GeNeDis manuscript No.
(will be inserted by the editor)

Blockchain For Mobile Health Applications

Acceleration With GPU Computing

Georgios Drakopoulos · Michail Marountas · Xenophon Liapakis · Giannis Tzimas · Phivos Mylonas · Spyros Sioutas

Received: date / Accepted: date

Abstract Blockchain is a linearly linked, distributed, and very robust data structure. Originally proposed as part of the Bitcoin distributed stack, it found a number of applications in a number of fields, most notably in smart contracts, social media, secure IoT, and cryptocurrency mining. It ensures data integrity by distributing strongly encrypted data in widely redundant segments. Each new insertion requires verification and approval by the majority of the users of the blockchain. Both encryption and verification are computationally intensive tasks which cannot be solved with ordinary off-the-shelf CPUs. This has resulted in a renewed scientific interest in secure distributed communication and coordination protocols. Mobile health applications are growing progressively popular and have the enormous advantage of timely diagnosis of certain conditions. However, privacy concerns have been raised as mobile health application by default have access to highly sensitive personal data. This chapter presents concisely how blockchain can be applied to mobile health applications in order to enhance privacy.

Keywords Blockchains · Digital health · Edge computing · Mobile computing · Mobile applications · Majority protocols · GPU computing

Mathematics Subject Classification (2010) 65Y05 · 68Q05 · 68Q10 · 68W10

G. Drakopoulos and P. Mylonas
Department of Informatics, Ionian University, Greece
E-mail: {c16drak, fmylonas}@ionio.gr

X. Liapakis
Interamerican SA, Greece
E-mail: liapakisx@interamerican.gr

G. Tzimas
Technological and Educational Institute of Western Greece, Antirrio Campus, Greece
E-mail: tzimas@teimes.gr

M. Marountas and S. Sioutas
Computer Engineering and Informatics Department, University of Patras, Greece
E-mail: {marounta, sioutas}@ceid.upatras.gr

1 Introduction

Perhaps the most well studied recent advent in the domain of distributed comouting and data structures is that of blockchain. The latter acts as a public or private ledger and from a structural perspective is a linear, distributed, and robust data structure in the sense that not only the insertion of new data requires special permission from its stakeholders, mostly but not necessarily ordinary netizens with a legitimate vested interest in a given blockchain, which is obtained from specially designed consensus protocols, but also the true netizen identities participating to a given block chain as well as data contained therein are strongly encrypted, typically with a public key scheme such as SHA-256. Additionally, the exact location of data insertion is decided on the basis of a secure hash function. Finally, when the number of netizens participating to a blockchain is large, typically in the thousands, it becomes difficult to hack or game it as any malicious changes become visible almost immediately.

Having the computational properties just described, a blockchain is an excellent data structure for securely storing large volumes of information for a very broad spectrum of purposes including but not limited to smart contracts, digital health information, smart city and smart infrastructure status, financial macrotransactions as well as gaming and social media microtransactions, and insurance information. In fact, the blockchain as a data structure was initially part of the Bitcoin stack as described in Nakamoto (2008) or as later explored in cryptocurrency surveys as for instance Antonopoulos (2014) or Antonopoulos (2017). Since then, however it took a life of its own with numerous parties developing some version of the original blockchain for their own purposes.

Blockchain is not the only recent computationally intensive development. Fields like numerical and distributed deep learning such as the training of multilyayer convolutional and recurrent neural networks, complex systems simulation such as brain networks and protein-to-protein interaction networks, as large scale social network analysis are notorious for their quick scaling. One response to the need for additional computational power was the development of hardware aiming at massive parallelism through special purpose GPUs along with the associated software which can take advantage of such specialized hardware and can orchestrate the appopriate sequence of computations to derive the desired result. Google TensorFlow, a low level framework whose primary unit is a tensor as explained among others in Abadi et al. (2016), namely a multidimensional array, belongs to this category.

The primary objective of this chapter is to concentrate and succintly present the ways TensorFlow and GPU computation in conjunction with blockchain can empower applications in the domains of digital health and insurance market. As a secondary objective, the computational capabilities and the dataflow model of TensorFlow are analysed.

