fix everything too early (standardisation): we are still at the stage of experiments, not of fully ready products. Is Algorand from Silvio Micali (MIT) [7], the next step?

Links:
[L1] http://www.ethereum.org
[L2] https://openzeppelin.org/

References:

[1] S. Haber, W. S. Stornetta: "How to time-stamp a digital document", Journal of Cryptology, January 1991, Vol. 3, Issue 2, pp 99–111 (first presented at CRYPTO '90). (see also patent US 5136647 A).

[2] J-J. Quisquater, Y. Desmedt: "The Chinese Lotto As An Exhaustive Code-breaking Machine", Computer, IEEE, Vol. 24, no. 11, p. 14-22 (1991). See also IETF RFC 3607.

[3] J. Benaloh, M. de Mare: "Efficient Broadcast Time-Stamping", TR from Clarkson University, 1991/1992.

[4] J-J. Quisquater, H. Massias, B. Preneel, B. Van Rompay: "TIMESEC final report", 1999, https://kwz.me/Xb (Cited in [5]).

[5] M. Une: "The Security Evaluation of Time Stamping Schemes: The Present Situation and Studies", Discussion Paper No. 2001-E-18, Institute for Monetary and Economic Studies, Bank of Japan, Tokyo.

[6] S. Nakamoto: "Bitcoin: A Peer-to-Peer Electronic Cash System, https://bitcoin.org/bitcoin.pdf

[7] J. Chen and S. Micali: "Algorand, The efficient and democratic ledger", eprint arXiv:1607.01341 (23 May 2017).

Please contact:
Jean-Jacques Quisquater
Crypto Group, Université catholique de Louvain, Belgium, and research affiliate at MIT
jjq@uclouvain.be

Permissioned Blockchains and Hyperledger Fabric

by Elli Androulaki, Christian Cachin, Angelo De Caro, Alessandro Sorniotti and Marko Vukolic (IBM Research, Zurich)

Blockchains can be defined as immutable decentralised ledgers for recording transactions that - depending on the system - are to various degrees resilient to malicious behaviour. Blockchain peers maintain copies of the ledger that consists of groups of transactions (blocks) linked together into a hash-chain. This effectively establishes total order among blocks and, consequently across transactions. Transactions have in recent years evolved to allow the execution of arbitrary logic, also known as smart contracts. In principle, a smart-contract is an application that operates on top of blockchain, which uses the underlying ordering of transactions (i.e., consensus) to maintain consistency of smart contract execution results across peers, now also referred to as execution replicas.

Blockchain networks, with the prominent example of Ethereum [L1], are typically public and open, i.e., anybody can participate without having a specific identity.

Permissioned blockchains have evolved as an alternative to open blockchains to address the need for running blockchain technology among a set of known and identifiable participants that are required to be explicitly admitted to the blockchain network. The concept behind permissioned blockchains is particularly interesting in business applications of blockchain technology and distributed ledgers, in which the participants require some means of identifying each other while not necessarily fully trusting each other.

In the world of business, permissioned blockchain systems often come across critical requirements (from a practical and regulatory perspective) for transactional security and privacy of business logic that is put on a shared ledger. In addition, commonly enterprise-purposed permissioned ledgers need to meet certain performance and scalability standards and/or comply with different cryptographic standards and practices, ultimately calling for modularity of crypto components.

Fabric [L2] is an open source project under the umbrella of Hyperledger [L3], a consortium hosted by Linux Foundation [L4] aiming to offer an enterprise-level permissioned blockchain platform. Fabric deals with all the aforementioned challenges, while offering support for execution of distributed applications (i.e., smart contracts or chaincodes in Fabric parlance) in general-purpose programming languages.

But, let's take a closer look to Fabric.

Technically, Fabric is a framework for executing (potentially non-deterministic) distributed applications in an untrusted environment. Fabric introduces execute-order-validate distributed execution paradigm, which effectively splits the traditional execution into pre-consensus (i.e., pre-ordering) execution and post-consensus validation. This separation facilitates a flexible trust model for execution of its smart contracts, also known as chaincodes, that is not impacted by the trust model considered by the underlying consensus mechanism. Beyond its novel replication approach, Fabric is best defined by the following features, which are novel in the blockchain context:

- A pluggable ordering service with multi-channel enablement. That is, Fabric supports state partitions, with each partition implementing total order semantics. Ordering service nodes (called orderers) impose total order on state updates (produced in the execution phase) using distributed consensus. The operation of orderers is logically decoupled from peers who execute chaincode and maintain the distributed ledger state. The consensus modularity goes beyond the possibility of plugging different ordering protocols in the byzantine fault-tolerant model [1], as, depending on the use case, different failure models can be assumed for orderers, such as simple crash fault-tolerant model or, in

future, the recently proposed XFT fault model [2].
- A flexible trust model for chaincode execution. A chaincode's deployers can specify the entities (or combination of entitites) that should be trusted to execute the deployed chaincode on a given channel. Chaincode deployers specify these entities by means of a policy, also referred to as endorsement policy, and can be completely independent from trust assumptions governing the ordering of transactions or the execution of other chaincodes.
- Parallelisation of chaincode execution, as not all chaincodes need to execute on all nodes.
- A modular and easily extensible membership framework. This constitutes the foundation of the permissioned nature of Fabric. Namely, as permissioned blockchains need to manage node (i.e., client, peer, orderer) identities, and access rights, membership services are a critical component of permissioned blockchains. Fabric allows for the definition and use of one or more membership abstractions, called membership service providers, each aiming to reflect an architecturally different membership management service, which is independent and securely reconfigurable. The default type of membership module supported by Fabric is compatible with X.509 certificates which are widely used by existing business membership systems.
- An access control enforcement mechanism to govern channel creation, channel participation, and administration, chaincode deployment, and chaincode execution.
- A highly efficient block dissemination mechanism from the ordering service to peers to ensure the system is able to sustain high volumes of peers, and transactions.
- A novel, two-phase smart-contract (or chaincode) deployment mechanism, to ensure that a maximum of one instance of a certain chaincode runs on each peer even if it is used to serve multiple channels.

Hyperledger Fabric V1 is due to be completed in June 2017, and it will constitute the first highly scalable permissioned blockchain platform combining the features listed above.

Links:
[L1] www.ethereum.org
[L2] www.hyperledger.org
[L3] github.com/hyperledger/fabric
[L4] www.linuxfoundation.org

References:

[1] C. Dwork, N. Lynch, L. Stockmeyer: "Consensus in the presence of partial synchrony", J. ACM, 35(2): 288–323, April 1988.
[2] S. Liu, et al.: "XFT: practical fault tolerance beyond crashes", OSDI 2016.

Please contact:
Elli Androulaki
IBM Research - Zurich, Switzerland
LLI@zurich.ibm.com

Bitcoin – Cryptocurrencies and Alternative Applications

by Aljosha Judmayer, Alexei Zamyatin, Nicholas Stifter and Edgar Weippl (SBA Research)

Exploring the real-world security of Bitcoin cryptocurrencies and alternative applications.

Bitcoin introduced a novel randomised consensus approach based on proof-of-work (PoW) which works with an unknown number of participants. The underlying concepts and techniques are collectively referred to as "blockchain". The first and still predominant use-case for blockchain technologies are cryptocurrencies.

In the context of the "Alternative Applications for Bitcoin (A2Bit)" project, we research how the fundamental principles and techniques of cryptocurrencies can be successfully applied to other problem domains, where replacing the reliance on a trusted third party can increase security, e.g., identity management and public key exchange.

Sovereignty regarding secret key management is the foundation of all security concepts based on blockchain technologies. As a first step, we performed the first large-scale empirical study to investigate how people perceive and experience the Bitcoin ecosystem in terms of security, privacy, and anonymity [1]. We surveyed 990 users of Bitcoin to determine the management strategies to protect their bitcoins and associated cryptographic keys. About half of the survey participants use exclusively web-based solutions. Also, many do not use all security capabilities offered by the Bitcoin management tool of their choice. Furthermore, they have significant misconceptions about how to remain anonymous and protect their privacy in the Bitcoin network. Twenty-two percent of the participants had already experienced monetary loss (lost bitcoins) due to security breaches and self-induced errors.

Today, more than 650 different cryptocurrencies are in circulation. The new cryptocurrencies provide additional features (e.g., Namecoin and Ethereum), alternative PoW algorithms (e.g., Litecoin and Dash), and new distributed consensus approaches [2]. The security of blockchains in a multi-PoW blockchain world has not yet been sufficiently studied.

A major challenge for introducing a new cryptocurrency is how to attract the interest of a critical mass of participants during the bootstrapping period. If not enough honest miners or mining pools join the new cryptocurrency at this crucial phase, the latter becomes vulnerable to dishonest miners and mining pools. Meanwhile, existing honest mining nodes do not have an incentive to split their effort to secure multiple PoW-based blockchains.

Alternative cryptocurrencies (e.g., Namecoin and Dogecoin) have opted for "merged mining", an approach that allows concurrent mining for multiple

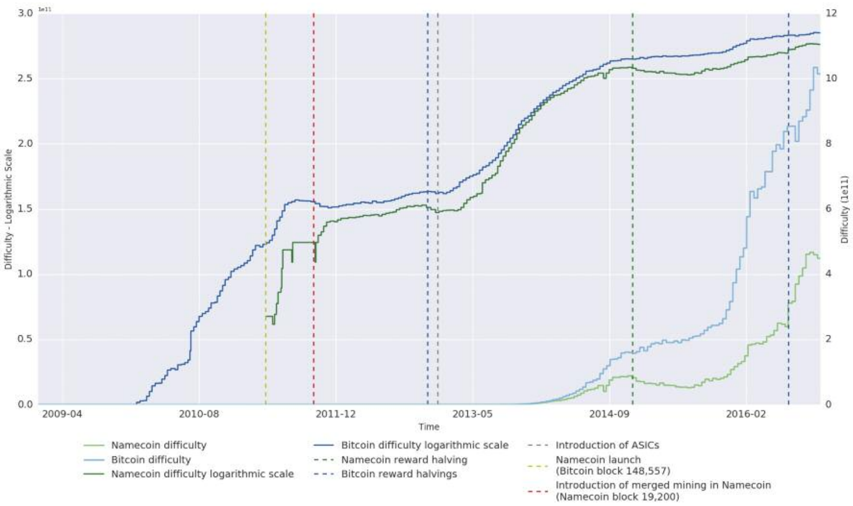

Figure 1: Difficulty development of Namecoin (green) and Bitcoin (blue) over time. Difficulty on a linear (light green/blue) and logarithmic scale (dark green/blue).

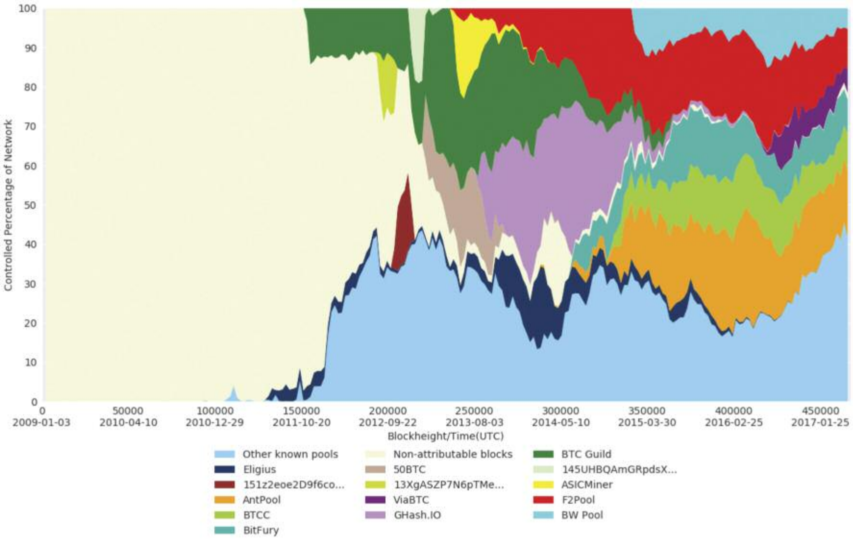

Figure 2: Distribution of Bitcoin blocks per pool over time. Each data point resembles the share among 2,016 blocks.

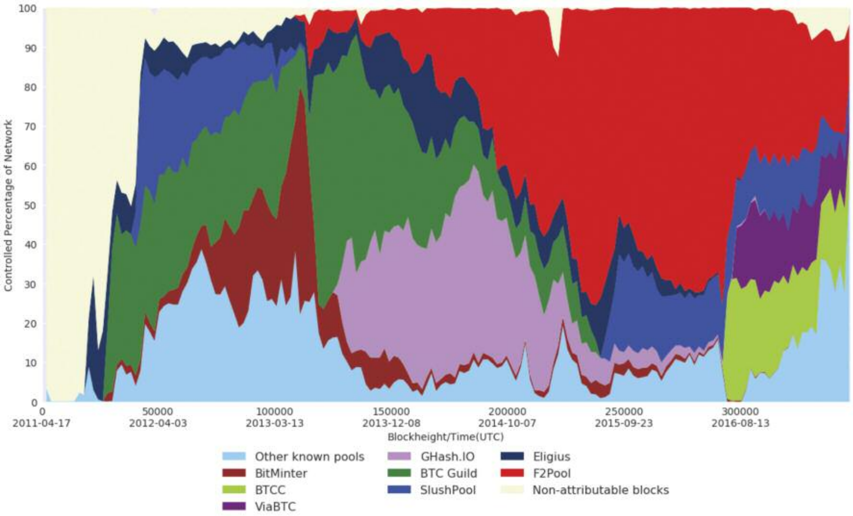

Figure 3: Distribution of Namecoin blocks per pool over time. Each data point resembles the share among 2,016 blocks.

blockchains without requiring additional PoW effort. That way, the mining power of an established (parent) cryptocurrency (e.g., Bitcoin) can contribute to increase the security of a new (child) cryptocurrency (e.g., Namecoin). In principle, this increases the security of the child cryptocurrency.

We performed a detailed analysis on two pairs of cryptocurrencies. Our findings indicate that through merged mining the child difficulty increases (see Figure 1). However, only a portion of the parent mining pools join merged mining. In Bitcoin, mining pools cannot collect a significant share of the processing power i.e., mined blocks (see Figure 2). In contrast, there are long periods where in child blockchains, some mining pools enjoy shares way beyond the theoretical limits for building a true distributed consensus (cf. Figure 3). The actual effects and implications for the mining ecosystem as well as appropriate defences are currently a work in progress.

The project A2Bit is a collaborative project of SBA Research, nic.at (the DNS registrar for .at), and the Austrian State Printing House (Österreichische Staatsdruckerei GmbH) supported by the Austrian Research Promotion Agency (FFG) under the BRIDGE Early Phase programme.

Links:

[L1] https://www.sba-research.org/a2bit/
[L2] https://kwz.me/Xt

References:

[1] K. Krombholz, A. Judmayer, M. Gusenbauer and E.R. Weippl: "The Other Side of the Coin: User Experiences with Bitcoin Security and Privacy", FC 2016.

[2] A. Judmayer, N. Stifter, K. Krombholz, and E.R. Weippl: "Blocks and Chains: Cryptographic currency technologies and their consensus systems", Morgan & Claypool Publishers, 2017

[3] A. Judmayer and E.R. Weippl: "Condensed Cryptographic Currencies Crash Course (C5)", ACM CCS 2016.

Please contact:

Aljosha Judmayer
SBA Research, Austria
+43 (1) 505 36 88
ajudmayer@sba-research.org

Identity Managenent on the Bitcoin Blockchain

by Daniel Augot (Inria, École polytechnique, and Université Paris-Saclay), Hervé Chabanne (OT-Morpho and Telecom Paristech) and William George (École polytechnique and Université Paris-Saclay)

We propose a way for users to obtain assured identities based on face-to-face proofing that can then be validated against a record on Bitcoin's blockchain. We obtain anonymity for users by making use of a scheme of Brands to store a commitment against which one can perform zero-knowledge proofs of identity and also enforce the confidentiality of the underlying data by letting users control a secret of their own. This way, users can gain access to services thanks to the identity records of our proposal.

We authenticate part of our identity with documents provided by third parties. These can be primary forms of identification like passports or driver licenses, issued by governments, but can be weaker, like bills provided by utility companies (banking, energy, phone). Our joint ongoing research between École polytechnique, Inria, and OT-Morpho (former Safran Identity and Security) consists in thinking of a blockchain as a platform for publishing such identity documents, taking advantage of the public availability, integrity and openness of the Bitcoin blockchain, while we also want to provide strong privacy for users. A natural idea, already proposed by MIT for academic diplomas [1], is to publish hashes of digitally signed certificates, using the "OP_RETURN" facility of Bitcoin transactions, which enables embedding 80 bytes of arbitrary data in a transaction. Our research is building and improving on this proposal, by considering digital certificates which do no reveal anything about their owner identity.

This can be achieved with Brands' certificates, and associated zero-knowledge proofs [2], which are as follows. Suppose an identity has n fields, $(X_1, \ldots, X_n)$, with an auxiliary random X_0, to prevent dictionary attacks. Let G be the group associated to the elliptic curve underlying Bitcoin signatures, which has 256-bit size (32 bytes). Knowing the DLREP of a given public h enables to make powerful zero-knowledge proofs. (see Figure 1).

Being in possession of the Discrete Logarithm Representation (DLREP) of h, the prover can authenticate by proving knowledge while revealing no fields, or, if required, may reveal one or several fields to the verifier. Moreover, the prover can also prove more complicated statements about her identity. This provides the user a tight control of divulged information, in a "PIMS" way [3]. Proof verification can be done by service providers, and by an intermediate service enabler (for single sign-on).

There are various ways for users to build h and convince identity providers of its validity, thanks again to Brands' proofs. The service enabler can then sign it, and h can be made public without revealing anything about its owner, except that a strong, validated, identity is blindly encoded in h. Also, the random X0 is not known to the identity provider, which thus cannot make fraudulent proofs.

Bitcoin mechanisms make it easy to insert such a h (32-bytes short) in the "OP_RETURN" field (80 bytes) of a transaction, by identity providers or utility services. Such a transaction being signed with the underlying Bitcoin mechanisms, this provides a proof that the issuer has accepted h from the user. Using the blockchain, the user can point to the transaction which contains its h, and use it to authenticate to a service provider. It is well known that Bitcoin has limited bandwidth and this problem can be alleviated by publishing roots of Merkle trees of users h's. Updating identities can be also done, and revocation seems easier using a public blockchain.

Using the Bitcoin blockchain offers several advantages. In particular, its robustness, openness, public availability, and the cryptographic platform it provides, make it easy to deploy a cryptographic solution, without heavy software engineering, and without relying on a central body for providing servers, bandwidth and availability. These features could help weak or failed states to issue identities.

We are also imagining ways to take advantage of the linkability of Bitcoin transactions. A user's proof may be linked to the certificate issuer's transaction, and/or, when convinced by the proof, the service provider could also publish an "accept" transaction, linked to the proof. A reputation can then be built, under the user's control. We are furthermore investigating the semantics of these linkability features.

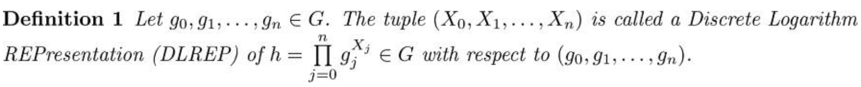

Definition 1 *Let $g_0, g_1, \ldots, g_n \in G$. The tuple $(X_0, X_1, \ldots, X_n)$ is called a Discrete Logarithm REPresentation (DLREP) of $h = \prod_{j=0}^{n} g_j^{X_j} \in G$ with respect to $(g_0, g_1, \ldots, g_n)$.*

Figure 1: Discrete Logarithm Representation of h.

References

[1] J. Nazaré, K. Hamilton, P. Schmidt: "Digital certificates project", online, source code available on https://github.com/digital-certificates, Consulted 2016, http://certificates.media.mit.edu.

[2] S. Brands: "Rethinking Public Key Infrastructures and Digital Certificates (Building in Privacy)", MIT Press, Cambridge, MA, USA, 2000.

[3] S. Abiteboul, B. André, D. Kaplan: "Managing your digital life", Commun. ACM, 58(5):32-35, April 2015.

Please contact:

Daniel Augot
Inria, Laboratoire LIX, École Polytechnique & CNRS UMR 7161, Université Paris-Saclay, France
daniel.augot@inria.fr

Hervé Chabanne
OT-Morpho, Télécom Paristech, France
herve.chabanne@morpho.com

William George
Laboratoire LIX, École Polytechnique & CNRS UMR 7161, Université Paris-Saclay, France
william.george@inria.fr

SpaceMint:
A Cryptocurrency Based on Proofs of Space

by Georg Fuchsbauer (Inria)

We introduce SpaceMint, a cryptocurrency that replaces energy-intensive computation underlying most of today's cryptocurrencies by "proof of space". Once set up, SpaceMint consumes very little energy, which will motivate regular users to participate in the mining process thereby truly decentralizing control over the currency.

How can we overcome Bitcoin's waste of electricity and tendency to concentration of control in the hands of a few by using a different commodity than computation? The idea of an electronic form of cash was first floated in the 1980s, but it has only seen wide-spread deployment in recent years. While earlier proposals relied on trusted institutions, such as banks, for the issuing of coins, Bitcoin drastically changed the economic model. Both creation and validation of coins are decentralised using a blockchain, which records all monetary transactions. Anyone who adds a new block to the chain is rewarded with freshly minted coins, but to do so, "miners" must solve a puzzle, which requires computational effort; a solution can therefore be considered a "proof of work" (PoW). The chances of mining the next block are proportional to a miner's invested computation. This way, PoW ensures distributed consensus in Bitcoin, and its security relies on no adversary gaining more computing power than the honest miners.

Although a market capitalisation of currently over 35 billion Euro has made Bitcoin the most successful electronic currency ever deployed, its expansion has come at a price. Its limited block size, which impedes scalability, has been widely discussed, but there are also concerns about long-term stability and sustainability, both directly stemming from the use of proofs of work. Bitcoin mining today is only profitable on specialised hardware, which implies high start-up costs for new miners and has resulted in a vast concentration of computing power in the hands of a few big players. This goes against the initial intent of decentralising control by letting small users benefit from spare CPU cycles to mine Bitcoin. From an environmental perspective, Bitcoin mining has led to a questionable waste of electricity in the order of hundreds of megawatts, most of it burnt in large-scale mining farms powered by application-specific integrated circuits (ASICs), which have no other use.