Table 1 Notation of this chapter.

symbol	meaning
$\triangleq$	Definition or equality by definition
$\langle s_k \rangle$	Sequence with elements s_k
$\lvert\langle s_k \rangle\rvert$	Sequence cardinality

The remaining of this work is structured as follows. In section 2 the relevant scientific literature regarding blockchain, GPU computing, and their applications is briefly reviewed. The properties of the blockchain as well as these of TensorFlow are described in section 3, whereas the blockchain applications in the domains of digital health and insurance are explored in section 4. The main findings of this chapter as well as possible future research directions are stated in section 5. Finally, table 1 summarises the notation of this work.

2 Previous Work

Blockchains were formally introduced in the seminal Bitcoin work of Nakamoto (2008). Their technological innovation and the potential to become a disruptive technology was explored among others in Barber et al. (2012) and in Cachin (2016). The combination of blockchains with the IoT and their applications to the mainstream industrial sector in conjunction with the upcoming digital transformations of Industry 4.0 are the focus of Miller (2018). Practical ways and the associated challenges to implement a blockchain over IoT and edge computing are shown in Zyskind et al. (2015). The financial prospects of Bitcoin in terms of wealth accumulation as well as the properties of Bitcoin versus the traditional fiat currency are the focus of a number of works, for instance Antonopoulos (2014), Antonopoulos (2017), Kosba et al. (2016), Swan (2015), and Böhme et al. (2015). The distributed implementation of blockchains is discussed in Abbas et al. (2018) and in Pass et al. (2017), whereas security aspects of the blockchains are treated in Puthal et al. (2018). A large number of blockchains besides the Bitcoin stack can be found in Underwood (2016).

Since the original public description of TensorFlow in Abadi et al. (2016) and in Abadi (2016) it was widely adopted from the deep learning community. In Matthews et al. (2017) a Gaussian process generator implemented with rudimentary TensorFlow operations is described in detail. For a new graph resilience metric based on paths see Drakopoulos et al. (2018b) along the lines of the regularization methodology of Kanavos et al. (2017). The advantages of and the ways for visualising the TensorFlow computations are Wongsuphasawat et al. (2018). For tensor applications in social network analysis such as multiway digital influence estimation see Drakopoulos et al. (2017), community structure discovery Drakopoulos et al. (2018a), and graph based k-means initialization Drakopoulos et al. (2016). Finally, for a genetic algorithm for clustering tensors containing linguistic and spatial data see Drakopoulos et al. (2019).

3 Parallelism and Blockchain

3.1 Blockchains

As their collective name suggests, from a structural point of view blockchains are, typically very long, linearly linked nodes. Each such node contains part of the postmarked information stored in the data structure along with some administrative information. The data stored in a blockchain can never be erased, although it can be updated provided all interested parties agree on that. Thus, both the original and the updated data are stored, making audits efficient.

Perhaps the most important advantages of selecting a blockchain scheme besides the increased security are the following:

- Blockchains support a very large volume of transactions which can take place almost simultaneously because of their very inherent distributed nature. Therefore, their stakeholders can perform any desired number of transactions within a very reasonable amount of time without worrying about the exact transaction execution time, which in certain cases may influence the transaction cost.
- The stakeholders of a given blockchain can stay informed of the global status of the blockchain in almost real time. Thus, not only can they perform transactions but they can also know their results almost immediately or at least at the moment the latter are actually executed.
- Blockchains, either public or private, offer full transparency since every participant to a given blockchain is free to validate any tranaction which took place within that blockchain. Additionally, the verification protocols are deliberately built so that verification be easy even for netizens with low computational resources, for instance a smartphone or a tablet. This reinforces the trust toward properly implemented and managed blockchains.
- In the case of a catastrophic loss, a properly implemented blockchain can at least partially rebuild itself from the segments stored at the computers of its stakeholders. This is feasible given the increased redundancy integrated into a blockchain.
- From a software engineering viewpoint, each blockchain node is a relatively simple construct and, therefore, it can be managed with little or no human intervention. Thus, a blockchain administrator is only required to control certain a few key parts of the data structure, making blockchains easy and inexpensive to maintain.
- Last but not least, any third parties and intermediaries are no longer necessary. The interested parties can directly communicate and get current quotes or any other vital pieces of information from each other.

Notice that blockchains are not immune to various sophisticated attacks, although the latter typically require considerable resources which are nowadays well within the capabilities of a dedicated hacker group or of a government agency. Although directly attacking the encryption protocols may not be a wise course of action, unless some knowledge of the private key is available, using a zero day exploit is.