The first proposed alternative to PoW in the mining process was "proof of stake", as used by Peercoin. There, a miner's chances to mine the next block are proportional to the amount of currency held by the miner. Unfortunately, there are attacks against such schemes that leverage precisely the fact that mining is "cheap", in that it requires no computational effort. Proof-of-stake-based currencies also suffer from a lack of participation, as for the system to function, sufficiently many currency holders must be online and mine. In order to separate mining of a currency from just holding it, an extrinsic commodity is needed, which for Bitcoin is computation.

SpaceMint [1] is a cryptocurrency proposal by researchers from MIT, IST Austria and Inria/ENS, which replaces PoW by proof of space. Instead of computing power, miners must invest disk space, and the amount of space dedicated to mining determines the chances of adding a block. To start mining, one must first initialise one's space, which for one terabyte takes about a day. Once this is done, miners only spend a fraction of a second per block mined. While miners are incentivised to invest in hard-disk capacity, this is a one-time cost, in contrast to the perpetual electricity expenditure for Bitcoin. SpaceMint mining does not use up resources, and hard disks can be repurposed, unlike Bitcoin mining equipment. Since almost everyone has unused disk space and SpaceMint can be mined at very low setup and maintenance costs, this will lead to well-distributed mining power.

Many cryptocurrencies, such as Litecoin or Ethereum, use PoW schemes that are less "ASIC-friendly" than Bitcoin in order to counter concentration of computing power; yet they all rely on consuming large amounts of energy. Permacoin is a currency that tries to claim back some utility via a concept called "proof of retrievability", which requires miners to store useful data while still solving PoW. Burstcoin is the only existing cryptocurrency that uses disk space as its main mining resource. However, as shown in [1], it succumbs to time/memory trade-offs, meaning that with some extra computation, miners can succeed using only a fraction of the prescribed memory. The system thus potentially degenerates to a PoW-based scheme with all the above-mentioned drawbacks.

SpaceMint creates a disincentive for any additional work via the concept of "proof of space", first introduced in [2]. It is an interactive protocol between a prover and a verifier, which needed to be adapted for the cryptocurrency setting. Furthermore, since creating a proof is easy (which inherently is not the case for PoW), miners can try to mine on many branches of the blockchain in parallel, which impedes fast consensus on the legitimate branch. Not using PoW also enables "grinding" attacks where deviating from the protocol can be beneficial. SpaceMint prevents such behaviour by specific design choices and a new blockchain format. Replacing work by space can thus make cryptocurrencies greener and more egalitarian.

References:

[1] S. Park et al.: "SpaceMint: A Cryptocurrency Based on Proofs of Space", Cryptology ePrint Archive report 2015/528 http://eprint.iacr.org/2015/528

[2] S. Dziembowski et al.: "Proofs of space", CRYPTO 2015

Please contact:

Georg Fuchsbauer, DI ENS, France
+33 1 4432 2082
georg.fuchsbauer@ens.fr

Coinblesk – A Real-time, Bitcoin-based Payment Approach and App

by Thomas Bocek, Sina Rafati, Bruno Rodrigues and Burkhard Stiller (University of Zürich)

The Communication Systems Group (CSG) of the University of Zürich has been exploring the use of blockchains in several application areas. The work concluded that for practical use, Bitcoin transactions should be gathered in a batch.

Generally, blockchains pave the path towards secure data storage in a decentralised manner. They are applicable to a wide range of application domains, such as financial technologies, public registries, and Internet-of-Things (IoT) [1]. As one of the most prominent blockchain examples, Bitcoin has attained large public and research interest, since it offers the first solution for a secure and fully decentralised crypto-currency. Thus, the Communication Systems Group (CSG) of the University of Zürich decided to focus research work on (a) real-time payments with Bitcoins [2, 3], which was trialled at the UZH Mensa [L1] and presented at public fairs [L2], (b) the use of blockchains within IoT, especially the supply chain in the pharmaceutical industry which is highly regulated, and (c) blockchain-based countermeasures for Distributed Denial-of-Service (DDoS) attacks by utilising Smart Contracts (SC).

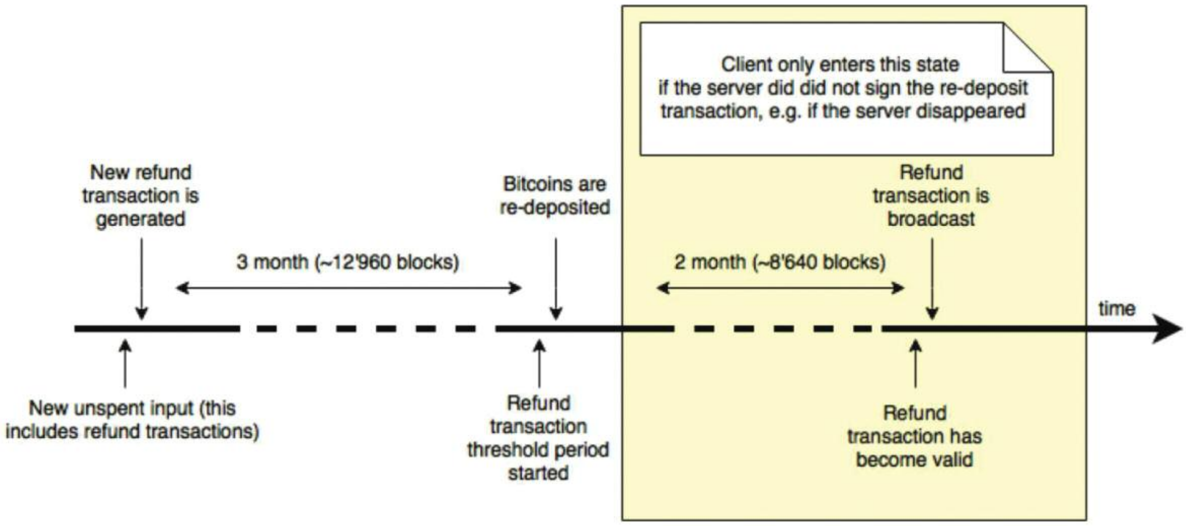

Figure 1: Coinblesk's refund transaction time-line.

Blockchain technology has become popular for multiple use-cases, such as IoT, crypto-currency, and security, because blockchains are inherently backed by Smart Contracts. They are defined as formalised protocols to facilitate, verify, or enforce the negotiation or performance of a contract. In this sense, Bitcoin, considered as the pioneer implementation of blockchains, and especially the Bitcoin Script, serve as the first SC for this crypto-currency. Besides theoretical work, the trial deployment of blockchains and their application-specific combination with SCs deliver valuable insights for distributed systems' operations. Specific blockchain benefits include a fully decentralised system operation, transaction transparency, immutability, and security combined with selected areas of legally binding interactions.

In this context the new Coinblesk approach [2, 3, L1, L2] belongs to the use-cases of crypto-currencies. It is an instant payment wallet with Bitcoins and minimal trust with the strategic goal to generalise and optimise its payment protocol to support other crypto-currencies, while maintaining security, privacy, and convenience as key. The CoinBlesk app for Android includes a Bitcoin payment server, where the seller and the buyer are able to handle Bitcoin payments. This safe and fast mobile payment method is contactless, using Near Field Communications (NFC) technology, without the need for swiping, signing, or PIN. To reach a transaction delay below one second, a multisig (multi signature) mechanism was designed such that the Coinblesk server cannot transfer funds without the signature of the client. Since sending every transaction immediately to the blockchain reveals the current limitations of Bitcoins, and the current fee of an average transaction is more than US $2, these transactions are batched and transaction fees are reduced by performing the clearing operation at the server, where the user can specify an amount stipulating when clearing should be made. Only once that amount is reached, is a transaction sent to the Bitcoin blockchain. Thus, if a transaction is cleared on the server (not yet sent to the Bitcoin blockchain) a virtual balance is maintained in order to acknowledge the payment within this one second limit.

This mechanism reduces the number of transactions – termed "batching transactions" – sent to the Bitcoin blockchain and, thus, lowers the average transaction fees of these transactions. The system has been built in such a way that the user can set that maximum amount, since only the user can determine the trust level to be reached. In turn, the system has to broadcast these batched transactions to the Bitcoin blockchain, e.g., if the user sets the limit at €100 and if the virtual balance reaches this value, all accumulated transactions are broadcast. This approach was chosen over the Lightning network's approach [L4], since its technical complexity is lower and more importantly it also works with transaction malleability. The current Coinblesk design can be optimised further, once transaction malleability is solved in the Bitcoin network or any another crypto-currency, such as Litecoin, which does not suffer from malleability, is used. However, as mentioned above, the Coinblesk app does not follow the fully trustless approach in such cases, since the Coinblesk server requires this minimal trust up to the amount specified by the user.

All funds deposited in Coinblesk are held at a 2-of-2 multisig address, which means that even if the operator of the Coinblesk server is intentionally malicious, he will never be able to steal a user's funds. In the case of a Coinblesk server hacking and private keys being stolen, the hacking could only be successful if hackers were able to gain access to the user's private keys as well in order to steal bitcoins. Also, if the Coinblesk server disappears, clients are no longer able to spend their bitcoins. This is a major problem, because Swiss law requires customers of a payment service to be able to gain full access to their funds in any situation, and especially if the operator of a payment system should become bankrupt – or in the case of the Coinblesk service, it might be hacked. Additionally, all Coinblesk clients need to trust that the system will not disappear.

Thus, the effective solution to this problem is a "refund transaction" as time-lined in Figure 1. A refund transaction is a pre-signed, time-locked transaction, which sends all client funds to an address, exclusively controlled by that client. Therefore, a refund transaction is automatically created by the Coinblesk app as soon as a new unspent output appears in the wallet – in particular, whenever bitcoins are received or a transaction is created. The app takes all the unspent outputs and creates a single transaction sending all bitcoins to an address of a private key that is derived from the client's private seed. The client signs this transaction and returns it to the server. The server checks that the transaction is in fact time-locked, signs it, and returns the transaction fully signed back to the client. Now, the client is in possession of a valid, fully signed refund transaction that becomes valid as soon as the time-lock expires. Thus, in case the Coinblesk server suddenly disappears, a client can broadcast the refund transaction and regain control over all their bitcoins.

In conclusion, the experience with the Coinblesk design and implementation as well as experience from other applications, such as the pharmaceutical supply chain [L3, L5], provides useful information about scalability, energy efficiency, ease-of-use, and some insights into customer acceptance. These results should be widely applicable in the blockchain world.

Links:

[L1] http://www.csg.uzh.ch/csg/en/news/Bitcoins.html
[L2] http://www.csg.uzh.ch/csg/en/news/ coinbleskatCeBIT.html
[L3] http://www.csg.uzh.ch/csg/en/news/kickstart-accelerator.html
[L4] https://lightning.network/lightning-network-paper.pdf
[L5] https://modum.io/

References:

[1] T. Bocek, B. Stiller: "Smart Contracts – Blockchains in the Wings", in: C. Linnhoff-Popien, R. Schneider, M. Zaddach (Edts.): "Digital Marketplaces Unleashed", Springer, 2017.

[2] A. D. Carli: "Protocol Improvements in CoinBlesk – A Mobile Bitcoin Instant Payment Solution", Master Thesis, Univ. Zürich, Department of Informatics, Communication Systems Group, Zürich, Switzerland, April 2016.

[3] R. Voellmy: "CoinBlesk, a Mobile NFC Bitcoin Payment System", Bachelor Thesis, Univ.Zürich, Communication Systems Group, Department of Informatics, Zürich, Switzerland, August 2015.

Please contact:

Thomas Bocek, Sina Rafati, Bruno Rodrigues, Burkhard Stiller
University of Zürich, Switzerland
[bocek¦rafati¦rodrigues¦stiller]@ifi.uzh.ch

Bitcoin Unchained

by Christopher Carr, Colin Boyd (NTNU), Xavier Boyen and Thomas Haines (QUT)

Bitcoin's distributed ledger is an innovative way of solving the double spending problem in a decentralised system. However, it causes incompressible transaction delays and incentivises consolidation of mining power. We ask, is it possible to eliminate these problems without losing the decentralised principles that Bitcoin was built on?

Over eight years have gone by since Bitcoin's deployment, and it is still going strong. While there are many explanations for its success, the innovative backbone structure – the blockchain -– which has inspired so many alternative systems, undoubtedly plays a leading role in this story.

Blockchains store the state of the transactions in the system. Users compete to form new blocks, which confirm both new and all existing transactions in the previous blocks. Those who create blocks first are rewarded with cash in the system.

Despite the blockchain innovation, there are some fundamental problems that lie in its design, which stem from the blockchain itself, and affect all similar systems.

Two major problems which are inherent to almost all blockchain models are:

1. Consolidation of power: Users are incentivised to form into groups to maximise their expected reward over time. Cartels formed in this manner are commonly referred to as mining pools.
2. Incompressible delays: All transactions have a delay before they can be considered confirmed within the system. In Bitcoin itself, this is exacerbated by block size restrictions, a source of heated debate within the community. Recently, almost all blocks have been full to capacity of transactions, and as of the time of writing have fees for posting transactions over 10 USD.

Previously, there has been a line of inquiry that looks at alternative ways of designing proofs-of-work to avoid mining pools. Miller, Kosba, Katz and Shi [1] create a proof-of-work system that allows for any pool member to cheat and reap all the rewards for themselves. Importantly, they show that a cheater can do this without any way of being caught, thus removing the incentive for mining pool formation. Lewenberg, Somplinsky and Zohar [2] design a system that allows for collections of transactions to be confirmed in such a way that overlapping blocks can be counted along with the transactions contained within them.

Our motivation stems from simultaneously addressing these two fundamental problems of consolidation of power and incompressible delays. In a joint research effort, which is a collaboration between the Norwegian University of Science and Technology [L1] and Queensland University of Technology [L2], we ask: "What happens if we remove blocks altogether?" Instead of collecting multiple transactions together, whenever you wish to create a transaction you simply reference two recent, existing transactions.

Once blocks are removed, we need a way of securing transactions against double spending. To achieve this, we look to the incentive mechanisms, and use these to promote the desired characteristics. We incentivise the collection of recent previous transactions by increasing the reward for doing so. This can also be thought of as a form of small blocks, but removing the enforced confirmation delay.

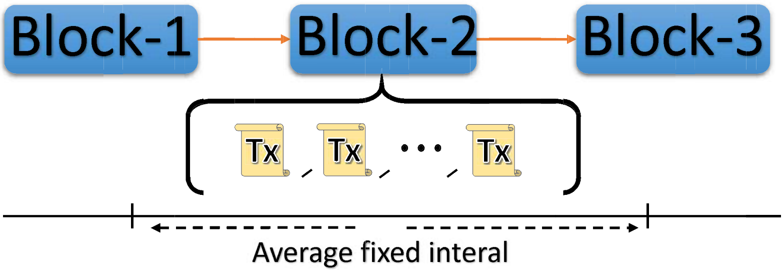

Figure 1: Blockchain model: Transactions (Tx) are collected together over some fixed average time interval and grouped into blocks, confirming the full group of transactions.

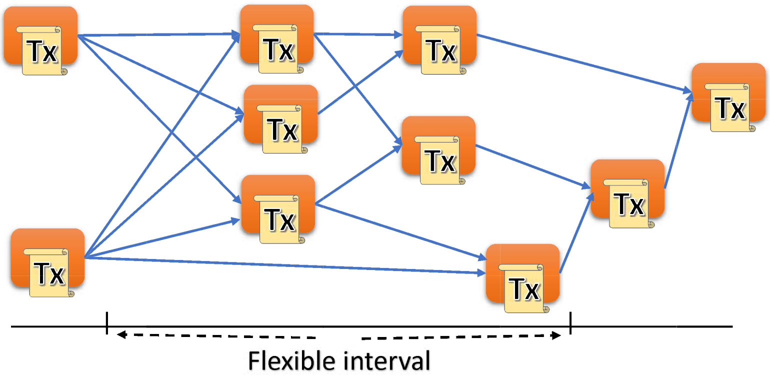

Figure 2: Blockchain free model: Transactions (Tx) are collected indivudually over a flexible time period and confirm previous transactions.

To highlight these aspects, Figure 1 shows a standard blockchain model, where transactions are collected together and formed into a block. Contrast this with Figure 2, which shows the block-less model, where transactions confirm only two previous transactions.

So far, we have developed a blockchain free system [3], and demonstrated the security of the system under the assumption of a majority of rational users. We show that the incentive mechanisms we put in place encourage transactions to finally group together at the head of the chain, where all previous transactions are confirmed from the leading transaction - a property we call convergence.

We believe this novel approach represents a large step forwards in tackling these highlighted blockchain problems. Our focus now is on addressing implementation decisions. The challenge is to select appropriate parameters that do not undermine the theoretical underpinnings. Our hope is that by designing and implementing a system in this way, we can get closer to the true ideal of a decentralised digital cash system.

Links:
[L1] http://www.ntnu.edu/iik/nacl-lab
[L2] https://kwz.me/Xd

References:

[1] A. Miller, et al: "Nonoutsourceable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions", ACM Conference on Computer and Communications Security 2015: 680-691.

[2] Y. Lewenberg, Y. Sompolinsky, A. Zohar: "Inclusive Block Chain Protocols", Financial Cryptography 2015: 528-547.

[3] X. Boyen, C. Carr, T. Haines: "Blockchain-Free Cryptocurrencies: A Framework for Truly Decentralised Fast Transactions", IACR Cryptology ePrint Archive 2016: 871 (2016).

Please contact:

Christopher Carr, NTNU, Norway
ccarr@ntnu.no

A Holistic Approach to Smart Contract Security

by Nicholas Stifter, Aljosha Judmayer, and Edgar Weippl (SBA Research)

Secure Execution of Smart Contracts (SESC) aims to identify and analyse security aspects of smart contracts and the platforms on which they execute from a holistic viewpoint. We focus on the long-term sustainability and security of smart contract infrastructures.

The advent of Bitcoin as a decentralised cryptocurrency has a fundamental impact on both practical applications and scientific research, reaching well beyond its immediate use-case as a form of currency. Many concepts that previously needed to rely on a trusted third party now become feasible as decentralised implementations, thanks to Bitcoin's underlying blockchain technology.

One such application relates to smart contracts i.e., "a computerised transaction protocol that executes the terms of a contract" [1]. Conceptually, smart contracts can be understood as program code that is executed for transacting parties. A blockchain-based smart contract platform serves as a decentralised arbiter to both verify and enforce the execution of these smart contracts based on the platform's defined rules. In practice, smart contract platforms may be more closely related to the field of trusted computing as they offer the ability to execute code with relatively high trust in a decentralised environment.

Smart contract platforms generally follow an open, permissionless model where anyone can deploy their own smart contract code and where both publishing and executing smart contracts incurs transaction fees. The expressiveness of the programming language and code used to define smart contracts plays an important role in such systems because it greatly influences what can and cannot be achieved. As an example, Bitcoin provides limited smart contract functionality because its transactions are governed by the execution of stateless scripts in a simple, non-Turing-complete stack-based language. Other platforms, such as Ethereum support complex and stateful Turing-complete contracts can cover a much wider range of application scenarios. The correct and secure execution of such smart contracts depends not only on the contract's code and its execution environment itself, but also on the underlying properties of the distributed system that actually facilitates such a decentralised smart contract platform [2]. Therefore, a holistic approach towards smart contract security which integrates all these aspects and their interactions is not only prudent, but necessary.

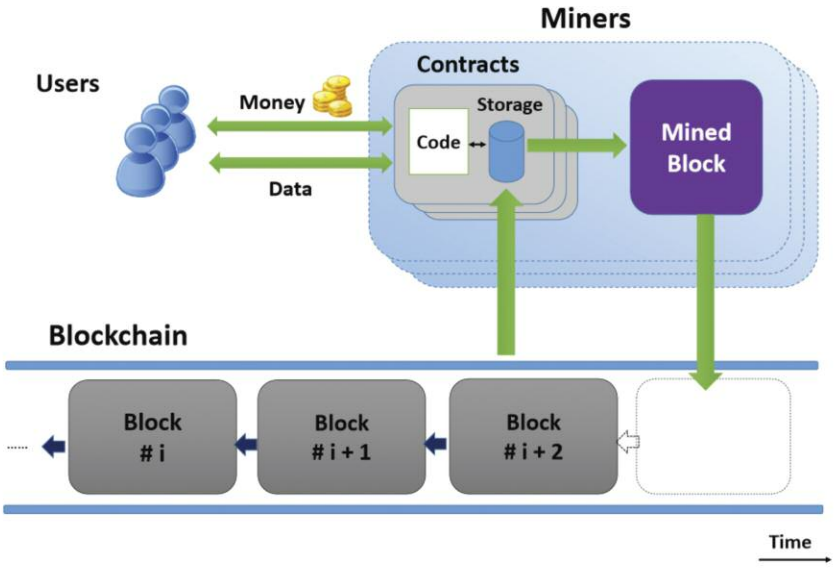

Figure 1: A blockchain-based smart contract according to [L2].

The difficulties and obstacles encountered when trying to ensure both the correctness and the security of more complex smart contract code are manifold, and it is not surprising that the recent history of decentralised smart contract platforms contains a number of serious security incidents [3]. Many of these incidents can, at least partially, be attributed to a lack of established paradigms and best-practices and in particular the complex interaction patterns of the different components and aspects that arise in decentralised smart contract platforms.

The "Secure Execution of Smart Contracts (SESC)" project aims to systemise available technologies and explore the emerging requirements for safely and reliably creating and maintaining smart contracts and their governing (blockchain) infrastructures in the long-term. This encases elements and approaches such as formal verification and automated analysis; security impact analysis on smart contract infrastructures on client side devices; and the applicability of container technologies. These elements outline how properties of the underlying distributed system may adversely affect both smart contracts and the hosting platform itself.