As with any new technology, blockchain management software is by no means error free. However, most known attacks so far take on a completely different approach akin to a brute force attack by relying on big botnet networks in order to take charge of a small or medium sized blockchain.

Yet another method, holistic in nature, for attacking a blockchain is through the use of control theory concepts. The current state, in any way that is estimated by the attacker, of a large blockchain is represented as a control vector $\mathbf{x}[n]$. Then a usually linear state space model is formulated as follows:

$$\begin{aligned} \mathbf{x}[n+1] &\triangleq \mathbf{A}\mathbf{x}[n] + \mathbf{b}u[n] \\ \mathbf{y}[n+1] &\triangleq \mathbf{c}^T\mathbf{x}[n+1] + du[n] \end{aligned} \tag{1}$$

If the attacker can insert an appropriate input sequence $\langle u_k \rangle$, then, depending on the modelling correctness, he may bring the entire system to an undersirable state. Of course, such a sequence may not exist or its cardinality $|\langle u_k \rangle|$ might approach infinity.

3.2 TensorFlow

Google TensorFlow is a low level programming framework based on the dataflow programming paradigm and using tensors, namely multidimensional arrays, as its primary data structure. Originally developed for simulating brain networks, it is a powerful tool for deep learning. It has official APIs for Python and C++, whereas unofficial APIs are being developed for a number of established programming languages. Moreover, it has computational kernels for CPUs, GPUs, and TPUs.

Besides the methods for elementary operations such as Kronecker and Hadamard tensor multiplication, minimum location, tensor reshaping, and tensor factorizations such as Kruskal and Tucker decompositions, TensorFlow has a number of numerical optimizers which are common in deep learning such as AdaGrad. Also, TensorFlow supports checkpoints, allowing the early termination of a training process.

Within a blockchain context, TensorFlow can accelerate numerical computations for hashing or encryption. Additionally, it can be used to train a neural network, recurrent, convolutional, autoencoding, or otherwise, which can predict the volume in the immediate future, so that a bursty load of transactions can be better balanced throughout the blockchain nodes. Moreover, similar networks can be built in order to predict which blockchain user will be the next to generate a chunk of data or will ask for a transaction verification, again for load balancing purposes. Finally, large deep learning networks can in theory be deployed in order to mount an attack on the encryption protocol used by a given blockchain, but to the best of the knowledge of the authors, no such use has been recorded.

4 Applications

The blockchain as a ledger structure, either public or private, because of its secure and distributed design is a place for storing sensitive data such as health condition and financial transactions. Additionally, as stated earlier any intermediaries are eliminated, at least in a higher level. Thus, any fees and premiums such as taxes or bank processing fees are also, in theory at least, automatically gone.

Regarding the digital health world, blockchain-based applications have an enormous potential. The following list contains some of the most prominent ones.

- The medical records of a netizen can be stored with safety in a blockchain and can be recovered only by the certified health professional who cure the netizen regardless of their location or whether they have cured her before.
- Blockchain can facilitate automated monitoring of selected biomarkers by smartphones and the measurements can be compared against personalized baselines.
- Netizens have much improved control over their personal records and their consent can be obtained under more transparent and clear conditions.
- Netizens can use micropayments or mobile payments in order to procure medicines, further protecting their privacy.

Concerning the growing insurance market, there is also a significant room for blockchain-based applications. Some indicative are the following:

- Netizens can search easier for attractive offers and can contact insurance agents directly in order to negotiate for even better offers. This can also be done through software agents configured to look specific offers or terms.
- Netizens and insurance agents can hold smart contracts such as property and vehicle electronic contract purchses in blockchains. At a later point, should the need arise, they can directly renegotiate contract terms which will be also recorded in the blockchain, provided the interested parties reach an agreement.
- Once smart contracts are recorded, ordinary shallow or deep learning algorithms can be run atop the blockchain in order to identify possible fraud cases.
- Blockchains simplify considerably payments and can even be combined with mobile payments. Payment records remain immutable and constitute proof that a payment indeed took place at the time indicated.
- Claims can be automatically verified by smartphones and other personal devices which are connected to the blockchain, reducing thus the administrative burden and the overhead.

At this point it should also be reminded that the general advantages of section 3 also hold in addition to those listed above.