One aspect of particular interest is what effects later parts of the development lifecycle of smart contracts and their governing infrastructures will have on security. While this can have a significant impact on both security and maintainability, the topic area has received little attention from developers and researchers alike. It is largely unclear how future approaches and solutions towards sustainability and scalability might influence smart contracts that are being deployed today. The current predominant paradigms render smart contract code difficult, if not impossible, to change once it has been deployed. Clearly, the topic of smart contract security is of paramount importance if such decentralised smart contract platforms are to gain widespread adoption. SESC will provide much needed insights into this relatively new problem domain. We will explore the fundamental requirements for long-term maintenance and sustainability of smart contract systems. Through SESC, we aim to both identify and address new application domain-specific problems, thereby enabling the community to take a more proactive stance towards smart contract security.

SESC is a collaborative project of SBA Research, Venionaire Capital, and handcheque. It is supported by the Austrian Research Promotion Agency (FFG) under the BRIDGE 1 programme and kicked off in January 2017.

Links:
[L1] https://kwz.me/Xc
[L2] http://eprint.iacr.org/2015/460.pdf

References:

[1] N. Szabo: “Smart Contracts”, 1994, https://kwz.me/XS

[2] C. Natoli and V. Gramoli: “The blockchain anomaly”, IEEE NCA 2016.

[3] N. Atzei and M. Bartoletti and T. Cimoli: “A survey of attacks on Ethereum smart contracts (SoK)”, International Conference on Principles of Security and Trust 2017.

Please contact:

Nicholas Stifter ,SBA Research
+43 (1) 505 36 88
nstifter@sba-research.org

Correctness of Smart Contracts for Consistency Enforcement

by Thomas Osterland and Thomas Rose (Fraunhofer FIT)

Smart contracts are a proposed mechanism to help maintain consistency among data and transactions. They are automatically triggered by the conduct of a transaction and they also function to safeguard transaction histories. A cascade of automatically initiated smart contracts could result in data errors and smart contracts interfering with one another, but correctness can be assured by means of model checking.

Blockchain technology was initially introduced as a transaction management technology that transfers centralised control to a distributed environment with new means of consistency enforcement [1]. Platforms with smart contracts extend the original blockchain protocol by a process automation to proactively maintain consistency among data, and in particular transactions, by enabling a full automation of agreements and the autonomous adherence of these agreements [2].

An appealing use-case of smart contracts is the documentation and clearing of micro-payments between parties in smart grid environments - for example, among members of a community that share a renewable energy network and e-mobility-oriented consumers. A variety of different small plants homogeneously distributed over the grid produces, controls and invoices the energy flow, in contrast to traditional electrical plants in large-scale grids that operate in a centralised fashion. All activities by entities in the grid can be managed as well as safeguarded by a smart contract.

The major benefit of using blockchain-enabled smart contracts for such grids is that participation on the network is permissionless and self-regulating. Someone who erects a wind turbine in their garden can directly participate in the global energy supply, independent from central trusted intermediaries. Blockchain technology allows the coequal participation of different parties on the smart grid without the requirement that they trust each other. Small contributions from individuals will be negotiated in the same way as large contributions by industrial companies. Smart contracts thus allow the elimination of intermediaries both in topological regard - for instance, large suppliers of electricity are replaced by individual smart contracts, and in regard to function - for instance, tasks like bookkeeping and payment are handled by smart contracts. In addition, a blockchain provides 100% uptime and its decentralised nature protects against catastrophes and acts of terrorism.

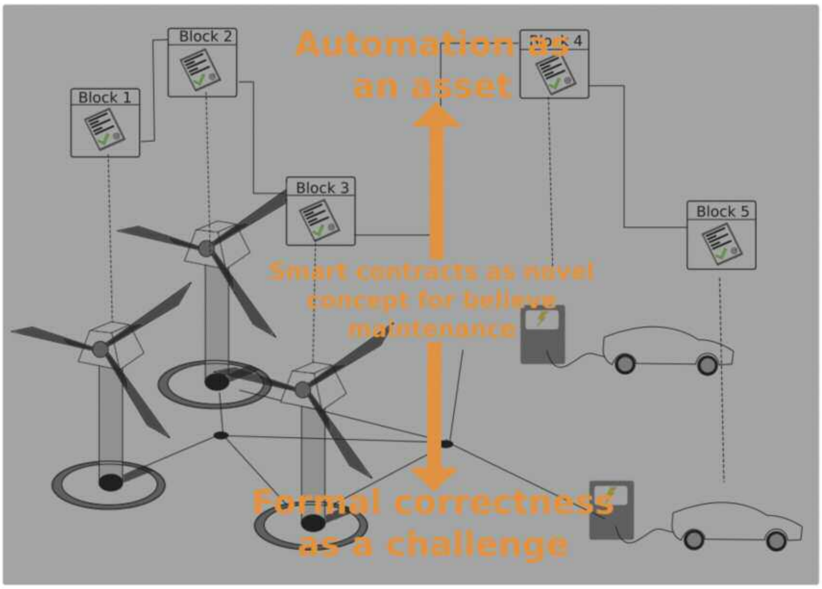
Block 1
Block 2
Block 3
Block 4
Block 5
Automation as an asset
Smart contracts as novel concept for believe maintenance
Formal correctness as a challenge

Figure 1: Chances and challenges of smart contracts.

While the protected execution of smart contracts is ensured by the blockchain protocol, there is no approach that ensures the correctness of the rules encoded in the smart contract.

As a consequence, the correctness of a smart contract must be ensured in advance, before formal instantiation in a blockchain. This is certainly important for developers, as well as suppliers and consumers that rely on the soundness of a smart contract. Moreover, it furnishes a source of trust for users because trust is maintained by algorithmic concepts. When proving the correctness of smart contracts, a model of the actual correct behaviour of a contract is necessary in first place. Determining whether a contract reacts correctly is not always as trivial as it seems, and proving it (automatically) means that the behaviour must be defined as conditions in a formal notation, for instance (temporal) first order logics.

The process of testing whether a given system satisfies these formally introduced conditions is called formal verification. In the example of the limited selling price a corresponding condition ensures that the selling price depends on the movements of the energy market and that it is not limited to a certain range.

One approach with a high degree of industry acceptance is model checking [3]. The idea is to analyse the state space of a program. A program state is the valuation of every variable of a program for a given execution step. Running a program means executing the program instructions consecutively and every instruction changes the storage and thus alters the program state. Certain (forbidden) states in the graph represent situations that contradict the desired behaviour of the contract. For instance, in the smart grid example, every state from which we cannot reach a state in which the selling price exceeds a specified limit. Model checking can then confirm whether forbidden states are reachable. Besides the checking of forbidden behavior, model checking can also test liveness properties, that describe desired behavior that must eventually occur.

To-reiterate, smart contracts certainly provide a powerful functional surplus for maintaining the consistency of transactions in applications governed by blockchain technology. However, the intended level of automation might cause cascading effects that have to be checked by formal methods of algorithmic proof. Model checking appears as a favourite candidate because of its balance among expressiveness of conditions to be verified and computational complexity for verification. In case of erroneous contracts, it provides counterexamples that support debugging and its application does not require expert knowledge.

References:

[1] J. McKeogh, R. Cheong. "Consensus: Immutable agreement for the internet of value," in KPMG, 2016

[2] S. Underwood. "Blockchain beyond bitcoin," in Commun. ACM 59, 2016, pp 15-17.

[3] K. Havelund , T. Pressburger. "Model Checking Java Programs Using Java PathFinder," in T.STTT (2000) 2:366, 2000, pp 366-381.

Please contact:

Thomas Rose, Thomas Osterland
Fraunhofer FIT, Germany
+49 2241 14 2798, +49 2241 14 3618,
thomas.rose@fit.fraunhofer.de
thomas.osterland@fit.fraunhofer.de

Implementation of a Blockchain Workflow Management Prototype

by Gilbert Fridgen (Fraunhofer FIT), Bernd Sablowsky (Norddeutsche Landesbank) and Nils Urbach (Fraunhofer FIT)

Blockchain technology offers huge potential to various industries and application areas. In a joint applied research project, Fraunhofer Institute for Applied Information Technology (FIT) together with Norddeutsche Landesbank (NORD/LB) identified (inter-company) workflow management as a promising application area and developed a Blockchain prototype for a documentary letter of credit in the international shipping business. In addition to the project's explicit outcome – a Blockchain prototype and strategic support for Blockchain innovation management – the joint project revealed important insights into the technology's applicability in the field.

Norddeutsche Landesbank (NORD/LB) regularly evaluates technology innovations regarding their strategic implications. In this context, NORD/LB, in cooperation with Fraunhofer FIT, started a Blockchain project to identify potential application scenarios and the implementation of a prototype. The aim of the project was to familiarise with Blockchain technology to obtain transferable results for future activities with Blockchain or Distributed Ledger Technology.

A Specific Workflow Management Use-case

Inter-company workflow management is a promising application area for Blockchain technology since requirements of transparency and traceability often demand cumbersome manual processes in today's businesses. In addition, regulatory requirements can make these processes even worse. Having set the objective to gain as many insights into this application field as possible, the interdisciplinary project team consisting of technology, business, and innovation experts started to analyse one specific use-case in this area. The team identified a banking use-case that is accompanied by plenty of paperwork, which is currently literally sent around the world, e.g., to gather signatures from all relevant process participants. Specifically, the project team implemented a Blockchain prototype of a documentary letter of credit. Simply speaking, a letter of credit is a payment instrument that guarantees an exporter the payment of its goods as long as it fulfils certain conditions such as submitting the correct documents to the corresponding banks. The simplified process is depicted in Figure 1. These documents make the process quite lengthy as they must be sent paper-based from one process participant to the other.

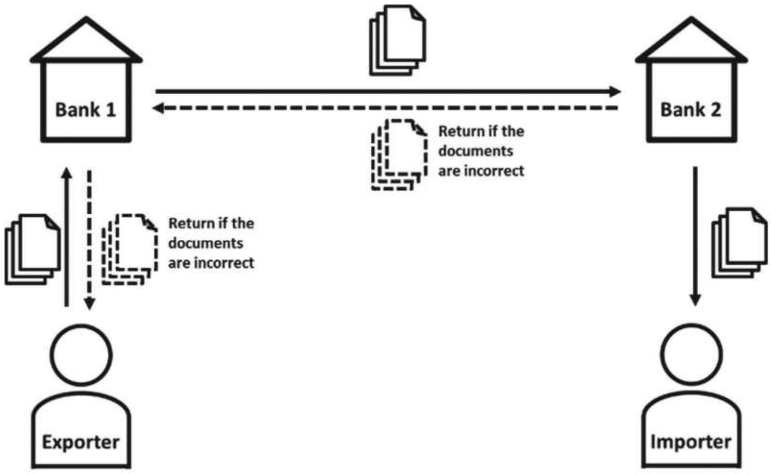
Bank 1
Bank 2
Return if the documents are incorrect
Return if the documents are incorrect
Exporter
Importer

Figure. 1: Document flow in a documentary letter of credit.

Analysing the Disruptive Potential of Blockchain

The project team first discussed the conceptual and technological features of prevailing Blockchain solutions and how these might improve the current state of the process. In particular, Blockchain's tamper-proof information 'storage' ability, thus auditability, could be a major advantage. For example, storing process-related documents in a secure cloud, and hashing their content can provide proof that the original documents have been digitalised and made available to all process participants. As a second step, standardised documents for a letter of credit could be treated fully digitally and would thus completely avoid paper-based documents. Avoiding paper-based documents can reduce costs associated with stationary, postage, and long transportation times.

Transportation time alone (per courier) from one process participant to another can currently take more than two weeks. Furthermore, digital, tamper-proof documents enable parallel processing within single process steps, since the second participant (e.g., bank 2) does not need to wait for the first one (e.g. bank 1) to finish their process step and send the documents on.

Using smart contracts can result in even more important process improvements. A smart contract is a computer program that executes a job if a predefined condition is met. For example, if a stock price drops below a certain level, the share is sold. Within the entire document process (Figure 1), diverse control procedures are necessary as the payment is bound to predefined conditions. Hence, currently bank employees must manually check if these conditions are met. A lot of this work is repetitive and can possibly be implemented using smart contracts with their aforementioned predefined conditions. For example, if a particular shipping date is crucial, this date might be implemented within a smart contract and checked against the date that the shipping is actually accomplished, e.g., when the exporter signs the shipping of goods in the harbour. Then the smart contract is automatically approved if this condition is met - and the next process step is triggered – or rejected if the condition is not met. Hence, a manual process is no longer necessary for such repetitive tasks and the whole processing time can be reduced significantly.

In line with these examples, the implemented prototype puts emphasis on process improvement and ease of use. We make use of the decentral nature of Blockchain systems and include participants of different companies, use smart contracts for process automation, and create a fully digitalised process.

The Project's Insights

NORD/LB together with Fraunhofer FIT identified a set of potential use-cases for a prototypical implementation that were purposefully different from prevailing cash system use-cases like Bitcoin. In particular, workflow management allows for insights into how to make use of Blockchain properties. Specifically, the tamper-proof 'storage' of information as well as process automation using smart contracts were discussed and implemented. Furthermore, the decentralised nature of Blockchain could allow the entire process to be developed in a direction that radically changes or even removes the role of intermediaries (here the banks in Figure 1). Of course, given the early stage of the technology's development and the current lack of standardisation and legislation this is a glimpse into the future of Blockchain technology and digitalisation, but it is neither far-fetched nor improbable.

Links:
https://kwz.me/Xy
https://kwz.me/XH

Please contact:
Gilbert Fridgen, Nils Urbach
Fraunhofer FIT, Bayreuth, Germany
+49 921 554711, + 49 921 554712
gilbert.fridgen@fit.fraunhofer.de,
nils.urbach@fit.fraunhofer.de

Bernd Sablowsky
Norddeutsche Landesbank, Hannover, Germany
+49 511 3615020
bernd.sablowsky@nordlb.de

Proofs of Work - the Engines of Trust

by Alex Biryukov (University of Luxembourg)

Customisable proofs of work and memory hard functions are investigated by the SnT&CSC/CryptoLUX team at University of Luxembourg.

Proofs-of-work (PoW) are at the core of most of the present day blockchains and cryptocurrencies. These are the tools that make large public distributed ledgers possible, since they replace the difficult to quantify and manage trust by hard to forge mathematical computations.

Proofs-of-work have been first proposed as a way to mitigate the spam problem and were later used by Nakamoto in the Bitcoin protocol. First blockchain proofs of work were often based on iteration of cryptographic functions (double SHA-256 in Bitcoin) until the result shows a special lucky number. Due to cryptographic properties of the function this is similar to winning a lottery and the lucky "miner" is rewarded with cryptocurrency. One of the smart decisions in Bitcoin was to make this winning chance adjustable depending on the available market of miners. However after the first cryptocurrencies started to gain popularity and value, the mining process entered into an arms race of mining hardware: from desktop computer mining, to GPU, FPGA and finally to ASIC mining.

Today Bitcoin mining is concentrated in the hands of about a dozen mining farm operators, who have access to optimised ASICs, cheap electricity (for example in China) and environments that facilitate cooling (in Scandinavia, for example). Bitcoin miners serve as validators of the transactions in the Bitcoin public ledger and it is miners who decide what will be included in the ledger. Thus mining centralisation goes against the democratic principles declared in the original Bitcoin whitepaper. Moreover current blocksize/SegWit debate demonstrates that Bitcoin is hostage to its mining conglomerates, who have made huge investments into Bitcoin "printing" hardware.

This situation is due to high parallelisation of the Bitcoin mining process: an ASIC full of SHA-256 cores is more than 30,000 times more energy efficient in Bitcoin mining than the general purpose CPU. In order to remedy this problem, memory-hard functions were proposed. Our CryptoLUX team [L1] has been working on a project to design democratic proofs of work and we have come up with two different designs. One is based on our team's memory-hard password hashing design called Argon2, which won the 2015 international password hashing competition (PHC). In [1] we propose building memory-hard, but easy to verify PoW, using Merkle hash tree on top of the Argon2 hash chain (Figure 1). Then we compute the hash of the tree root together with a nonce and apply Fiat-Shamir's method to produce queries about random locations in the Argon2 hash chain.

If attackers try to cheat and store only a fraction of the Argon2 chain, they are very likely to be caught as they will not be able to demonstrate the knowledge of the proper Argon2 chain elements with their correct paths in the Merkle tree. This scheme is called MTP and can be instantiated with any memory-hard hash function. In another work [2] we followed a completely different pass, using the well-studied generalised birthday problem. In this problem given k lists of n-bit numbers one is asked to find a set of elements, one per list, such that the XOR of all the numbers is zero. The best currently known algorithm

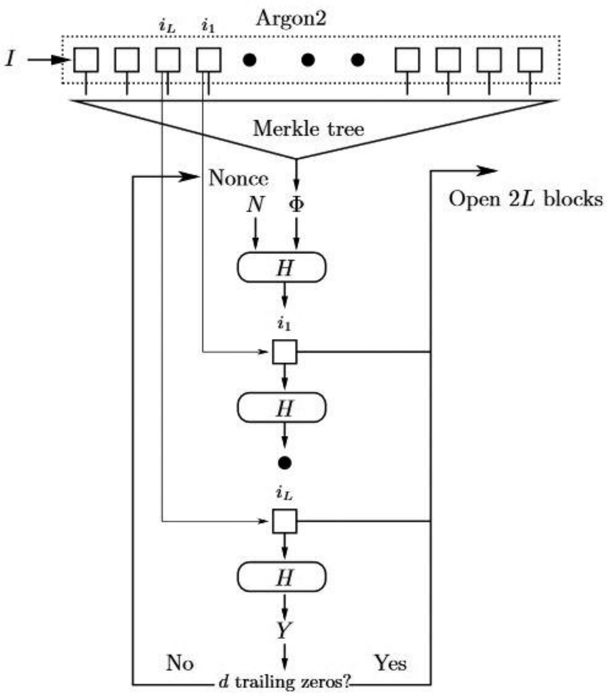

Figure 1: Merkle-tree based Proof-of-Work with light verification.

Proof-of-Work puzzle

was developed by D. Wagner, and in order to find the solution it needs to store the numbers in the lists and to constantly sort the new resulting lists (which is memory-hard). An improvement in Wagner's algorithm would be an important breakthrough for cryptography. On the other hand, once the solution is found, its correctness is very easy to verify. We use this problem (with some important hardening modifications, e.g., algorithm binding) to build a new memory-hard proof of work function that we call Equihash. This function is now being used in one of the popular cryptocurrencies - Zcash. So far it holds the ASIC resistance promise and is mined on CPUs and GPUs.

One of the main concerns with proof-of-work based cryptocurrencies is their waste of energy. Indeed the amount of electricity being burned just for Bitcoin is approaching the energy consumption of a country like Denmark. It is a valid question whether this is a justified price to pay for the running of a trustless public ledger. While electricity can be very cheap in some places (e.g., an old hydroelectric power plant in the middle of a rural area) and green forms of energy might make electricity cheap in the future, many researchers have been wondering if it is possible to avoid the energy waste. In our team we have started to explore "greener" alternatives, such as proofs-of-stake consensus protocols, which involve economic and game-theoretic reasoning or distributed ledgers based on Byzantine fault tolerance (BFT), which have high throughput in terms of transactions but require permissioned blockchains due to trust and scalability issues.

Link:
[L1] https://www.cryptolux.org

References:

[1] Alex Biryukov, Dmitry Khovratovich: Egalitarian Computing. USENIX Security Symposium 2016: 315-326.
[2] Alex Biryukov, Dmitry Khovratovich: Equihash: Asymmetric Proof-of-Work Based on the Generalized Birthday Problem. NDSS 2016.

Please contact:

Alex Biryukov
University of Luxembourg
+352 46 66 44 6793
alex.biryukov@uni.lu

Design Requirements for a Branched Blockchain Merging Algorithm

by Arthur Melissen (Storro B.V.)

Current blockchain technology as used by Bitcoin [1] and others uses a consensus model to determine the head of the chain and lets divergent branches starve, causing information loss. Information loss can be problematic for alternative blockchains in which connectivity could be reduced for longer periods of time and a consensus model is infeasible.

In order to support sparsely connected clusters of machines working on independent branches of a blockchain, divergent branches must be able to be merged with each other by a merging algorithm that carries the consensus of all parties involved.

We discuss a set of requirements that a blockchain merging algorithm must achieve, and show how such an algorithm can be used to facilitate parallel disconnected operations in an alternative blockchain.

Bitcoin's blockchain model allows an unknown number of participants to use its chain without a central authority. Because any client may append data at any time, branching the chain is both implicit and inevitable. In order to avoid merging branches or dealing with conflicting transactions, Bitcoin lets its users choose a main branch based on an estimated amount of work done on competing branches. Because Bitcoin is a system that registers financial transactions, it has a great need for consistency by distributed consensus in a relatively short amount of time. The process of distributed branch selection is effective in achieving this, but artificially limits the network to a single branch and thereby introduces a scaling bottleneck. Alternative blockchains may be designed to perform other tasks, such as publicly verifiable calculations, snapshotting filesystems, anonymous and

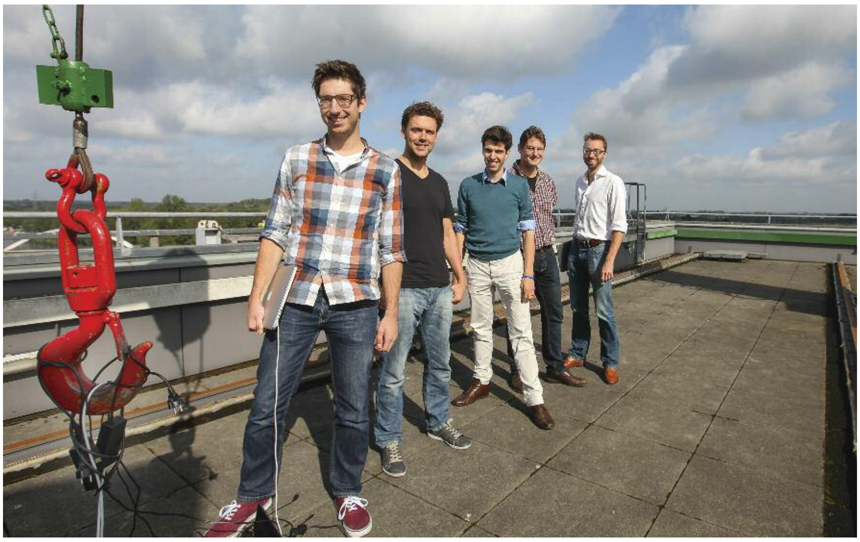

Team at Storro.

auditable voting, distributed notary systems to provide proof-of-possession of documents, and other distributed and anonymous services.

Many such alternative blockchains may not have a need for the same strong demands on consistency that Bitcoin does, and may require performance that cannot be achieved by distributing a single chain across all nodes in the entire network. For instance, a distributed filesystem might register its updates in different branches for individual servers and only periodically synchronise the filesystem by merging each server's branch with the other branches using a branch merging algorithm.

At Storro, we design and use blockchain models where it is not achievable to maintain a hierarchy of branches or even an ordering between different branches. It might be the case that a client working on one branch is not even aware of all other branches. In such an environment it is necessary for the merging algorithm to function between any pair of branches, independent of their lineage. Merging might be done opportunistically and asynchronously between intermittently connected clients and may fail due to network outages or power failures.

In effect, this means that the merging operations between semi-connected and unreliable clients is a chaotic process, but must still provide an eventually consistent state between all clients involved.

The simplest merging scenario is that of a branched blockchain consisting of only two branches. It should be clear that merging two different branches A and B should provide the same merged result state on both clients. This means that the merging operation must be deterministic and commutative. Since the algorithm is deterministic, it is an intuitive notion that the merging algorithm should most likely be automated, and not require any human intervention.

For blockchain networks of more than two branches to reach consensus, we must ensure that the result of the merging operation between branches *A* and *B* merged with *C* yields the same result as the merge between branch *A* and the result of a merge between *B* and *C*. This requirement equates to the traditional definition of an associative operation.

Just like with Bitcoin's double-spending attacks, inconsistencies may occur when one naively attempts to combine the results of different branches by adding their individual states together. However, unlike Bitcoin, which can select a main branch and simply ignore any others, a merging algorithm needs to provide a consistent solution for all conflicts that may arise when merging branches. Because the resolution of such conflicts will be a part of the result of the merging operation, the conflict resolver algorithm has the same requirements as its parent merging algorithm: It must be deterministic, commutative and associative. How this is expressed in a specific blockchain model is left to the designer. An online voting system for example may decide that for conflicting votes from the same voter id only the most recently timestamped is valid, a hash function is applied to select a winner deterministically, or conflicting votes are discarded altogether.

Given these requirements, what would a simple example merging algorithm look like for an alternative blockchain? In the case of a distributed filesystem, we can specify the result *R* of a merging operation between two branches *A* and *B* as follows:

Any files which are present in both *A* and *B* and contain the same content will be present in *R*.

Any newly created files which are present in only one branch will be present in *R*. Any files which are deleted in one or more branches and are not changed in any other branch are removed from *R*.

Conflict resolution: If a file is created or changed in both branches, the file with the larger content will be present in *R*. If the content size is equal, the content is compared and the higher value is chosen as the content in *R*.

This algorithm is deterministic, and by selecting the largest file in either version it is also commutative and associative. Applying this algorithm to a set of clients which share a blockchain representing filesystem snapshots should eventually provide a consistent state across all clients.

In contract, a merging algorithm that fails to meet any of these criteria, for instance by not being deterministic (by performing a random action upon merging), or not being commutative (by always preferring the local changes over remote ones), or not being associative (by not using methods that transitively lead to the same result over larger sets of branches) will relinquish its promise in providing consensus in the network after any amount of merging operations.

Any blockchain that should lead to eventual deterministic consensus among all nodes will need a merging algorithm that is deterministic, commutative and associative. At Storro we use these findings to design many of our merging algorithms with these properties in order to employ them in our proprietary decentralised blockchain models.

Reference:

[1] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System", http://www.Bitcoin.org, 2008

Please contact:

Arthur Melissen, Michel Eppink
Storro B.V., The Netherlands
arthur@storro.com,
michel@storro.com

Blockchain – Attack on and Chance for the Public Sector

by Christian Welzel (Fraunhofer FOKUS)

Until recently, digitalisation in the public sector was characterised largely by making existing processes faster or more efficient. Blockchain instead is challenging established public structures. Currently publically controlled functions for interaction could be organised completely privately, which requires a repositioning of the state. At the same time, Blockchain provides a technological approach that can be used by the public sector itself to improve transparency and trust.

There's been much media hype recently around Blockchain technology. More often than not, the technology is talked about in terms of its potential to revolutionise industry or even the entire internet. Whether or not this is pie in the sky remains to be seen, but there's no doubt that Blockchain technology has underpinned many disruptive ideas in recent years.

The methods behind Blockchain - peer-to-peer networks, hashing, Merkle trees etc. - are not new; it's their combination that results in innovation. Generally speaking, four types of blockchain can be identified, which may be distributed between the two dimensions of read access and write access (see Figure 1). Public blockchains, which can be read by anyone, focus on external effects like transparency or participation. Private blockchains can only be read by a restricted user group and focus on internal effects like process optimisation or collaboration.

Blockchain for the Public Sector

Blockchain is particularly important for the public sector, which is responsible for ensuring social coexistence according to common rules. In many cases the state and the public administration act as an intermediary to regulate and oversee transactions and processes. For this reason many states maintain several registers to manage ownerships: for properties or cars, for instance. In addition notaries guarantee transfer of ownership. In other cases the state takes on a role as trusted third party, which confirms the authenticity of a document or an identity. Thus, a technology that aims to replace third parties by cryptographic functions has direct implications for the state and the public sector in general.

It is therefore not surprising that some countries have been dealing with this technology for many years. In 2007 Estonia established the Keyless Signature Infrastructure (KSI) [1], a technology that uses many aspects of today's Blockchain solutions to ensure the integrity of medical documents. And in 2015 Estonia started its eResidency program that provides a Blockchain-based notary service, e.g., for business contracts or birth certificates.

Depending on the context, Blockchain technology can enable more effective processes, solve partial subproblems or fundamentally change existing work flows. In 2016, the UK Government Office for Science published a study on use-cases, which has received much attention [L2]. A pilot project to pay UK benefits was

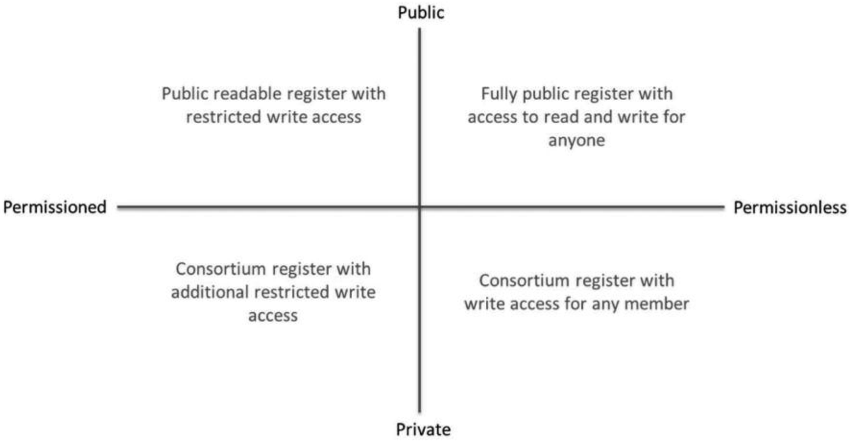
Public
Public readable register with restricted write access
Fully public register with access to read and write for anyone
Permissioned
Permissionless
Consortium register with additional restricted write access
Consortium register with write access for any member
Private

Figure 1: Types of Blockchain categorised according to read and write access [L1].

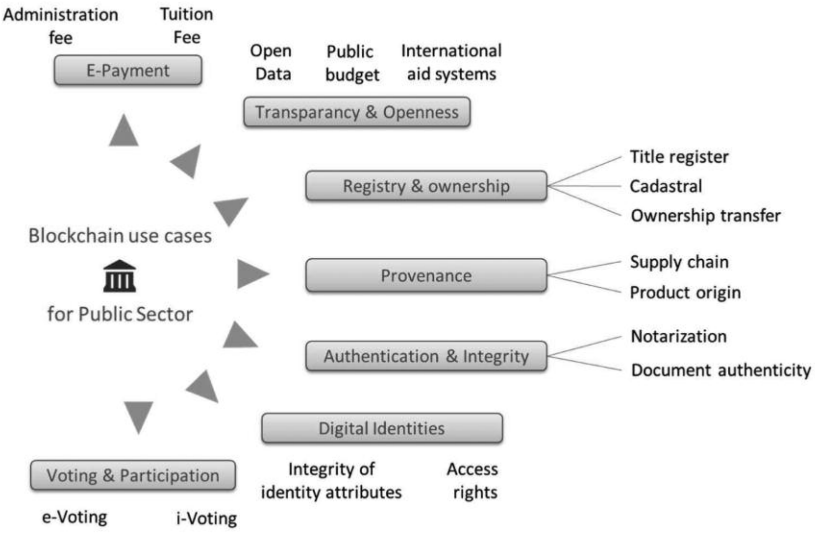

Figure 2: Common discussed and tested use-cases for Blockchain technology in the public sector globally [L1].

launched in the same year but also sparked privacy concerns.

Many countries, including Switzerland, USA, Sweden, Ghana and Georgia, are experimenting with Blockchain technology. Dubai announced a Blockchain strategy with the goal of processing all government documents digitally via a Blockchain infrastructure by 2020. The motivation to use Blockchain technology ranges from the prevention of corruption or misuse, to improved transparency (see Figure 2).

The Role of the State

It is important that governments take an interest in Blockchain technology. On the one hand, the state acts as a regulator, e.g., for cryptocurrencies like Bitcoin. The specific challenge for regulation lies in the decentralised nature of Blockchain. If responsibilities are missing, traditional regulation mechanisms seem to reach their limits. For this reason, some argue that Blockchain is resistant to regulation.

Aside from the issue of regulation, the main question for states and policy makers will be how to deal with the technology in practical terms. The range of possibilities extends from prohibition, through to tolerance, acceptance, allocation of public funding for its development or even application.

Four Recommendations:

1. *Monitor Blockchain technology and identify regulatory needs.*
   Since national regulation has only a limited impact here, a European or better an international debate on the topic is necessary, as with other digitisation questions.
2. *Apply Blockchain technology and develop best practice examples.*
   Experience can only be gained by testing. The public administration should therefore analyse its own processes and consider how Blockchain technology might be implemented within them.
3. *Push standardisation forward.*
   Currently, the development of Blockchain is characterised by proprietary interfaces. Standardisation, which is urgently needed, has just begun.
4. *Actively shape the further development of Blockchain technology.*
   There are a number of open questions about the ethical and societal implications of Blockchain technology. Governments should allocate funding to research in this area.

Links:

[L1] http://www.oeffentliche-it.de/publikationen

[L2] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf

Reference:

[1] Buldas A., et. al.: “Keyless Signatures Infrastructure: How to Build Global Distributed Hash-Trees.” In: Riis Nielson H., Gollmann D. (eds) Secure IT Systems. NordSec 2013.

Please contact:

Christian Welzel, Competence Center Public IT, Fraunhofer FOKUS, Germany
+49 30 3463 7377
christian.welzel@fokus.fraunhofer.de
Twitter: @OeffentlicheIT

How Distributed Ledgers Can Transform Healthcare Applications

by Angelica Lo Duca, Clara Bacciu, Andrea Marchetti (IIT-CNR)

The use of distributed ledgers (DL) in healthcare applications presents the opportunity to create new interoperable and secure environments, from which both medical professionals and patients can benefit.

Within the Institute of Informatics and Telematics of the National Research Council in Pisa, a new working group on distributed ledgers (DL) is being set up. The main objective of the group is to study DL technology and implement DL-based solutions for different scenarios, such as traceability of products and health care applications (HCAs). In this paper we focus on HCAs and we illustrate how they can be transformed with the introduction of DL.

For many years, electronic health records (EHRs) about patients’ health care have been stored in different HCAs, which essentially are very heterogeneous and are not designed to manage multi-institutional and lifetime medical records (Figure 1). In fact, different health institutions do not usually share a common HCA, and EHRs associated to the same patient and stored by different institutions do not refer each other. This means that it is not possible to extract a patient’s history simply from their EHR. In some cases, even within the same institution, two EHR associated to the same patient are not connected. Thus, a patient moving from one institution to another, must register each time to the local HCA, provided by the current institution, in order to build his/her local EHR. In addition, the patient should provide the current institution with the previous resignation sheet, which contains all the information related to the patient’s previous recovery. This process, which often is tedious, could also introduce errors and lose potentially useful information about the patient. As a consequence, medical professionals may end up working with incomplete information.

DL can potentially transform HCAs as they provide new opportunities for health IT systems, by guaranteeing interoperability among heterogeneous applications in a secure way [1]. In fact, an architecture based on DL does not require node homogeneity: the only requirement is that a node of the system can be identified in a secure way, i.e., it is equipped with a pair public/private key. Security provided by DL through authentication, confidentiality and accountability, allows patients’ information to be protected from unauthorised access. Both doctors and patients could benefit from the use of DL in HCAs: on the one hand, doctors could access more information about a patient, such as their health history, laboratory results, prescribed medications, and even information extracted from personal bio sensors. On the other hand, it would also allow patients to access their full EHR, anywhere and anytime and it would remove the need for individuals to register to a new HCA for each hospitalisation. Furthmore, the whole health sector may benefit from the adoption of DL, through economic savings and increased efficiency of the health IT system.

A variety of companies have already entered the ecosystem, trying to capitalise on opportunities created by DL applied to HCAs [2, L1, L2, L3]. All the existing models in this field separate databases where all EHRs are stored from the database which specifies how to access the data, i.e., the DL [3]. Figure 2 illustrates how DLs can be used in HCAs: a data lake is used to store all the patients' information (laboratory results, doctors' prescriptions and so on). The DL is exploited to store all the pointers to the data lake in a secure way. Patients and providers (e.g., institutions, doctors, laboratories) are equipped with public/private keys, which allow them to authenticate each other. Patients can give partial and time-limited view of their EHRs to one or more providers through access control policies, regulated, for example, by means of smart contracts. Once the medical exam is ready, the provider signs it with a digital signature and sends it back to the client. If the client accepts it, the exam is added to the DL.

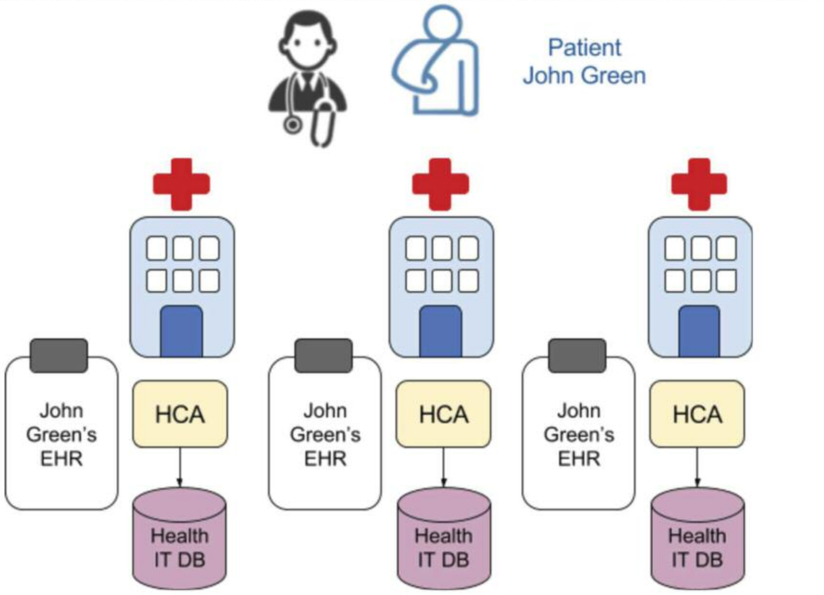

Figure 1: Traditional management of EHRs.

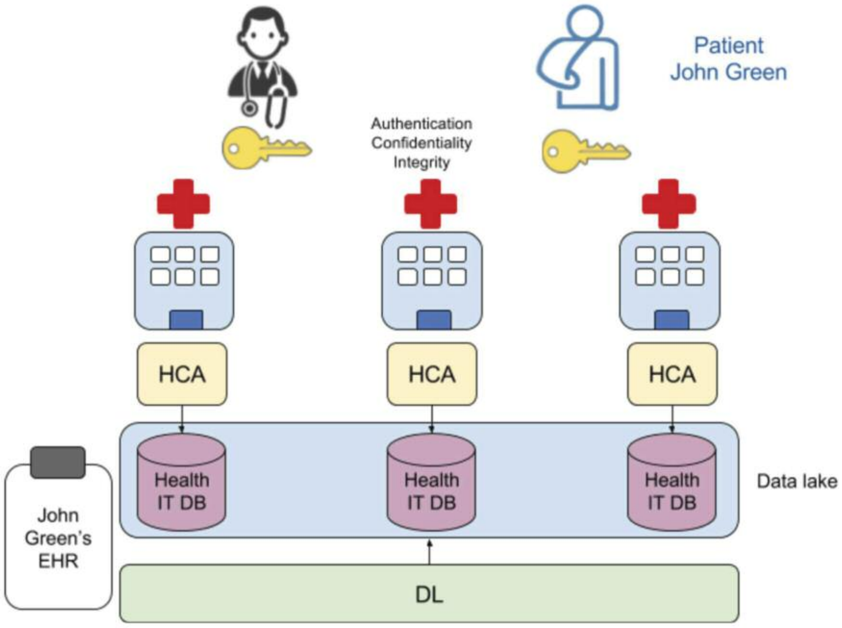

Figure 2: How the use of DL can transform the management of EHRs.

DLs undoubtedly represent a real opportunity for HCAs because they do not require any changes to the existing health IT databases, which are still maintained separated. Only a new database, i.e., the DL, should be added to the network, containing only the pointers to the original health IT databases. The DL would play two roles: to connect all the existing health IT databases and to guarantee privacy and integrity of EHRs.

Links:

[L1] https://deepmind.com/applied/deepmind-health/
[L2] https://gem.co/health/
[L3] https://blockchainhealth.co

References:

[1] M. Mettler: "Blockchain technology in healthcare: The revolution starts here", in e-Health Networking, Applications and Services, Healthcom 2016, pp. 1-3, IEEE, 2016.

[2] A. Ekblaw et al.: "A Case Study for Blockchain in Healthcare: 'MedRec' prototype for electronic health records and medical research data", Whitepaper, 2016. http://dci.mit.edu/assets/papers/eckblaw.pdf

[3] L. A. Linn, M. B. KOO: "Blockchain for Health Data and Its Potential Use in Health IT and Health Care Related Research" https://www.healthit.gov/sites/default/files/11-74-ablockchainforhealthcare.pdf

Please contact:

Andrea Marchetti
IIT-CNR, Italy
+39 050 315 2649
andrea.marchetti@iit.cnr.it

Blockchain-enabled Intelligent Asset Exchange for a Circular Economy

by Ioannis Askoxylakis (FORTH), George Alexandris (Bournemouth University) and Giorgos Demetriou (Ecole des Ponts Business School)

As the notion of circular economy gains momentum, intelligent assets increasingly form the backbone of sustainable ecosystems. Although these assets can supply the necessary knowledge for fueling the value drivers of a circular economy, the generated value will be significantly amplified by allowing third parties to manage them and profit from better asset utilisation. However, for an ever-changing networked environment consisting of numerous assets, ownership needs to be dynamic, granular and adaptable in order to maximise gains. Blockchain-based mechanisms can effectively serve this need by enabling transfer of asset ownership directly between parties participating in the circular economy while introducing trust, efficiency and automation in asset exchange contracts.

The term "circular economy" refers to an economy that is restorative and regenerative by design, aiming to keep resources at their highest utility and value at all times. Its value drivers include extending the useful life of finite resources, maximising the utilisation of assets and creating new use cycles for end-of-life assets. Simultaneously, assets leverage advances in IoT, thus creating an emerging class of "intelligent assets" [1] governed by three underlying attributes enabling circularity: location, condition and availability. These properties sufficiently describe the state of the asset for operating under a certain role in the ecosystem, but at the same time raise two fundamental questions:

- How do these properties affect the value generated by the asset?
- Who defines the role(s) of the asset?

Given that a value of an object is usually tightly coupled with its role in a system, these questions cannot be answered independently. Furthermore, in a complex system like a smart city, it is unlikely that a single possible answer for every asset exists. We thus introduce the idea of an exchange, where different entities can acquire a stake in an asset and operate it for their own profit. The entity operating the asset can define the asset's role (within certain boundaries set by a global system owner) according to what it judges to be in its best interest. This leads us to two crucial concepts: asset ownership and asset control.

Transitioning from "Intelligent Asset" to "Intelligent Property"

Characterising an asset as a property, requires the owning party to (a) provide a universally accepted testimony of the asset's ownership and (b) to have exclusive access to the asset itself. Both requirements can be achieved by registering the asset as a digital asset on the blockchain. Furthermore, all historical location, condition and availability properties which affect the asset's value will be visible on the blockchain and signed by the respective asset owner at that point in time, thus guaranteeing that the data is trustworthy. Properties can be further augmented with asset data of particular interest to prospective buyers such as damages, alterations, repairs etc. Access to assets can be granted by com-

Real-world action	System equivalent
Municipality commissions a new smart asset, e.g., a smart lamppost	Municipality registers a new asset, e.g., Lamppost_123 on the blockchain together with a snapshot of all its relevant properties for circularity. Municipality signs the transaction with its private key.
Municipality assigns the maintenance of the asset to Party A	Municipality adds a transaction to the blockchain with a snapshot of Lamppost_123s properties and stating that Party A is now the owner of the asset. Municipality signs the transaction with its private key
Lamppost_123 acknowledges Party A as its new owner	Lamppost_123 receives a pre-defined number of confirmation blocks of the blockchain stating that Party A is the new owner
Party A accesses Lamppost_123	Party A generates an access token for Lamppost_123 The token is signed with Party A's private key and presented to Lamppost_123 Lamppost_123verifies the token with Party A's public key and verifies that Party A is still its owner by checking the latest blockchain copy
Party A transfers ownership of Lamppost_123 to Party B	Party A adds a transaction to the blockchain with a snapshot of Lamppost_123s properties and stating that Party B is now the owner of the asset. Party A signs the transaction with its private key

Table 1: Sample sequence of actions for an intelligent asset blockchain.