5 Conclusions

The twofold epicenter of this chapter was the blockchain applications in the domains of digital health and insurance market and the ways Google TensorFlow, a low level

computational framework for computationally intensive applications, can be used to accelerate the associated computations. The blockchain has numerous applications in the domains of medical healthcare and insurance. Moreover, it reinforces the privacy and transparency conditions and, thus, help establish a viable and scalable market.

Further research directions include the development of extensively tested blockchain management systems so that most, if not all, zero day exploits are eliminated. Moreover, given the recent advances in quantum computing which make large scale brute force attacks feasible, stronger cryptographic schemes should be sought in order to protect the sensitive personal data stored in blockchains.

Acknowledgements We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

Abadi M (2016) TensorFlow: Learning functions at scale. ACM SIGPLAN Notices 51(9):1–1

Abadi M, et al. (2016) TensorFlow: A system for large-scale machine learning. In: OSDI, vol 16, pp 265–283

Abbas N, Zhang Y, Taherkordi A, Skeie T (2018) Mobile edge computing: A survey. IEEE Internet of Things Journal 5(1):450–465

Antonopoulos AM (2014) Mastering Bitcoin: Unlocking digital cryptocurrencies. O'Reilly Media, Inc.

Antonopoulos AM (2017) Mastering Bitcoin: Programming the open blockchain. O'Reilly Media, Inc.

Barber S, Boyen X, Shi E, Uzun E (2012) Bitter to better - How to make Bitcoin a better currency. In: International Conference on Financial Cryptography and Data Security, Springer, pp 399–414

Böhme R, Christin N, Edelman B, Moore T (2015) Bitcoin: Economics, technology, and governance. Journal of Economic Perspectives 29(2):213–38

Cachin C (2016) Architecture of the hyperledger blockchain fabric. In: Workshop on Distributed Cryptocurrencies and Consensus Ledgers, vol 310

Drakopoulos G, Gourgaris P, Kanavos A, Makris C (2016) A fuzzy graph framework for initializing k-means. IJAIT 25(6):1–21

Drakopoulos G, Kanavos A, Mylonas P, Sioutas S (2017) Defining and evaluating Twitter influence metrics: A higher order approach in Neo4j. SNAM 71(1)

Drakopoulos G, Gourgaris P, Kanavos A (2018a) Graph communities in Neo4j: Four algorithms at work. Evolving Systems DOI 10.1007/s12530-018-9244-x

Drakopoulos G, Liapakis X, Tzimas G, Mylonas P (2018b) A graph resilience metric based on paths: Higher order analytics with GPU. In: ICTAI, IEEE

Drakopoulos G, Stathopoulou F, Kanavos A, Paraskevas M, Tzimas G, Mylonas P, Iliadis L (2019) A genetic algorithm for spatiosocial tensor clustering: Exploiting TensorFlow potential. Evolving Systems DOI 10.1007/s12530-019-09267-8

Kanavos A, Drakopoulos G, Tsakalidis A (2017) Graph community discovery algorithms in Neo4j with a regularization-based evaluation metric. In: WEBIST

Kosba A, Miller A, Shi E, Wen Z, Papamanthou C (2016) Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In: IEEE symposium on security and privacy, IEEE, pp 839–858

Matthews DG, et al. (2017) GPflow: A Gaussian process library using TensorFlow. The Journal of Machine Learning Research 18(1):1299–1304

Miller D (2018) Blockchain and the Internet of Things in the industrial sector. IT Professional 20(3):15–18

Nakamoto S (2008) Bitcoin: A peer-to-peer electronic cash system

Pass R, Seeman L, Shelat A (2017) Analysis of the blockchain protocol in asynchronous networks. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, pp 643–673

Puthal D, Malik N, Mohanty SP, Kougianos E, Yang C (2018) The blockchain as a decentralized security framework. IEEE Consumer Electronics Magazine 7(2):18–21

Swan M (2015) Blockchain: Blueprint for a new economy. O'Reilly Media, Inc.

Underwood S (2016) Blockchain beyond Bitcoin. Communications of the ACM 59(11):15–17

Wongsuphasawat K, et al. (2018) Visualizing dataflow graphs of deep learning models in TensorFlow. Transactions on visualization and computer graphics 24(1):1–12

Zyskind G, Nathan O, et al. (2015) Decentralizing privacy: Using blockchain to protect personal data. In: SPW, IEEE, pp 180–184