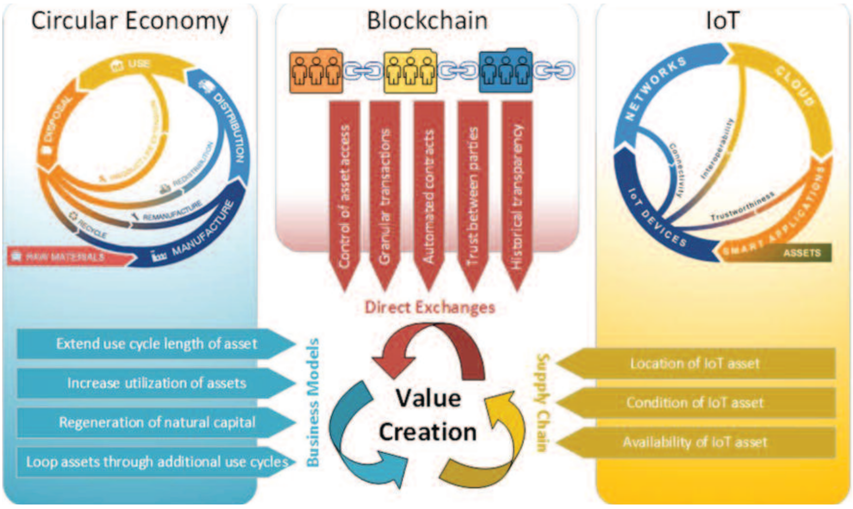

Figure 1: Interplay of value drivers (left-side circular economy diagram adopted from [2]).

bining the blockchain with owner-signed access tokens as described in Table 1.

Blueprint for a blockchain implementation for exchanging intelligent assets in smart cities

In the case of a smart city, the global system owner is the city's municipality which has ultimate authority over all assets and is also responsible for setting the rules of the blockchain. The municipality, its intelligent assets and all potential owners of assets form a blockchain network. In this context, asset owner means the party which is responsible for and has a stake in operating the asset. An asset can either be obtained for operation/maintenance on behalf of the municipality (i.e., the municipality pays the owning party for services provided), or it can be leased from the municipality (i.e., the owning party pays the municipality a rent for the asset). An owning party controls the asset and is entitled to transfer its ownership to another party via a contract. The parties can also specify rules governing the contract which will be enforced automatically by the asset itself in the future under certain conditions. It should be noted that the actual payment of funds between parties is independent of the asset transfer, and may or may not be part of the same blockchain.

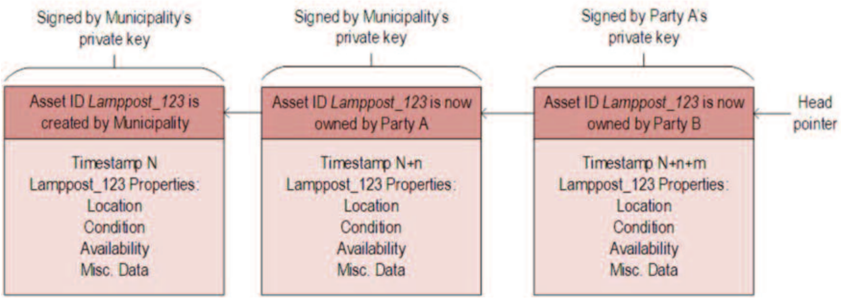
Signed by Municipality's private key
Signed by Municipality's private key
Signed by Party A's private key
Asset ID Lamppost_123 is created by Municipality
Asset ID Lamppost_123 is now owned by Party A
Asset ID Lamppost_123 is now owned by Party B
Head pointer
Timestamp N
Lamppost_123 Properties:
Location
Condition
Availability
Misc. Data
Timestamp N+n
Lamppost_123 Properties:
Location
Condition
Availability
Misc. Data
Timestamp N+n+m
Lamppost_123 Properties:
Location
Condition
Availability
Misc. Data

Figure 2: Simplifed blockchain.

Figure 2 shows the resulting simplified blockchain resulting from the above scenario.

The proposed design is part of ongoing research taking place in the context of the Horizon 2020 project CyberSure [3] that is coordinated by Computer Emergency Response Team of the Foundation for Research and Technology-Hellas, in collaboration with the Cybersecurity Research Centre of Bournemouth University and the Circular Economy Research Centre of Ecole de Ponts Business School.

References:

[1] Intelligent Assets: unlocking the circular economy potential, Ellen MacArthour Foundation, 2016
[2] https://kwz.me/Xx
[3] http://www.cybersure.eu/

Please contact:

Ioannis Askoxylakis, FORTH, Greece
+30 2810 391723
asko@ics.forth.gr
www.forthcert.gr

Blockchain and Autonomous Institutions

by Mariusz Nowostawski (NTNU)

Blockchain technology is relatively young, although the underlying cryptographic mechanisms have been known in computer science for some time. At NTNU we are interested in questions such as "Will blockchain and smart contracts make society better? Does this mean that we must have trust in the code, and not in humans? Could Blockchain technology be used as a powerful cyber weapon?"

The main focus of Bitcoin is to provide a global store of value, a state of account and a medium of exchange. In other words, the primary objective of Bitcoin-like systems is to offer a currency. In contrast, the main focus of Ethereum is to provide a global decentralised computational fabric that can facilitate interactions between humans, and algorithms expressed as non-mutable, verifiable code.

Blockchain technology has inspired many social scientists and economists. Technology enthusiasts consider the underlying mechanisms of social organisation, governance, finance and law. Researchers started investigating the implications of global, decentralised and non-FIAT monetary systems. For example, some researchers argue that automation and decentralisation may make the concept of the state entirely obsolete.

Blockchain technology, or more broadly, "distributed ledger technology", enables many novel decentralised applications. Ranging from simple digital tokens representing currency, through to digital assets management, and audit trails, to the ability to establish decentralised institutions [1]. The property of being decentralised means that a trusted third party is no longer required in many scenarios that would previously have required one. This has far-reaching social implications.

The Blockchain research group at NTNU explores the rich and evolving space of digital currencies. There is currently a lively marketplace for over 2,000 digital currencies in addition to Bitcoin, even though most people have probably never heard of any of them. Tracking, anonymization, and forensic readiness of cryptocurrencies are the main collaboration points between NTNU and Norwegian and European law enforcement agencies. The research group in NTNU investigates several issues related to the current state of the technology. Among them, the researchers consider scalability and resilience of off-chain transactions in the Lightning Network. Lightning Network is the proposed decentralised payment solution that would enable unconstrained scaling of the transactions in the system. Other areas of inter-

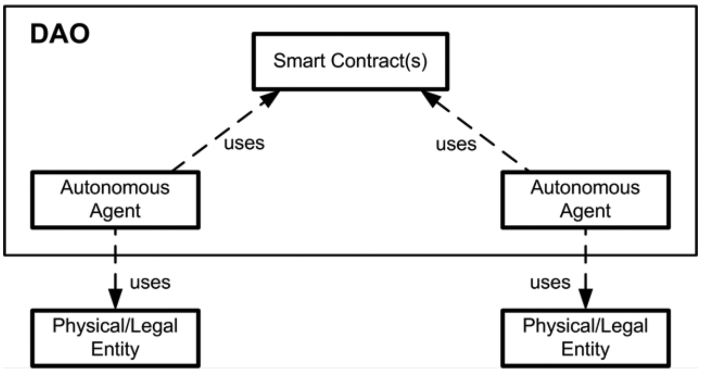

Figure 1: Next generation of institutions will integrate humans, AI subsystems, Decentralized Autonomous Organisation (DAO) and smart contracts, to provide automated, verifiable and efficient workflows.

est are in the value chain and in supply chain management. The lead in this area belongs to San Francisco-based startup SKUchain [2], which has provided a few research projects to collaborators in NTNU.

Research on ensuring consistency of computations conducted by a distributed network of nodes is important and often considered the most urgent problem in the field. Multi-party computations, order-preserving hashing and automated verifiable code generations are all active research areas.

Advocates of private blockchains say that a private blockchain is to a public blockchain what the intranet is to the internet. But this analogy is misleading: the internet is qualitatively, as well as quantitatively, different to an intranet, which serves a completely different purpose. Public blockchains are qualitatively different to permission-based blockchain deployments. Instagrams, Facebooks and Googles cannot happen on intranets – they can only happen on the Internet. Similarly, there are things that can only happen through open public blockchains. Bitcoin does not make sense as a private blockchain experiment. However, state-owned currencies make perfect sense as private blockchains. So, what's the difference between a private and a public blockchain?

Nobody truly understands yet what the consequences of blockchain technology will be. The blockchain innovation is based on openness, different trust models, and new ways of control. Independent, non-trusting parties can interact with each other through an open, trusted, and verifiable communication and value-exchange fabric.

Private, permission-based blockchains are about new ways of control; new ways of doing old things, where one group extracts value out of another. A state-issued digital currency will use private, state-owned blockchain, and it will be possible to track who transacts with whom and when. Unlike cash, private blockchain technology will allow the state to have full knowledge about the way money and value flows in the economy, down to the single penny. With a flick of a switch, criminal wallets can be disabled, and funds ceased and burned. This technology is like a weapon and can be used to restrict personal freedom and to increase control - or to the contrary, it can be used to enhance personal freedom.

Blockchain technology has the potential to have a significant impact on society and it is of paramount important to understand how it can be used and misused. It is a matter of trust!

References:

[1] M. Nowostawski, C. K. Frantz: “Blockchain: The emergence of autonomous decentralised Institutions”, SOTICS 2016. http://kwz.me/XB

[2] Skuchain Brackets - Blockchain Technology for Collaborative Commerce. https://www.skuchain.com/

[3] C. K. Frantz, M. Nowostawski: “From institutions to code: Towards automated generation of smart contracts”, FAS*W, 2016. DOI: 10.1109/FAS-W.2016.53

Please contact:

Mariusz Nowostawski
NTNU, Norway
mariusz.nowostawski@ntnu.no

Self-Sovereign Identity Framework and Blockchain

by Rieks Joosten (TNO)

The ability to use identities in many different digital contexts is vital for doing electronic business transactions. Such identities are hard to come by, in particular when the transaction involves international parties that do not necessarily trust each other (yet). The Dutch Techruption project has taken on the challenge of specifying a self-sovereign identity framework (SSIF) that aims to solve this problem, and to build demonstrators that show its practical use, for businesses, consumers and governments. Blockchain technologies are used for critical parts, such as storing commitments to attestations and revocation events.

The Techruption Blockchain Project is a public-private partnership project in the Netherlands, within which large corporates, small companies, startups and scientific institutions collectively create disruptive technological innovations around distributed ledger (blockchain) technologies (DLT). DLTs are particularly useful in business and governance situations that involve multiple parties that do not necessarily trust one another to negotiate and execute electronic business transactions. In many cases such transactions require the ability to establish and validate identities and identity attributes, or to check whether or not they have been revoked.

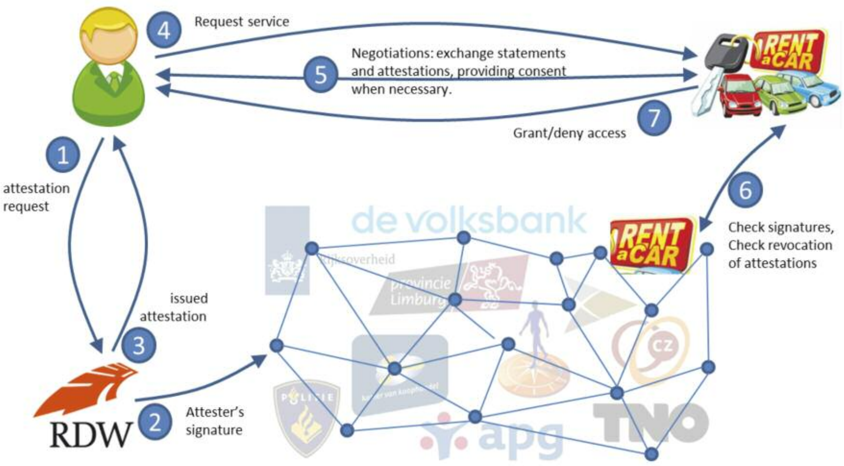
4
Request service
5
Negotiations: exchange statements and attestations, providing consent when necessary.
7
Grant/deny access
1
attestation request
issued attestation
3
RDW
2
Attester's signature
6
Check signatures, Check revocation of attestations
RENT aCAR
de volksbank
TNO
apg

Figure 1: Techruption Blockchain Project.

Seven participants of the project (Accenture, APG, Brightlands, Chamber of Commerce, De Volksbank, Rabobank, and TNO) are developing a self-sovereign identity framework (SSIF) for the creation, validation and revocation of such identities that can be used in conjunction with blockchain technologies and the (disruptive) applications that are enabled by such technologies. The goal is to specify, validate and ultimately build a trustworthy, open digital infrastructure for self-sovereign identities that is secure, decentralized, open source, supports privacy (e.g., GDPR compliance) in multiple roles, and lacks a single point of failure or large information honey-pot. We aim to follow well-established requirements for user-centric identity systems [1], [L1].

The SSIF has a terminology and method (based on the DEMO models for (business) transactions [2] and the Networked Risk Management model [L2]) that a business party can use to specify all information it needs to construct a valid argument for deciding whether or not it should commit to a proposed business transaction.

One prerequisite for an argument to be valid is that the meaning (semantics) of every statement must be defined. Using semantic web models (e.g., RDF(S)) for mapping statement-representations to their corresponding meaning allows the use of semantic business standards such as UBL, open data, data from “things” (from IoT frameworks) as well as personal data.

The other prerequisite for valid arguments is that the truth of all statements must be established. The SSIF assumes that the “truth” of a statement is subjective, i.e., a decision by the party that uses the statement in an argument. One of the most important concepts in the SSIF is the “attestation”, i.e., a statement (with well-defined semantics) that is signed by some party (the “attestor”) that attests to the truth of other statements.

The value of attestations is that parties that trust them (or the attestors) can simplify processes, such as the onboarding process of banks that are heavily burdened by KYC regulations. We believe there is new business to be found in the issuance of solid attestations, and that smart contracts can be designed that facilitate such businesses.

Another crucial role of DLTs is to register events by which attestations are revoked. It is easy to envisage the benefits of knowing whether or not a person is still an employee of a company, if a passport or driving licence has been revoked, a father still has parental authority, or a certificate is still valid. The distributed nature of DLTs allow parties to query a single (but redundant) endpoint for revocations of attestations by all parties, rather than having to query a specific endpoint for individual parties.

Several ideas still need to be developed. One such idea is the SSIF wallet, which is seen as an embodiment of the “self-sovereignty” aspect of the framework. It is envisaged as a container for statements and attestations that can act as a proxy of its owner in the negotiation and execution of electronic business transactions. This means that the proliferation of statements and attestations is controlled by the user, and their use is subject to the user’s consent. While the idea itself is not new (we have seen InfoCards, attribute-based credentials, etc.), the interfacing with and use of BLTs is.

Ultimately, the project will provide a solid framework for self-sovereign identity that can be used in combination with DLTs, with a firm conceptual underpinning, that reuses existing technologies, is easy to install and maintain from a business perspective, is an enabler for disruptive business ideas, and has at least one working prototype to demonstrate its viability.

Links:

[L1] https://kwz.me/Xv
[L2] https://kwz.me/Xw

References:

[1] K. Cameron, R. Posch, and K. Rannenberg: “Appendix d. proposal for a common identity framework: A user-centric identity metasystem”, The Future of Identity in the Information Society (2009): 477.

[2] J. Dietz: “Understanding and Modelling Business Processes with DEMO”, LNCS, vol 1728, 1999.

Please contact:

Rieks Joosten
TNO, the Netherlands
+31622901317
rieks.joosten@tno.nl

Distributed Access Control Through Blockchain Technology

by Damiano Di Francesco Maesa, Laura Ricci (Università di Pisa) and Paolo Mori (IIT-CNR)

We defined a distributed access control system on top of blockchain technology. The underlying idea is to properly represent the access rights of the subjects in the blockchain in order to easily allow their enforcement at access request time. By leveraging blockchain advantages we can add new desired properties, such as auditability, to the access control system. To prove the feasibility and validate the proposed approach we developed a proof of concept implementation and performed some relevant experiments.

A blockchain is a distributed, always available, irreversible, and tamper-resistant public database where the control over data and its evolution is distributed among a variable set of peers. Blockchain technology does not require the existence of trust relationships among the system's users. Consequently, it employs a distributed consensus algorithm to allow the users to agree on immutable and auditable append-only operation without requiring interaction with a trusted third party. We are interested in the auditability of the data stored in the blockchain, since the blockchain can be used as a publicly verifiable proof that the data existed at the time it was saved in it.

Access control systems are meant to regulate the access to critical or valuable resources. Several access control models, i.e., ways of defining the policies expressing the rights of subjects to access resources, have been defined, and here we focus on attribute-based access control (ABAC) policies. An ABAC policy combines a set of rules expressing conditions over a set of attributes paired to the subject, to the resource or to the environment. The rules must be satisfied accordingly in order for the access right to be granted. A well-known policy language to express ABAC policies is the eXtensible Access Control Markup Language (XACML), defined by the OASIS consortium.

Our proposal exploits blockchain technology as the base framework on top of which we build an ABAC system. The first step for defining our distributed blockchain-based access control system is to store the access control policies in the blockchain. Depending on the underlying blockchain, different technical solutions can be adopted. If the blockchain allows for arbitrary data storage, then we can save the policy directly on it. On the contrary, if the underlying blockchain has strict space constraints, e.g., Bitcoin [1], we should adopt more complex solutions, such as storing links to the policies in the chain while the complete policies are stored elsewhere (e.g., in distributed hash tables (DHT)). This is possible as long as the storage system remains tamper-proof and guarantees data availability, and the linking to the policies is unique and tamper-proof as well. Another option is the one provided by the Ethereum blockchain [2], which allows smart contracts to be represented and run. In this case, these smart contracts can be exploited to properly encode the policies themselves in executable format.

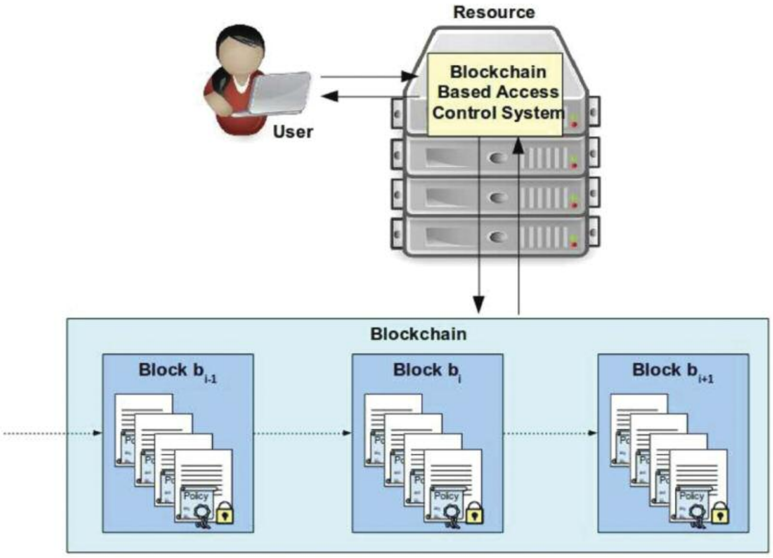

Figure 1: Architecture of the blockchain based access control framework.

The next step of our approach is to define the policy enforcement architecture, i.e., to define the set of components required to perform, at access request time, the policy evaluation against the current access context. If the policy is not stored in the blockchain in executable format, the architecture of the enforcement system is similar to the XACML reference one [1].

As an example, in [3] we described a preliminary prototype of a blockchain-based access control system which exploits the Bitcoin blockchain and XACML policies. In this case, since Bitcoin was not designed to store arbitrary data, we defined a customised strategy to compress the XACML policies and we exploited the OP_RETURN script op code and MULTISIG transactions to store them in the chain. The attributes required at access request time to perform the policy evaluation can be retrieved from traditional attribute providers (e.g., Lightweight Directory Access Protocol (LDAP) services). However, we envisage that attributes could be stored and managed exploiting the blockchain as well. If the policy is stored in the blockchain in executable format, i.e., through smart contracts, most of the policy enforcement architecture is embedded in the blockchain itself. Such smart contracts represent self-evaluating policies that can be queried directly and transparently at access request time. The attributes required for the evaluation of the policy are encoded in the blockchain as smart contract as well. We have developed a proof-of-concept implementation of this approach on the Ethereum blockchain, demonstrating the feasibility of our proposal.

The evaluation of a blockchain-based access control policy could be performed by a party which is not trusted by the resource owner or by subject of the request who, instead, would like to be guaranteed against malicious or erroneous policy evaluations. For instance, the party that evaluates the policy and enforces the result could maliciously force the system to deny an access although the policy would have granted it.

Blockchain technology can be exploited to address this problem as well, In fact, having the policies and the attributes publicly available through the blockchain, allows any user to know at any time the policies that are applicable to its access request and the related access context. This allows distributed auditability, detecting parties that fraudulently alter the rights granted by the enforceable policies.

References:

[1] S. Nakamoto: “Bitcoin: A peer-to-peer electronic cash system”, http://bitcoin.org/bitcoin.pdf (2008).

[2] G. Wood: “Ethereum: A secure decentralised generalised transaction ledger”, Ethereum Project Yellow Paper 151 (2014).

[3] D. Di Francesco Maesa, P. Mori, L. Ricci : “Blockchain Based Access Control”, Proc. of the 17th IFIP International Conference on Distributed Applications and Interoperable Systems (2017).

Please contact:
Damiano Di Francesco Maesa,
Dipartimento di Informatica,
Università di Pisa, Italy
damiano.difrancescomaesa@for.unipi.it

Blockchain Ensures Transparency in Personal Data Usage: Being Ready for the New EU General Data Protection Regulation

by Uwe Roth (Luxembourg Institute of Science and Technology, LIST)

The new EU General Data Protection Regulation (GDPR) [1][L1], which will come into effect in 2018, demands transparency as one of the main principles for the collection, processing, storing and transfer of personal data. Transparency ensures that individuals can enforce their legal rights: to withdraw consent for their personal data to be processed or to request that their data are erased. At the Luxembourg Institute of Science and Technology we filed a patent, based on blockchain technology, that guarantees transparency in the context of files that are exchanged in a shared data pool. It guarantees that access by partners to specific files can be traced without a central entity.

Closed consortia who provide and exchange personal data between their partners need to understand the impact that the new General Data Protection Regulation (GDPR) of the European Union will have on their processes. In fact, it demands that the data controller must provide to an individual, upon request, information about the transfer of their personal data to third parties, third countries or international organisations.

The legal counterpart and single point of contact for a person that provided their personal data will not be the consortium as a whole, but a single data controller partner who has collected the private data and published it in the data pool. This contractual partner needs to ensure that the processing of personal data is legal, and may be sued if it fails to fulfil its legal obligations.

In the absence of a central logging facility that acts as a trusted third party, and without sophisticated access to appropriate management solutions, it can be virtually impossible to trace which partner organisations have accessed individual data sets. It can, therefore, be extremely difficult to provide relevant information to an individual who requests that their data be deleted. As a consequence, the processing of personal data inside a consortium might, in the future, be criminalised or result in a fine.

In 2016, The Luxembourg Institute of Science and Technology, LIST [L2], filed a patent application that addresses these new demands with a solution based on the latest file-distribution, blockchain and encryption technologies. In a purely decentralised peer-to-peer environment of equal partners, without requiring any centralised instance or further authorisation steps, the solution empowers the provider of data to trace the access to the data by partners in the distributed and shared data pool. This trace of access is without any doubt and cannot be denied.

The solution is based on a file distribution network of encrypted files. It creates redundancy to increase availability and to improve download speed (Figure 1). The blockchain network is used to log the access to the files which will allow every access of the data by any

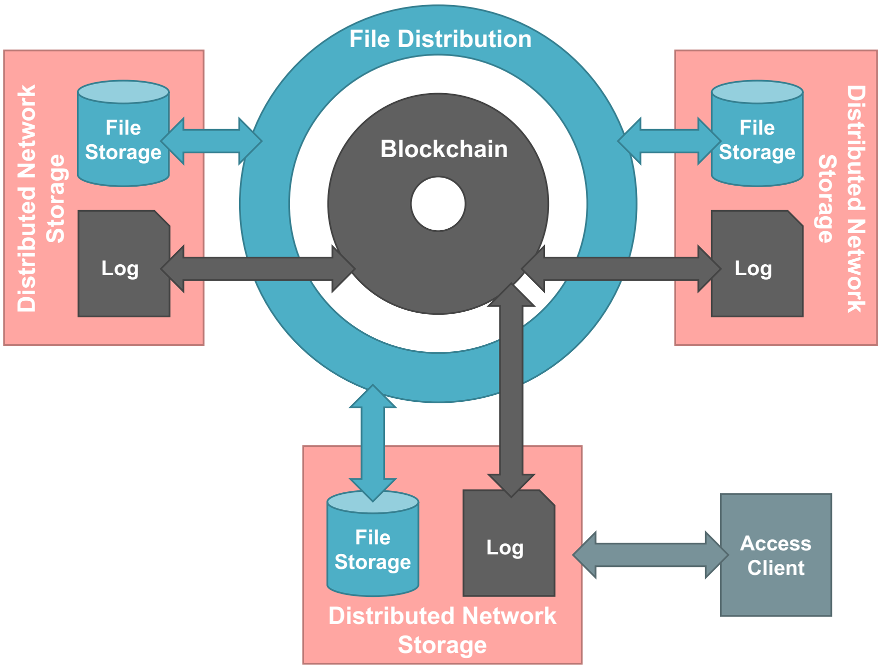

Figure 1: Clients access their local network storage that acts as a node in a decentralised peer-to-peer network.

partner in the network to be traced. In our solution, we abstract from the use of the underlying blockchain technology (Figure 2) in an ISO-OSI [2] layered way: Starting from the lowest blockchain-network layer that is only used to broadcast a number of bytes amongst the network, each new layer on top of this adds another level of complexity, e.g., adding point-to-point communication, adding end-to-end-encryption, or adding messaging. This allows us to make our approach independent from the specific blockchain implementation that is finally used, and replace it without the need to re-implement the entire solution.

We are currently in the process of developing a first proof-of-concept that includes all the elements described in the patent, and, at the same time, searching for partners that can provide use-cases that are relevant for specific markets. For example, an clinical research consortium that exchanges genomics data between their partners, which may be international. An alternative use-case is the tracing of access, not only to personal data inside a shared data pool, but also to copyright protected data that triggers payment obligations. In addition to logging access to files, we are currently investigating an approach to log access to any type of service, including data inside a distributed database.

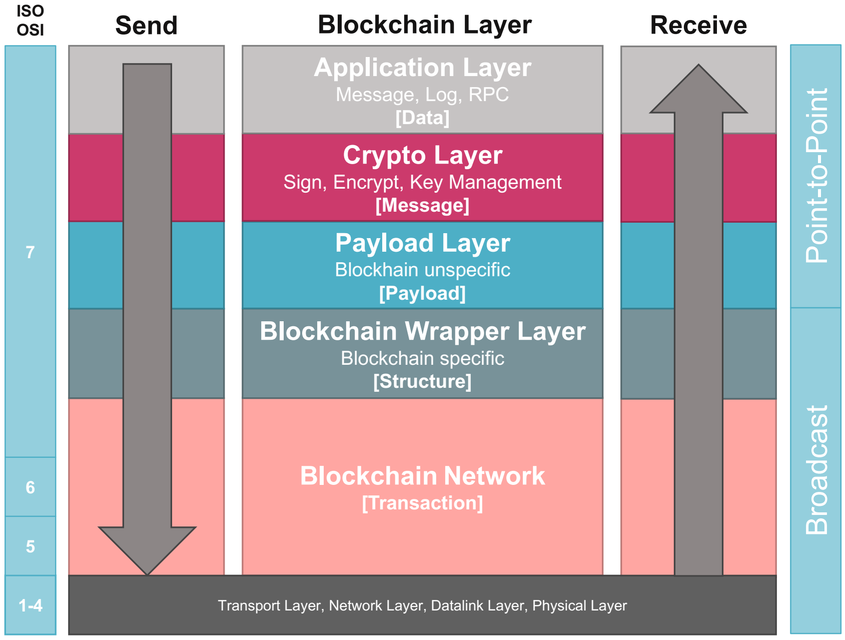

Figure 2: Usage of the blockchain in a layered architecture, in relation to the ISO-OSI model.

Links:
[L1] http://www.eugdpr.org
[L2] http://www.list.lu

References:

[1] GDPR, Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), online: http://data.europa.eu/eli/reg/2016/679/oj

[2] H. Zimmermann: "OSI Reference Model – The ISO Model of Architecture for Open Systems Interconnection", IEEE Transactions on Communications, vol. 28, no. 4, April 1980, pp. 425–432.

Please contact:

Uwe Roth
Luxembourg Institute of Science and Technology, LIST
uwe.roth@list.lu

Flexible Transparency: A Privacy Enabler in Blockchain Technologies

by Maria Christofi (Trusted Labs) and Aline Gouget (Gemalto)

Use blockchain technologies while keeping the control of the transparency you provide!

An important property of blockchain technologies is that trust is distributed between all blockchain network members. When anyone can be a member of this network, it is about a permissionless model, whereas when access to the network is allowed only to entities that belong to a defined consortium, it is a permissioned model. We focus here on one key topic for blockchain technologies: the sharing of trust beyond the choice of the consensus protocol.

One main characteristic of blockchain technology is that it enables maintenance of a continuously-growing list of ordered records, called blocks, on which there is a consensus. Each block is timestamped, linked to previous blocks using cryptography and contains one or several "transactions" that have been "verified" by the blockchain members. "Transaction" should be understood in a broad sense to be adapted to many use-cases. For example, a Bitcoin transaction is related to the digital money transfer using pseudonyms, but a transaction can also be a simple exchange of assets, described into a smart contract, or it can be almost "everything", e.g., a hash of data for internal logs or proof of anteriority [1]. Then, depending on the use-case, data fields have to be categorised, i.e., data fields publicly shared, data fields shared only within the blockchain consortium, data fields shared only with identified entities and data fields that have to be kept confidential (and not shared with anyone). The meaning of "verified" must also be taken in a very broad sense, as illustrated by the three cases below.

Data protection can have different meanings depending on the use-case, e.g., authentication of the data source, authentication of people accessing it, securitisation of the data transport or of the storage. Moreover, different notions of privacy properties may have to be supported depending on the use-case and on the data fields' classification.

Let's have a look at three use-cases that illustrate the shades of "transparency" that can be provided by blockchain technologies.

Case 1: Bitcoin transaction.

Data is shared publicly in a permissionless ledger. This is an example with almost maximal transparency by a full delegation of the verification steps to the blockchain network. Indeed, all the steps of the transaction verification (i.e., amount/balance verification, signature authenticity verification, detection of double-spending) as well as the insertion into the ledger are done by the blockchain network. Regarding privacy properties, only the pseudonymity of payer/payee is claimed to be supported. However, this privacy property is quite weak and it is possible to extract some information by analysing the graph of transactions, e.g., [2]. For Case 1, the ledger is auditable with independently-verifiable proof-of-time, at any time by anybody.

Case 2: Smart contracts with confidential formula based on zero-knowledge proof of knowledge techniques, e.g., [3].

This use-case can be relevant for both permissionless and permissioned networks. In both models, only the stakeholders involved in the smart contract have access to the exact formula that will be used to compute the final amount of the transaction. The rules of the smart contract are shared inside the blockchain network but the formula that will be used to compute the final amount of the contract is kept confidential; this formula is embedded into the smart contract in an encrypted form. Then, when the smart contract is executed, the payer and/or the payee reveal the final amount of the transaction while proving that this amount is compliant with the encrypted formula contained within the smart contract, without disclosing the formula. This privacy feature is important, especially for business relationships where the rules may change from one customer to another. All these features could be formally verified to increase the assurance and trust in transaction verification. For Case 2, the ledger is also auditable at any time by any member of the blockchain network.

Case 3: Records of private financial transactions by disclosing only hash of the transaction data.

This use-case is more suitable for use with a permissioned model where the blockchain network can take responsibility for authenticating the data source. Then, it is guaranteed that only data provided by authenticated sources have been added to the permissioned ledger while transaction data remain confidential (except for the stakeholders involved in the transaction). However, financial data could be shared a posteriori, e.g., to a judge to enable transaction verifications and anchoring into the ledger.

These use-cases illustrate that it is not necessary to choose between full worldwide transparency and transparency only within a consortium, but several levels of transparency can be considered depending on the use-cases needs and especially concerning the classification of data fields, possibly using a hierarchy of consortia. The aim of our work is to provide recommendations for designing customised transparency depending on the use-case.

References:

[1] A. Buldas, A. Kroonmaa and R. Laanoja: "Keyless Signatures' Infrastructure: How to Build Global Distributed Hash-Trees", Cryptology ePrint Archive, Report 2013/834, 2013.

[2] D. Ron et A. Shamir: "Quantitative analysis of the full Bitcoin transaction graph", Financial Cryptography, 2013.

[3] A. E. Kosba, et al.: "Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts", IEEE Symposium on Security and Privacy 2016, 2016.

Please contact:

Maria Christofi, Trusted Labs, France
maria.christofi@trusted-labs.com

Aline Gouget, Gemalto, France
aline.gouget@gemalto.com

CWI Joins the Dutch National Blockchain Coalition as a Founding Member

by Eric Pauwels (CWI)

Earlier this year, as part of the government's Digital Agenda, the top-level ICT Team set up by the Dutch Ministry of Economic Affairs instigated the formation of a National Blockchain Coalition. This coalition is a joint initiative of over 20 organisations – including Centrum Wiskunde & Informatica (CWI) – active in government and research, as well as the financial, health, logistics, and energy sectors. With this NBC initiative, the Netherlands aims to become one of the international leaders in blockchain development and applications.

The founding partners of the National Blockchain Coalition (NBC) anticipate that blockchain technologies will open up new and more efficient digital value transactions, including micro-payments, and will therefore have a huge impact on services including administration, healthcare, finance, energy and logistics networks [1]. As a consequence, the coalition foresees positive effects on the autonomy of citizens, transparency of operations and cyber-security, as well as a significant reduction in administrative overheads.

The first action line in the agenda primarily focuses on the development of "digital identities", which would allow persons and legal entities, but also services, and even objects or devices to autonomously engage in efficient but trusted digital transactions. To realise this vision, various technological solutions need to be developed. In addition to the obvious need for technological solutions, the coalition is also exploring what is required to remove legal obstacles or perception-based objections that would hamper wider acceptance.

At CWI, researchers are interested in various aspects of blockchain technologies and applications. At the fundamental level there are deep questions about the computational foundations of trust and consensus in distributed peer-to-peer networks. Developing algorithms of provable performance that address the balance between decentralisation and permission-less access to the system on the one hand, and scalability and tamper-resistance on the other, remains an active area of research for cryptologists [2].

Another topic of interest concerns quantum-proof versions of blockchain technologies. The wider roll-out and adoption of these technologies will most likely coincide with the emergence of functional and commercially viable quantum computers and sophisticated quantum software. Failing to adequately address these technological innovations might rapidly render large-scale investments in blockchain-based information infrastructure worthless and threaten to put broad swathes of society and industry at risk of fraud and malicious interference.

Research in formal methods at CWI is also highly relevant, in particular, with respect to the use of smart contracts. As smart contracts embedded in blockchain are immutable and automatically triggered by transactions, there are new incentives to develop efficient mathematical methods that can be used to formally prove their correctness.

On the application side there is strong interest in using blockchain technologies to enable micro-payments in peer-to-peer markets. As partners in the European ERA-Net project GRID-FRIENDS, CWI researchers are involved in the design of a decentralised energy coordination infrastructure at Schoonschip, a new housing development in Amsterdam. The ambitious goal of this innovative building project is to create a community of about 40 floating family dwellings that together constitute a quasi-autarchic energy microgrid. The houses are furnished with solar panels, heat pumps as well as batteries to temporarily store electricity. The batteries come equipped with planning and optimisation algorithms developed at CWI that coordinate the exchange of surplus energy among neighbours in the microgrid. As a result, there is a pressing need for a decentralised but tamper-resistant platform to keep track of all the ensuing micro-transactions. CWI researchers are investigating how blockchains can be harnessed to address these and related issues, such as automated negotiation and preference elicitation [3].

In summary, the instigation of a Dutch National Blockchain Coalition testifies to the strategic importance and urgency assigned to the development and roll-out of blockchain technologies by top-level ICT team in the Netherlands. Researchers at CWI are actively investigating both fundamental and application-oriented aspects of various blockchain technologies.

The National Blockchain Coalition partners include: ABN AMRO, ING, Volksbank, Nationale Nederlanden, Rotterdam Port Authority, Enexis, Alliander, the Dutch Royal Notarial Association, Brightlands and the Ministries of Economic Affairs, Infrastructure and the Environment, Security and Justice, and Interior and Kingdom Relations. On the knowledge side they are joined by TU Delft, Tilburg University, Radboud University, Centrum Wiskunde & Informatica (CWI), NWO and TNO. The social perspective is contributed by ECP | Platform voor de Informatie Samenleving (Platform for the Information Society). On 20 March 2017, Minister of Economic Affairs of the Netherlands Henk Kamp received the action agenda of the National Blockchain Coalition.

Links:
https://kwz.me/XE
http://www.grid-friends.com
https://kwz.me/XK

References:

[1] Distributed Ledger Technology: beyond block chain. Report UK Government Office for Science. 19 Jan 2016 (https://www.gov.uk/government/publications/distributed-ledger-technology-blackett-review)

[2] A. Narayanan et. al. “Bitcoin and Cryptocurrency Technologies. A Comprehensive Introduction”. Princeton University Press, 2016

[3] T. Baarslag and M. Kaisers: “The Value of Information in Automated Negotiation. A Decision Model for Eliciting User Preferences”, Proc. of 16th Int. Conf. On Autonomous Agents and Multi-Agent Systems (AAMAS 2017).

Please contact:
Eric Pauwels, CWI, The Netherlands
+31 (0)20 592 4225
eric.pauwels@cwi.nl

Blockchain Lab – Design, Implementation and Evaluation of Innovative Business and Process Models

by Gilbert Fridgen, Wolfgang Prinz, Thomas Rose and Nils Urbach (Fraunhofer FIT)

Blockchain is considered to be enabling technology that is going beyond the Bitcoin crypto currency. It replaces centralised transaction management by the distribution of transactions across a network of nodes with different methods for consensus finding. This major change of governance may change sectors of our societies far beyond digital currencies. Fraunhofer FIT established a Blockchain Lab in 2016 in order to explore its impact. It will serve as an experience lab for technical components, implementation platforms, application prototypes and blueprints for novel governance, process, and business models.

A blockchain is essentially an electronic ledger for digital records, events or transactions maintained by the participants in a distributed computer network. This distribution of transaction management across a peer-to-peer network of interested parties plus new forms of consensus finding to preserve global consistency will allow significant changes to well-established service process and governance patterns. A blockchain may be used not only to distribute transaction management, but also to automate processes, rules and organisational principles. Using smart contracts, consistency rules may be attached to each transaction. They specify what has to be checked in a transaction and which follow-up activities have to be triggered.

This built-in automation will allow the re-engineering of many processes and the elimination of intermediaries and agencies, as long as information consistency is safeguarded by smart contracts compliant to auditing requirements. It is increasingly being used in a number of commercial and administrative fields, to distribute the management of transactions across an open network.

Blockchains also enable fundamental organisational changes in governance, so they may be characterised as a disruptive innovation that breaks up established business models. For instance, a blockchain might be used to maintain estate property records in a peer-to-peer

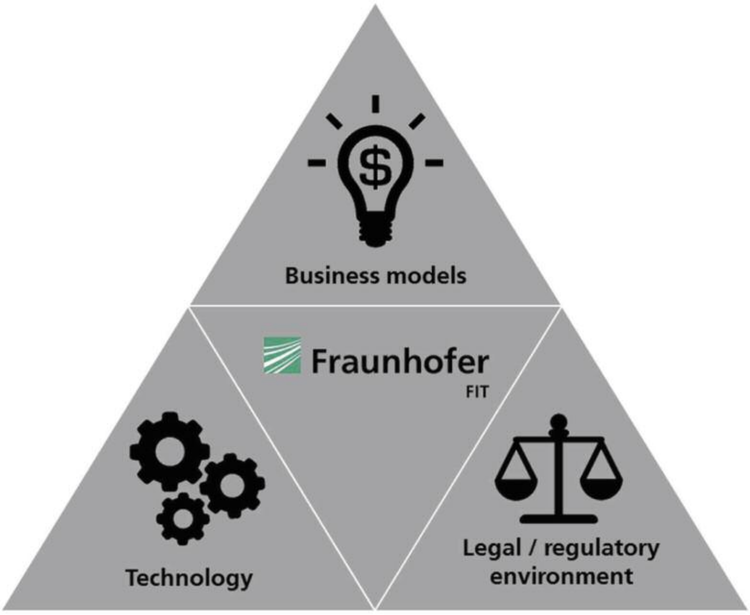

Figure 1: The multi-disciplinary approach of the Blockchain Lab.

network eliminating institutional overhead. Generally speaking, several different classes of use-cases have been identified, conceivably triggering novel service processes and governance. Providing, auditing and preserving provenance information is, for instance, an important service today that is vital for a broad range of industries.

Thus, blockchain technology may have many different effects not just on the processes but also on the structures of governance, which may significantly change the distribution of tasks among the actors involved in a process. The new roles and governance changes directly raises the question of new business models for the new value chain established by re-engineering the process.

As a consequence, the different dimensions of blockchain technology require a multi-disciplinary approach to exploit the potential capabilities of distributed transaction management combined with novel consensus methods. This will be implemented in the Blockchain Lab established at Fraunhofer FIT in 2016. It will serve as an experience lab for technical components, implementation platforms, application prototypes as well as blueprints for novel governance and business models. It is a multi-disciplinary unit rooted in three of FIT's research departments: Cooperation Systems for consensus methods, Decision Support for new governance and business models, and Information Systems for innovative applications. We will also look at the legal aspects of blockchain applications. Our aim is to showcase the state of the art in this fledgling research area using practical, integrative applications.

The work of the Blockchain Lab will be based on three cornerstones: business model, technology and legal / regulatory environment (see Figure 1). Business models will be developed together with individual partners or industry groups, focusing on potential analysis, evaluation and the design of disruptive solutions. Technology will focus on providing a development platform that includes different blockchain systems (P2P network, validation server etc.), implementation of blockchain solutions and evaluation of blockchain concepts. The relevant legal aspects and regulations are taken into account in our evaluation of business models and blockchain systems.

We will focus on short development cycles, which are quick to implement, in close cooperation with our development partners. We aim to develop functional applications quickly and to build marketable products through iterative improvements. To build these bespoke systems we will organise workshops, bring together consortia of industrial partners and conduct R&D projects covering all necessary development steps.

Link:
https://www.fit.fraunhofer.de/de/fb/cscw/blockchain.html

References:

[1] A. Narayanan, et al.: "Bitcoin and Cryptocurrency Technologies – A Comprehensive Introduction", Princeton University Press, 2016.

[2] V. Schlatt, A. Schweizer, N. Urbach, G. Fridgen: "Blockchain: Grundlagen, Anwendungen und Potenziale, Projektgruppe Wirtschaftsinformatik des Fraunhofer-Instituts für Angewandte Informationstechnik, 2016.

Please contact:

Gilbert Fridgen, Fraunhofer FIT,
Bayreuth Germany,
+49 921 55 4711,
gilbert.fridgen@fit.fraunhofer.de

Wolfgang Prinz, Fraunhofer FIT,
Sankt Augustin, Germany,
+49 2241 14 2730,
wolfgang.prinz@fit.fraunhofer.de

Thomas Rose, Fraunhofer FIT, Sankt
Augustin, Germany,
+49 2241 14 2798,
thomas.rose@fit.fraunhofer.de

Nils Urbach, Fraunhofer FIT, Bayreuth,
Germany,
+49 921 554-712,
nils.urbach@fit.fraunhofer.de

Machine Learning in IoT for Autonomous, Adaptive Sensing

by Frank Alexander Kraemer, Nattachart Tamkittikhun and Anders Eivind Braten (NTNU)

Inevitably, there will be a huge number of sensor devices within the internet of things (IoT) - but how can we possibly manage to optimise each and every one of them? Our answer is to treat them as autonomous units, much like robots. To this end we have been experimenting with different approaches to find out how constrained devices can benefit from machine learning, so that they can operate optimally.

Sensor devices are often situated in heterogeneous environments that change over time, for instance by changing location or variations in the weather. This is critical for their operation: Many sensor devices use energy harvesting, like solar energy, to sustain their operations, and their energy budget is critical to achieving their goals. This requires a high degree of optimisation. One of the characteristics of the internet of things (IoT), however, is its expected scale in terms of the number of devices. Therefore, the task of optimising IoT sensors or individually oversee their operation, cannot be performed manually. This leaves us with two options: Either over-dimensioning the system, for instance by investing in larger solar panels or batteries, or reducing the duty cycle of sensor devices to save energy, which effectively means to sense less frequently and send less data. In either case, systems do not operate optimally. This was also our experience within a smart city sensing project [1], where we used a static sensing approach. Sensing the emission data every six minutes worked adequately during the summer, but the solar panels could not provide enough energy during the dark winter in the Nordic areas, which eventually caused the sensor devices to shut down.

The experiences within this smart city project motivated our approach of autonomous and adaptive sensing in the ART project: Instead of looking at sensor devices as simple and constrained sources of data, we see them as autonomous agents, much like robots. Throughout their operation, they have to constantly plan ahead and make decisions based on the changing environment and what they have observed so far. Possible mechanisms for this include different machine learning techniques, applied in combination with each other.

To verify such an approach, we established a lab for autonomous sensors, which consists of an array of sensor nodes called Waspmotes, an off-the-shelf sensing system from Libelium driven by an 8-bit microcontroller. They communicate via LoRaWAN to a backend. Since the sensor nodes as well as the network are fairly constrained, the question is how machine learning can be applied in such a scenario. One solution is a centralised approach, in which machine learning is applied as part of the device management. Instead of just collecting data and monitoring key performance parameters such as sending frequency and battery level, the backend also learns from the received metadata and calculates optimised sensing strategies.

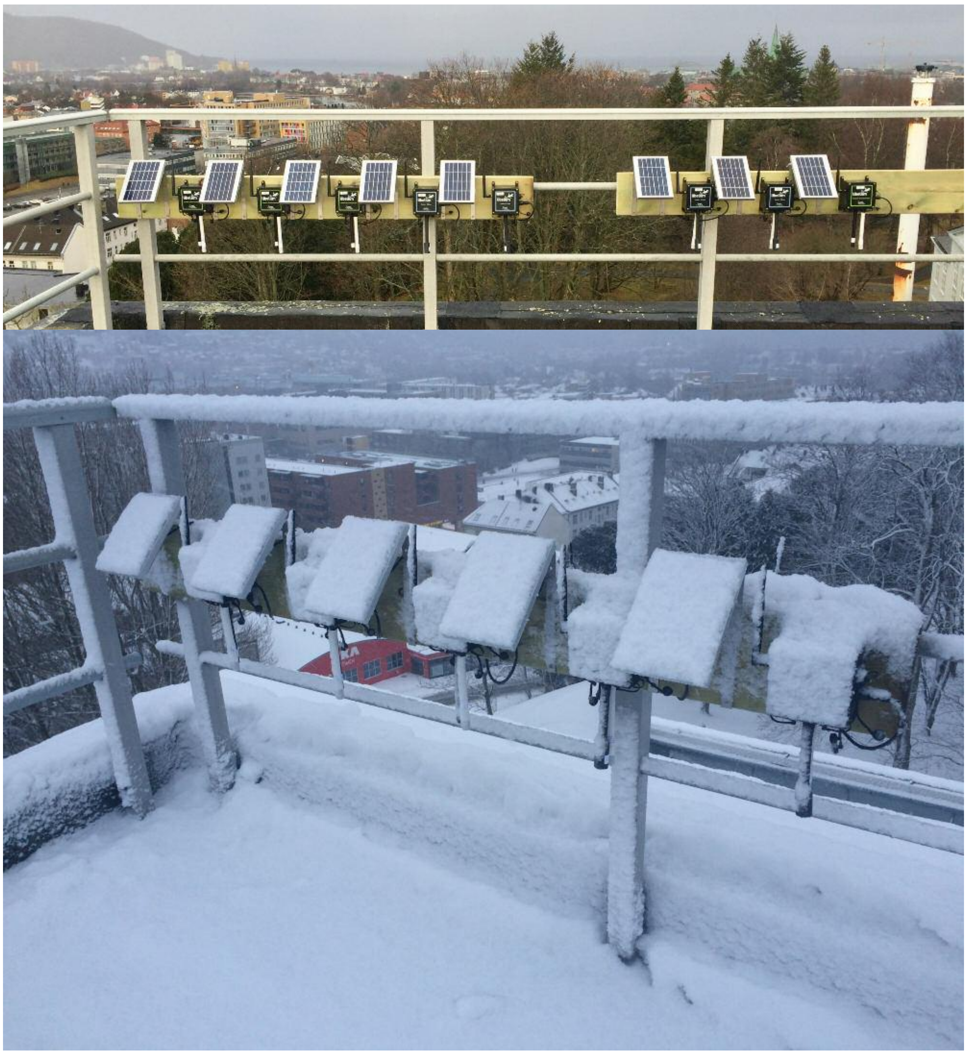

Figure 1: Sensor devices need to constantly adapt and plan ahead to maintain optimal operation in variable environmental conditions.

These sensing strategies are sent to the sensors every hour, and provide a guideline for how often data should be acquired to achieve a good balance between energy consumption and the required data rate.

This extended form of device management learns over time how the harvested solar energy depends on the current weather conditions. It also learns how the energy consumption changes with different sensing modes by considering the development of the battery level over time. Using weather forecasts, a planning algorithm predicts the resulting battery profile for different sensing strategies. The goal is both to keep the battery from being drained, and utilise the harvested energy to maximise the quality of data that is sensed and delivered. The SINet project follows a similar approach, but focuses on managing intermittent network connectivity.

The preliminary results are very encouraging. Already few features about the solar position (azimuth, zenith) and a couple of weather characteristics (cloudiness, precipitation, symbolic weather) are good indicators for the expected solar intake for the next day. We are experimenting with different machine learning techniques, including k-nearest neighbours and neural networks. In the given setting, the currently selected algorithms are less significant than the availability of sufficient training data. We are therefore investigating how autonomous sensing systems can be bootstrapped, i.e., how sensor nodes can start their operation even if little or no previous data exists for the prediction.

Another important question is how the system can perform its task in a less centralised way, i.e., considering autonomy on sensor level. To avoid single points-of-failure, sensors must be able to learn from insights gained at a global level, i.e., in the cloud, but still be able to act locally. For IoT applications, this implies a paradigm shift: machine learning methods should not only help analyse data collected by IoT nodes, but also help them to make optimal decisions about their own operation so that they can act autonomously.

Links:

SINet Project: https://sinet.item.ntnu.no
ART Project: https://ntnu.edu/iik/aas

Reference:

[1] Ahlers, D., Driscoll, P., Kraemer, F. A., Anthonisen, F., & Krogstie, J. (2016). A Measurement-Driven Approach to Understand Urban Greenhouse Gas Emissions in Nordic Cities. Norsk Informatikkonferanse NIK.

Please contact:

Frank Alexander Kraemer, Norwegian University of Science and Technology, NTNU, Norway
kraemer@ntnu.no

Cache-aware Roofline Model in Intel® Advisor

by Leonel Sousa and Aleksandar Ilic (INESC-ID)

Researchers from INESC-ID, Instituto Superior Técnico, University of Lisbon proposed a set of fundamental Cache-aware Roofline models, which provide a simple and intuitive way of visually representing the limits of parallel processing on multi-core processors.

As computing systems evolve towards complex multi-core designs with deep and diverse memory hierarchies, improving the performance and optimising the execution of real-world applications become of fundamental importance. In high-performance computing environments, it is crucial to determine which hardware resources represent the main execution bottlenecks that limit the application performance, especially when deciding on the most adequate software optimisation technique to be applied.In this process, simple but insightful models are particularly useful, since they provide the means to quickly and easily assess the main characteristics of the architectures and the features of the applications.

To support this decision process, researchers from INESC-ID, Instituto Superior Técnico (IST), University of Lisbon, Aleksandar Ilic, and Leonel Sousa, together with Frederico Pratas, PhD from IST, now with Imagination Technologies, proposed a set of fundamental Cache-aware Roofline models [1,2], which provide a simple and intuitive way of visually representing the limits of parallel processing on contemporary multi-core processors. These Cache-aware Roofline models evaluate how key micro-architectural aspects, such as accessing different functional units or different memory hierarchy levels, affect realistically achievable upper-bounds for performance, power consumption and energy-efficiency on a given multi-core architecture.

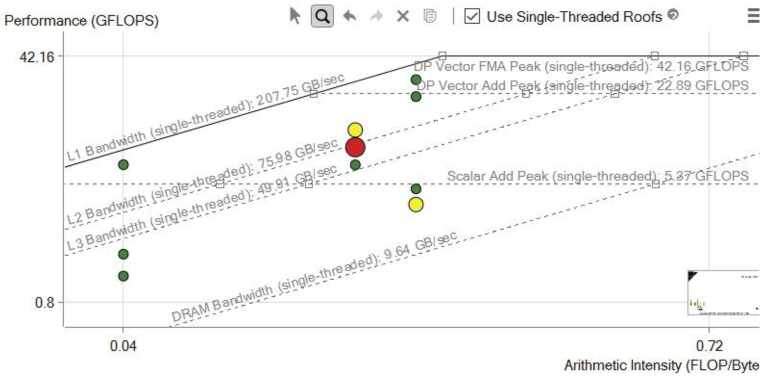

Figure 1: Cache-aware Roofline in Intel® Advisor.

In 2017, a team of Intel software developers (led by Zakhar Matveev, Roman Belenov and Philippe Thierry) successfully integrated the performance Cache-aware Roofline model as an official feature of Intel® Advisor, which is part of the Parallel Studio XE suite (Intel's main application development framework) [3,4]. Within Intel® Advisor, the process of building the roofline plots and in-depth application characterisation are fully automated with respect to the hardware platform where the applications are executed. The support for a wide range of Intel devices is also provided, which covers all contemporary Intel CPU micro-architectures (from Nehalem to Skylake) up to massively parallel coprocessors (e.g., Intel Xeon Phi Knights Landing).

A brief overview of the Cache-aware Roofline in Intel® Advisor

The performance Cache-aware Roofline is plotted with the X axis as arithmetic intensity (measured in FLOPs/Byte) and the Y axis as the performance in GFLOPs/Second, both in logarithmic scale. Before collecting data from a specific application, the Intel® Advisor automatically runs a set of quick benchmarks to measure the hardware limitations of the used processor, which it then plots as lines on the chart, called roofs (see Figure 1). The horizontal lines represent the number of floating point computations (of a given type) that the underlying hardware can perform in a given span of time. The diagonal lines are representative of how many bytes of data a given memory hierarchy level can deliver per second.

Each dot represents a loop or function in the program, and its position in the Roofline plot indicates performance and arithmetic intensity. The size and colour of the dots in Intel® Advisor's Roofline chart indicate how much of the total program time a loop or function takes: small, green dots take up relatively little time, so are likely not worth optimising; large, red dots take up the most time, so they are the best candidates for optimisation, especially those with a large gap to the topmost attainable roofs. In general, the further a dot is from the topmost roofs, the more room for improvement there is. For example, the Scalar Add Peak represents the maximum possible performance without taking advantage of vectorisation, as indicated by the next roof up being the Vector Add Peak.

Where can I get Intel® Advisor with Cache-aware Roofline?

As stated in the Intel early access program: The Intel Advisor offers a great step forward in memory performance optimization with a new vivid Advisor "Roofline" bounds and bottlenecks analysis. Cache-aware Roofline is currently a feature of Intel® Advisor beginning officially with version 2017 Update 2, which is part of the Parallel Studio XE suite (Cluster Edition and Professional Edition) [3].

References:

[1] A. Ilic, F. Pratas, and L. Sousa: "Cache-aware Roofline model: Upgrading the loft," IEEE Computer Architecture Letters, vol. 13, n. 1, pp. 21-24, 2014.
[2] A. Ilic, F. Pratas, and L. Sousa: "Beyond the Roofline: Cache-aware Power and Energy-Efficiency Modeling for Multi-cores," IEEE Transactions on Computers, vol. 66, n. 1, pp. 52-58, 2017.
[3] Intel. (2017) Intel® Advisor Roofline. http://tiny.cc/1i82ly

Please contact:

Leonel Sousa and Aleksandar Ilic, INESC-ID, Portugal
leonel.sousa@tecnico.ulisboa.pt,
aleksandar.ilic@tecnico.ulisboa.pt

Lightweight Random Indexing for Polylingual Text Classification

by Alejandro Moreo Fernandez, Andrea Esuli and Fabrizio Sebastiani (ISTI-CNR)

Researchers from ISTI-CNR, Pisa (in a joint effort with the Qatar Computing Research Institute), have undertaken an effort aimed at producing more accurate and more efficient means of performing poly-lingual text classification, i.e., automatic text classification in which classifying text in one language can also leverage training data expressed in a different language.

Multilingual Text Classification (MLTC) is a text classification task in which documents are written each in one among a set L of natural languages, and in which all documents must be classified under the same classification scheme, irrespective of language. This scenario is more and more frequent, given the large quantity of multilingual platforms and communities emerging on the Internet.

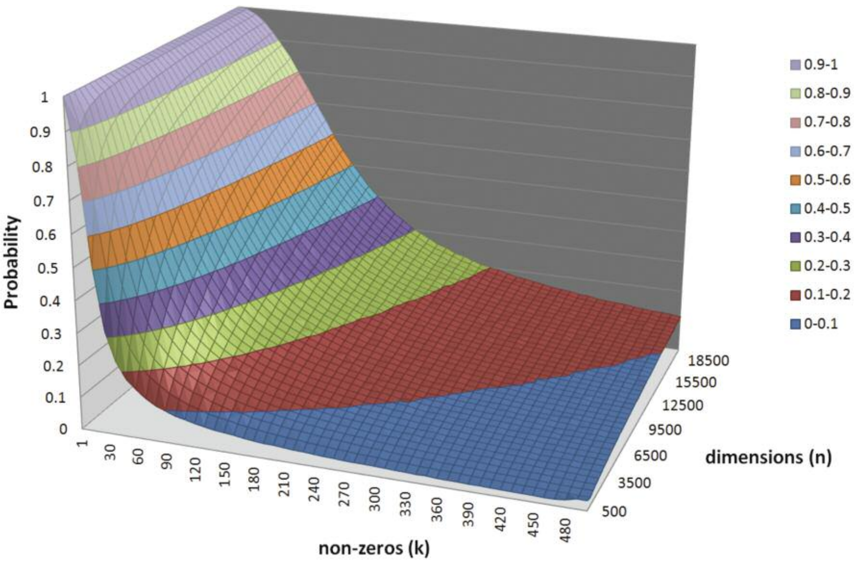

Figure 1: Variation in the probability of orthogonality of two random index vectors as a function the number of non-zero values in the random index vector and its dimensionality.

There are two main variants of MLTC, namely Cross-Lingual Text Classification (CLTC) and Polylingual Text Classification (PLTC). In CLTC we assume that for one or more of the languages in L there are no training documents; the task thus consists of classifying the test documents expressed in these languages by leveraging the training documents expressed in the other languages. In PLTC, which is the focus of this work, we assume (differently from CLTC) that for each language in L there is a representative set of training documents; PLTC consists of improving the accuracy of each of the |L| monolingual classifiers by also leveraging the training documents written in the other (|L|-1) languages. This task is receiving increased attention in the text classification community also due the new challenge it poses, i.e., how to effectively leverage polylingual resources in order to infer a multilingual classifier and to improve the performance of a monolingual one.

The obvious solution, consisting of generating a single polylingual classifier from the juxtaposed monolingual vector spaces, is usually infeasible, since the dimensionality of the resulting vector space is roughly |L| times that of a monolingual one, and is thus often unmanageable. As a response, the use of machine translation tools or multilingual dictionaries has been proposed. However, these resources are not always available, or are not always free to use.

One machine-translation-free and dictionary-free method that had never been applied to PLTC before, is Random Indexing (RI). RI is a context-counting model belonging to the family of random projection methods, which produces linear projections into a nearly-orthogonal reduced space where the original distances between vectors are approximately preserved. RI thus delivers semantically meaningful representations in a reduced space, and can be viewed as a cheaper approximation of Latent Semantic Analysis. We have analysed RI in terms of space and time efficiency, and have proposed as a result a particular configuration of it (that we have dubbed Lightweight Random Indexing -- LRI). LRI is designed so that the orthogonality of the projection base is maximized, which causes sparsity to be preserved after the projection (see Figure 1). The orthogonality of random index vectors plays an important role for features that are shared across languages: if their corresponding random index vectors are orthogonal with respect to all the other vectors, the information they contribute to the process is maximized, instead of being diluted by other less informative features.

We have run experiments on two well-known public benchmarks, Reuters RCV1/RCV2 (a comparable corpus -- i.e., documents are not direct translations of each other, but are simply about similar topics) and JRC-Acquis (a parallel one -- i.e., each text is available in all languages thanks to the intervention of human translators); for both benchmarks, we addressed five languages (English, Italian, Spanish, French, German). These experiments have shown LRI to outperform (both in terms of effectiveness and efficiency) a number of previously proposed machine-translation-free and dictionary-free PLTC methods that we used as baselines, including other more classical instantiations of random indexing.

Link:
http://http://jair.org/papers/paper5194.html

Please contact:
Fabrizio Sebastiani, ISTI-CNR, Italy
+39 050 6212892, fabrizio.sebastiani@isti.cnr.it

Real Flight Demonstration of Monocular Image-Based Aircraft Sense and Avoid

by Péter Bauer, Antal Hiba, Bálint Daróczy, Márk Melczer, Bálint Vanek (MTA SZTAKI)

The Institute for Computer Science and Control (SZTAKI) of the Hungarian Academy of Sciences (MTA) continues active research in the field of aircraft sense and avoid since six years. The sense (see) and avoid capability is a crucial ability in integrating unmanned aerial vehicles (UAVs) into the national airspace. SZTAKI focuses on the development of a monocular camera based on-board system, which processes image data in real-time and initiates collision avoidance if required. Capabilities of the system are continuously tested in real flight applying small UAVs testbeds.

Sense and avoid (S&A) capability is a crucial ability for future unmanned aerial vehicles (UAVs). It is vital to integrate civilian and governmental UAVs into the common airspace. At the highest level of integration, Airborne Sense and Avoid (ABSAA) systems are required to guarantee airspace safety. In this field, the most critical question is the case of non-cooperative S&A for which usually complicated multi-sensor systems are developed. However, in case of small UAVs the size, weight and power consumption of the onboard S&A system should be minimal. Monocular vision based solutions can be cost and weight effective therefore especially good for small UAVs. These systems basically measure the position (bearing) and size of intruder aircraft (A/C) from the camera image without range and intruder size information. This scale ambiguity makes the decision about the possibility of mid-air collision (MAC) or near mid-air collision (NMAC) complicated. However, [1] points out that the relative distance of an intruder from an own A/C (when it crosses the camera focal plane), called closest point of approach (CPA), well characterizes the possibility of collision.

The starting point of our research was to develop a relatively small onboard camera system, which can make all calculations onboard and in real-time to decide about the possibility of collision. The decision information can be sent to the autopilot of the aircraft to make an evasive maneuver if required. The developed system equipped with two HD cameras and a Tegra TK1 GPU board can be seen in Figure 1. After constructing the hardware system, the goal of theoretical developments was to estimate the time to the closest point of approach (TTCPA – the time remaining until aircraft reach the closest point), closest point of approach (CPA), and the direction of the intruder aircraft at CPA from a series of measured parameters of the intruder image in the camera, such as size and position. CPA is understood as the ratio of the smallest intruder distance relative to the own aircraft and the characteristic size of the intruder. As the absolute distance cannot be estimated from monocular images because of the scale ambiguity of perspective projection, this relative

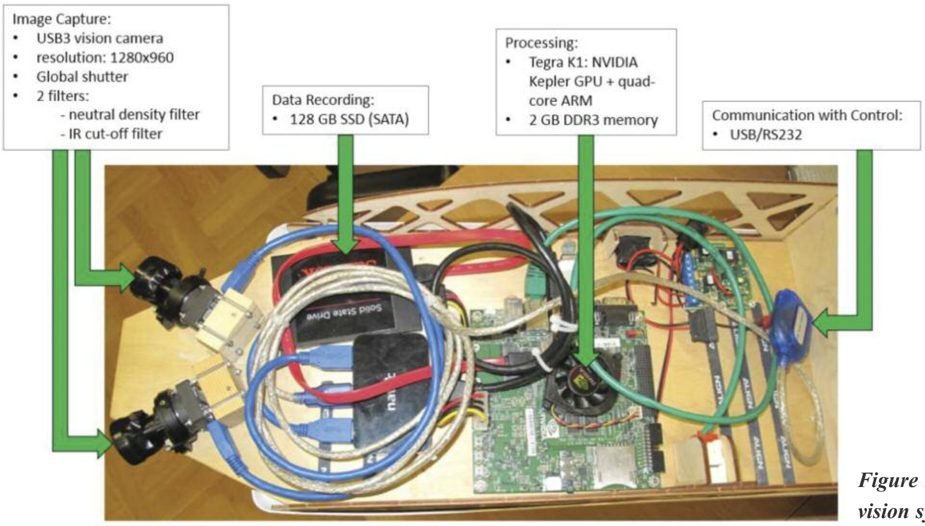

Figure 1: The onboard vision system.

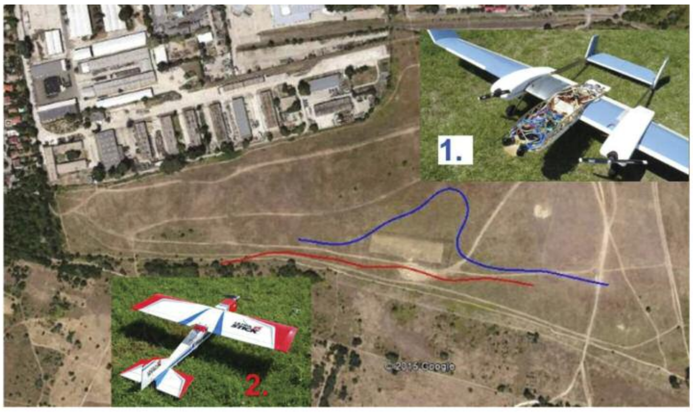

Figure 2: Aircraft photos and trajectories plotted over the airfield.

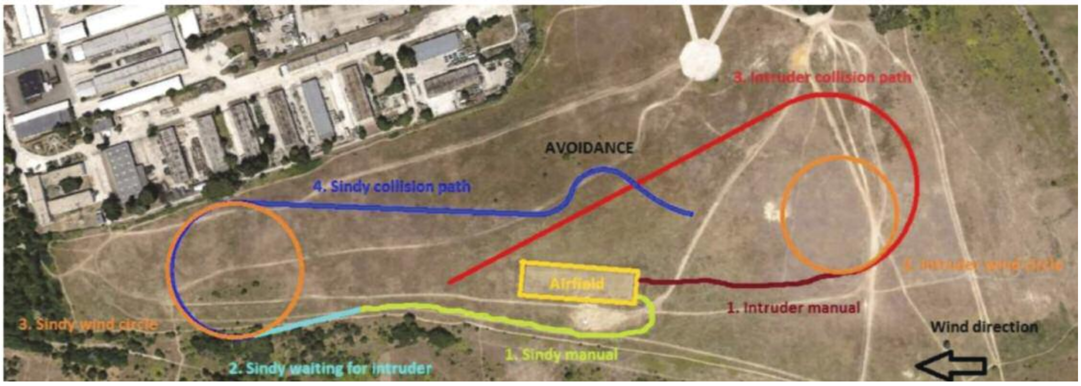

Figure 3: Non-parallel collision trajectory design.

ratio should be used to decide about the possibility of collision. In practice, this relative CPA value is enough to decide because one can avoid any intruder coming closer than a given multiple of its characteristic size. The details of the theoretical developments can be found in [2].

After the theoretical developments and hardware-in-the-loop test runs, real flight demonstrations are done in a critical UAV to UAV situation where the intruder aircraft (aircraft 2 in Figure 2.) is a small (about 1m wingspan) one. This makes detection and identification extremely difficult and decreases the time available to make a decision. The current state of real flight-testing is to make collision encounters with parallel aircraft trajectories. In this case there is definitely a point when the aircrafts are closest to each other. Several real flight tests have been conducted for these cases setting 15m or 30m distance between the trajectories. The 15m case is the NMAC where avoidance is required (and is shown by the blue trajectory in Figure 2.), the 30m is the clear case. Decision thresholds for TTCPA and CPA are set to distinguish NMAC and clear situations. About 25 NMAC and 25 clear flight runs have been done. Currently the system works with homogeneous sky background with a success rate of 80-100% regarding avoidance in NMAC and a false alarm rate of 40-50% in the clear case. This latter means unrequired avoidance when the intruder is far from us and that the tuning was done focusing on safety. Ground-based and onboard video recordings of the flight tests can be seen on our Youtube channel [L1].

The case of non-parallel trajectories requires a special design algorithm, which guides the aircraft along the required trajectories to meet at the same point at the same time. This task has two challenges: first, the starting points of the aircraft are uncertain as they take-off under manual control, and the autopilot is started only later. Second the wind disturbance should be measured and the designed tracks corrected accordingly. Currently successful hardware-in-the-loop (HIL) simulations are carried out resulting in collision tracks with and without wind. Video recordings of these experiences are shown on our Youtube channel [L1]. Illustrative non-parallel trajectories are plotted in Figure 3 showing manual take-off, wind measuring circles, designed trajectories and the avoidance also based-on HIL simulation.

The future goal is to test the non-parallel collision trajectories in real flight, which first requires the implementation of safe and guaranteed communication between the aircraft as they have to exchange data (for example the autopilot starting points and trajectory parameters).

Link:
[L1] http://yt.vu/c/UCQMpnOuOMCiodDKQw8_hf5A

References:

[1] S. D. B: "Reactive Image-based Collision Avoidance System for Unmanned Aircraft Systems", Master's thesis, Australian Research Centre for Aerospace Automation, May 2011.

[2] P. Bauer, A. Hiba, J. Bokor: "Monocular Image-based Intruder Direction Estimation at Closest Point of Approach", in Proc. of ICUAS'17, Miami, FL, USA, June 2017.

Please contact:

Peter Bauer, Senior Research Fellow, Project Coordinator
MTA SZTAKI, bauer.peter@sztaki.mta.hu

Use-cases Covered by an Enhanced Virtual Research Environment

by Valerie Brasse (IS4RI and euroCRIS)

The purpose of virtual research environments (VREs) is to support research stakeholders throughout the research life-cycle. In order to describe the expectations and requirements at each stage of the life-cycle, high-level use-cases have been developed in the VRE4EIC project, which has developed a Europe-wide interoperable VRE to facilitate innovation and collaboration between multidisciplinary research communities.

The goal of a virtual research environment (VRE) system is to provide support at each stage of the research life-cycle. The computing requirements will vary depending on the expected support, but commonalities can be identified.

Based on the first requirements collection in the VRE4EIC project, use-cases have been elaborated to express the possible ways in which users interact with the system. Use-cases are useful as they provide:

- Context to the requirements, expressing the requirements in terms of a concrete objective that the user wants to realise,

- Test scenarios for acceptance testing, i.e. a 'black-box' view of the system, what the user inputs and requires as output, unaware of the system's internal behaviour,
- Advocacy material to demonstrate to potential VRE users and developers what the e-VRE is expected to do.

In our approach, use-cases have been expressed at two levels of granularity:
- Low level use-cases, which assemble requirements in a coherent sequence,
- High-level use-cases are expressed as an orchestration of low level use-cases.

In the first wave of use-case elicitation (a second wave is expected by the end of the project), 59 low level use-cases and 19 high-level use-cases have been identified and described. For example, at the beginning of the research lifecycle, the first stages are 'Ideas, Partners, Proposal writing', where ideas come up and partners get together and write proposals to get funding on research projects. On their side, funding agencies set up work programmes and funding calls, and allocate funding to the selected proposals submitted by the researchers.

The unique high-level use-case defined at this stage of the research lifecycle is named 'Manage funding calls'. The associated actor is a 'funding agency', and the sequence of steps includes the suggestion to have the funding agencies publish information about themselves and their funding calls, then the agencies should retrieve the submitted proposals. Finally, the use-case provides generic requirements related to the e-VRE user interface.

The Research process part of the research lifecycle is itself divided into four stages:

The first stage called 'Simulate, experiment, observe' is covered by four high-level use-cases, namely: 'Create a dataset', 'Create a dataset from an instrument', 'Manage an instrument', and 'Communicate'. These use-cases are related to the data acquisition and the collaborative work done at this stage.

Still in the research process, the second stage deals with 'Managing data'. It is covered by five use-cases related to the use and transformation of data, namely 'Access services and data from e-VRE', 'Manage data', 'Manage research projects', 'Transform data', and 'Query data'.

The third step in the research process is called 'Analyse data'. Five use-cases also cover it. They are 'Access services and data from e-VRE', 'Annotate data', 'Compare datasets', 'Process data', and 'Communicate'. They are related to the analysis of data in context, and in a collaborative way.

The last step in the research process, and the last step in the research lifecycle, are about 'Sharing data' and the research results obtained from it, generally in the form of a scientific 'Publication'. Three use-cases cover this part: 'Access services and data from e-VRE', 'Publish dataset', and 'Communicate'.

Some of the use-cases, such as 'Access services and data from e-VRE' and 'Communicate', are recurrent, being used at several stages. This highlights the main characteristics of a VRE as an enabler of services and data access, and communication or collaboration amongst the research stakeholders.

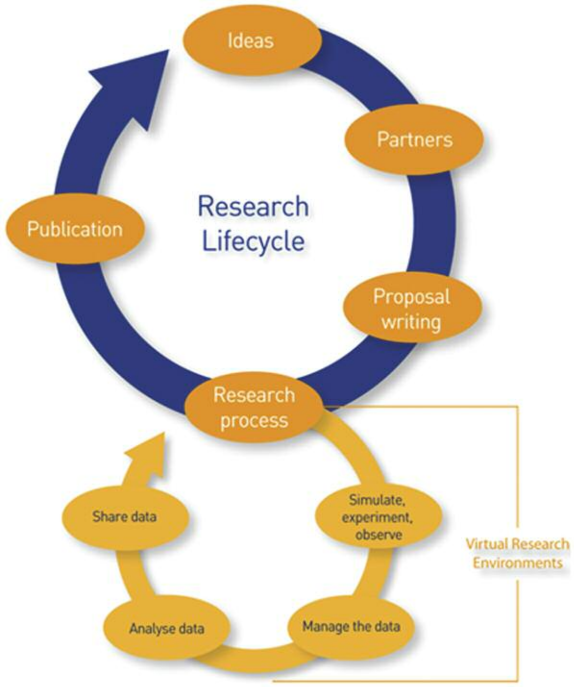

Figure 1: The use-cases that were developed for the target e-VRE platform cover the full research lifecycle.

As the project now goes into development, some of these use-cases will be implemented and available for acceptance testing in a further phase, in the context of enhancing current infrastructures for them to become VREs.

The following documents are available on demand:
- Y. Yin (ed.): "State-of-the-art and user requirements analysis. Deliverable D2.1 of the VRE4EIC Project.
- V. Brasse (ed.): "Use case report - First version. Deliverable D2.3 of the VRE4EIC Project.

Both available on demand by the project coordinator (see http://www.vre4eic.eu).

The project is currently carrying out a use-cases survey, which is targeted at anyone interacting with research data. Its purpose is to find out what researchers require in order to help them share and use data from multiple disciplines. Anyone involved in academic research activities can participate in this survey, since they might (potentially) share and/or use (open) research data. The results of this survey will be used to develop and further specify the requirements of the VRE developed in the VRE4EIC project, and will benefit future scientific research activities.

Links:
http://www.vre4eic.eu/
Survey: https://www.vre4eic.eu/publications/news/113-vre4eic-online-survey-to-evaluate-use-cases-for-a-virtual-research-environment

Please contact:
Valerie Brasse,
euroCRIS Executive for Projects, IS4RI Managing Partner
vbrasse@is4ri.com

ElasTest

ElasTest: A Cloud-based Platform for Testing Large Complex Distributed Software Systems

The ElasTest project, with an € 5M of EU funding under the Research and Innovation Action from the Horizon 2020 program, kicked-off in January 2017. The goal of ElasTest is to increase software quality by reducing the complexity of testing large distributed software systems in the Cloud. The project is led by the Spanish University Rey Juan Carlos (Prof. Francisco Gortázar) and involves the Technische Universität Berlin, the Consiglio Nazionale delle Ricerche, the Zurich University of Applied Sciences, the Fraunhofer FOKUS, the IMDEA Software Institute and the following industrial partners Atos Spain, IBM, Naevatec, and Relational.

The demand for larger and more interconnected software systems is constantly increasing, but the ability of developers to satisfy it is not evolving accordingly. The most limiting factor is the software validation, which typically requires very costly and complex testing processes to ensure the software is free of errors and complies with requirements. The ElasTest project aims at offering a flexible open source testing platform for rapid and accurate end-to-end testing that can significantly improve the efficiency and effectiveness of the testing process and the overall quality of modern applications, including web, mobile, real-time video communications, and Internet-of-Things.

The ElasTest cloud platform will be released as Free Open Source Software and has already started creating a Community of users and contributors who will help us in our endeavor for transforming ElasTest into a worldwide reference in the area of large software systems testing and guaranteeing the long term platform sustainability. To join us visit our community website http://elastest.io/

The CNR group, led by Antonia Bertolino (ISTI-CNR, Pisa), will contribute to the project with two important missions: leading the continuous research scouting in cloud testing and coordinating the experimental validation. The quantitative project objectives include reducing the time-to-market, increasing the quality of the software product, reducing the possibility of failures and improving the confidence and satisfaction of both end users and developers.

More information:
http://elastest.eu/
http://www.isti.cnr.it/events/2017/ELASTEST-Pressrelease.pdf

13th EUROPEAN COMPUTER SCIENCE SUMMIT

www.informatics-europe.org/ecss

ECSS 2017

LISBON - PORTUGAL
23-25 OCTOBER

Organized by

INFORMATICS EUROPE

UNIVERSIDADE NOVA DE LISBOA

NOVALINCS
LABORATORY FOR COMPUTER SCIENCE AND INFORMATICS

Sponsored by acm Europe Council

INFORMATICS CONNECTS!

ECSS is the forum to debate the trends and issues that impact the future of Informatics in Europe.

Join leading decision makers and hear from renowned speakers about the challenges and opportunities for Informatics research and education in an increasingly interconnected world.

Conference Chairs:
Luís Caires, Department of Informatics, Universidade NOVA de Lisboa
Lynda Hardman, Informatics Europe, CWI and Utrecht University

Program Chairs:
José Luiz Fiadeiro, Department of Computer Science, Royal Holloway University of London
Markus Püschel, Department of Computer Science, ETH Zurich

Pre-Summit Workshop For Deans and Department Heads (October 23)

Co-located meetings:
ERCIM and Academia Europaea Informatics Section

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Call for Proposals

Dagstuhl Seminars and Perspectives Workshops

Schloss Dagstuhl – Leibniz-Zentrum für Informatik (LZI) is accepting proposals for scientific seminars/ workshops in all areas of computer science, in particular also in connection with other fields.

If accepted the event will be hosted in the seclusion of Dagstuhl's well known, own, dedicated facilities in Wadern on the western fringe of Germany. Moreover, the Dagstuhl office will assume most of the organisational/ administrative work and the Dagstuhl scientific staff will support the organizers in preparing, running, and documenting the event. Due to subsidies the costs are very low for participants.

Dagstuhl events are typically proposed by a group of three to four outstanding researchers of different affiliations. This organizer team should represent a range of research communities and reflect Dagstuhl's international orientation. More information, in particular details about event form and setup as well as the proposal form and the proposing process can be found on

http://www.dagstuhl.de/dsproposal.

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is funded by the German federal and state government. It pursues a mission of furthering world class research in computer science by facilitating communication and interaction between researchers.

Important Dates

- Proposal submission: October 15 to November 1, 2017
- Notification: End of January 2018
- Seminar dates: Between mid 2018 and mid 2019.

IFIP TC6's Open Digital Library and Conferences

by Harry Rudin

IFIP Networking 2017 took place in Stockholm at KTH, the Swedish Royal Institute of Technology, from 12-17 June. The conference's aim was "The future of networking: Defined by software, driven by data and designed for all" and consisted of workshops, regular lectures and poster sessions. As is our aim in the IFIP Technical Committee 6 (Communication Systems), the papers are published in our open Digital Library, http://dl.ifip.org. There is no printed version of the conference papers; the Digital Library serves as the conference record and papers were made available several days before the conference opening. Access is free to everyone.

The best paper award was made for a paper on achieving low latency and high throughput for networks using a modified Explicit Congestion Notification (ECN) scheme through the use of active queue management (AQM). The title of the paper is "Alternative Backoff: Achieving Low Latency and High Throughput with ECN and AQM" and is written by Naeem Khademi, Grenville Armitage, Michael Welzl, Sebastian Zander, Gorry Fairhurst and David Ros.

Nearly all of IFIP TC6's conferences are published in the Digital Library; there are nine of these at the moment. Other IFIP conferences are also available via the Digital Library; there are over 150 of these. Do have a look!

IFIP TC6's Committee had their meeting alongside the conference. A change in the chairmanship took place: Aiko Pras from the University of Twente, in The Netherlands is followed by Burkhard Stiller from the University of Zurich in Switzerland.

Please contact:
Harry Rudin
hrudin@sunrise.ch

Call for Participation

ACM NanoCom 2017

Washington DC, USA, September 27-29, 2017

The main goals of the 4th ACM International Conference on Nanoscale Computing and Communication (ACM NanoCom 2017), are to increase the visibilty of this growing research area to the wider computing and communication research communities as well as bring together researchers from diverse disciplines that can foster and develop new paradigms for nanoscale devices. Due to the highly inter-disciplinary nature of this field of research, the conference aims to attract researchers and academics from various areas of study such as electrical and electronic engineering, computer science, biology, chemistry, physics, mathematics, bioengineering, biotechnology, materials science, nanotechnology, who have an interest in computing and communications at the nanoscale.

Keynotes

The following keynotes are presented:

- "The Living Computing Project - How Can I Make A Cell Compute?" by Douglas Densmore, Associate Professor of Electrical and Computer Engineering, Director of Cross-disciplinary Integration of Design Automation Research (CIDAR) Group, Boston University.

- "Redox: A Modality to Bridge Biological and Electronic Communication" by Gregory F. Payne, Professor of Bioengineering, Fischell Department of Bioengineering, University of Maryland.

- "Graphene and Related Materials for Photonics and Optoelectronics" by Andrea Ferrari, Professor of Nanotechnology, Director of the Cambridge Graphene Centre; Director of EPSRC Centre for Doctoral Training in Graphene Technology; Head of the Nanomaterials and Spectroscopy Group, University of Cambridge, UK:

More information:
https://nanocom.acm.org/

New 3D FleX-ray Lab at CWI

In May, demissionary Secretary of State of the Netherlands Sander Dekker opened the ‘FleX-ray Lab’ at CWI. With the new state-of-the-art CT scanner of this lab, it will be for the first time possible to look inside objects in 3D during the scanning process, and to adjust or zoom in while scanning. Thanks to realtime data processing and adjustment the scanner is able to retrieve more useful information from the scans – faster, with less harmful X-ray dose and in colour – than with current technologies. The new techniques can be used for medical imaging, quality control in the food industry and, for instance, for the restauration of antique masterpieces. The scanner has been developed in collaboration with X Ray Engineering (a spin-off of Ghent University), research institute Nikhef and its spin-off ASI. The lab is supervised by Joost Batenburg, group leader of the Computational Imaging group at CWI, internationally leading in the area of new mathematical image reconstruction techniques. CWI will make the research data and real-time software available as open source. See: https://www.cwi.nl/research/groups/computational-imaging

The opening of the FleX-ray Lab. F.l.t.r: Denis van Loo (XRE), Sophia Coban (CWI), Hans Roeland Poolman (ASI), Secretary of State Sander Dekker, Joost Batenburg (CWI), Els Koffeman (Nikhef), Jos Baeten (CWI), and Peter van Laarhoven (CWI).

18.8 Million Euro for Quantum Software Research

The Ministry for Education, Culture and Science in the Netherlands has awarded a Gravitation grant for large-scale research on quantum software in May. The grant of 18.8 million euro unites researchers from QuSoft, CWI, Leiden University, QuTech, TU Delft, UvA and the VU in pursuing state of the art research programmes in this new field. It allows the consortium to pioneer quantum software for, for instance, small quantum computers and a quantum internet. It will develop protocols for quantum communication, and for a new type of quantum-secure cryptography. These methods can be tested on quantum hardware that is developed in parallel in Delft and Leiden, and a quantum network between Amsterdam, Delft, Leiden and The Hague.

New ERCIM Board Members

Dieter Fellner, Director of the Fraunhofer Institute for Computer Graphics Research IGD and Professor of Computer Science at TU Darmstadt replaces Matthias Jarke on the ERCIM AISBL Board and the ERCIM EEIG Board of Directors as representative of Fraunhofer-Institut.

Before Dieter Fellner took office in Darmstadt in 2006, he has held academic positions at the Graz University of Technology, the University of Technology in Braunschweig, the University of Bonn, the Memorial University of Newfoundland, and the University of Denver, Colorado. He is still affiliated with the Graz University of Technology where he chairs the Institute of Computer Graphics and Knowledge Visualization he founded in 2005. Dieter Fellner is also CEO of the Fraunhofer Austria Research GmbH since November 2008. Since January 2016, Dieter W. Fellner serves as chairman of the Fraunhofer ICT Group and as member of the Fraunhofer Presidential Council.

Dimitris Plexousakis, Director of FORTH - ICS and Professor of Computer Science, University of Crete, succeeds Constantine Stehpanidis on the ERCIM EEIG Board of Directors.

Dimitris Plexousakis is a Professor and former Chair of the Department of Computer Science, University of Crete and a Researcher at the Institute of Computer Science, FORTH in Greece. He is heading the Information Systems Laboratory of FORTH-ICS. In 2017 he took office as Director of FORTH-ICS. Dimitris has served as representative of FORTH-ICS on the ERCIM AISBL as chair of the ERCIM Science Task Group since 2012.

We would like to express the warmest thanks to Matthias Jarke and Constantine Stephanidis for their leadership, and their enthusiastic participation and active contributions over many years.

ERCIM – the European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development in information technology and applied mathematics. Its member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

ERCIM is the European Host of the World Wide Web Consortium.

Consiglio Nazionale delle Ricerche
Area della Ricerca CNR di Pisa
Via G. Moruzzi 1, 56124 Pisa, Italy
http://www.iit.cnr.it/

Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
http://www.cwi.nl/

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
http://www.fnr.lu/

Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
http://www.ics.forth.gr/

Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
http://www.iuk.fraunhofer.de/

INESC
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, nº 378,
4200-465 Porto, Portugal

Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
http://www.inria.fr/

I.S.I. – Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
http://www.isi.gr/

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
http://www.ntnu.no/

SBA Research gGmbH
Favoritenstraße 16, 1040 Wien
http://www.sba-research.org/

SICS Swedish ICT
Box 1263,
SE-164 29 Kista, Sweden
http://www.sics.se/

Magyar Tudományos Akadémia
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
http://www.sztaki.hu/

University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
http://www.cs.ucy.ac.cy/

Universty of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Banacha 2, 02-097 Warsaw, Poland
http://www.mimuw.edu.pl/

VTT Technical Research Centre of Finland Ltd
PO Box 1000
FIN-02044 VTT, Finland
http://www.vttresearch.com

Subscribe to ERCIM News and order back copies at http://ercim-news.ercim.eu/