(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
24 December 2020 (24.12.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/255158 A1

(51) International Patent Classification:
G06F 21/10 (2013.01) *G06F 11/18* (2006.01)

(21) International Application Number:
PCT/IN2020/050501

(22) International Filing Date:
05 June 2020 (05.06.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
201941024497 20 June 2019 (20.06.2019) IN

(71) Applicant: **INDIAN INSTITUTE OF TECHNOLOGY MADRAS (IIT MADRAS)** [IN/IN]; The Dean, Industrial Consultancy & Sponsored Research (IC & SR), Indian Institute of Technology Madras, IIT Post, Chennai, Tamil Nadu 600 036 (IN).

(72) Inventors: **JANAKIRAM, Dharanipragada**; Department of Computer Science and Engineering, IIT Madras, Chennai, Tamil Nadu 600 036 (IN). **GUPTA, Himanshu**; Department of Computer Science and Engineering, IIT Madras, Chennai, Tamil Nadu 600 036 (IN).

(74) Agent: **KUMAR, G. Arun** et al.; K&S Partners, Intellectual Property Attorneys, Door No. 15 (Old No. 3) Postal Colony 4th Street, West Mambalam, Chennai, Tamil Nadu 600 033 (IN).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) Title: COLOSSEUM, A SCALABLE AND BYZANTINE-FAULT TOLERANT CONSENSUS MECHANISM FOR PERMISSIONED SYSTEMS

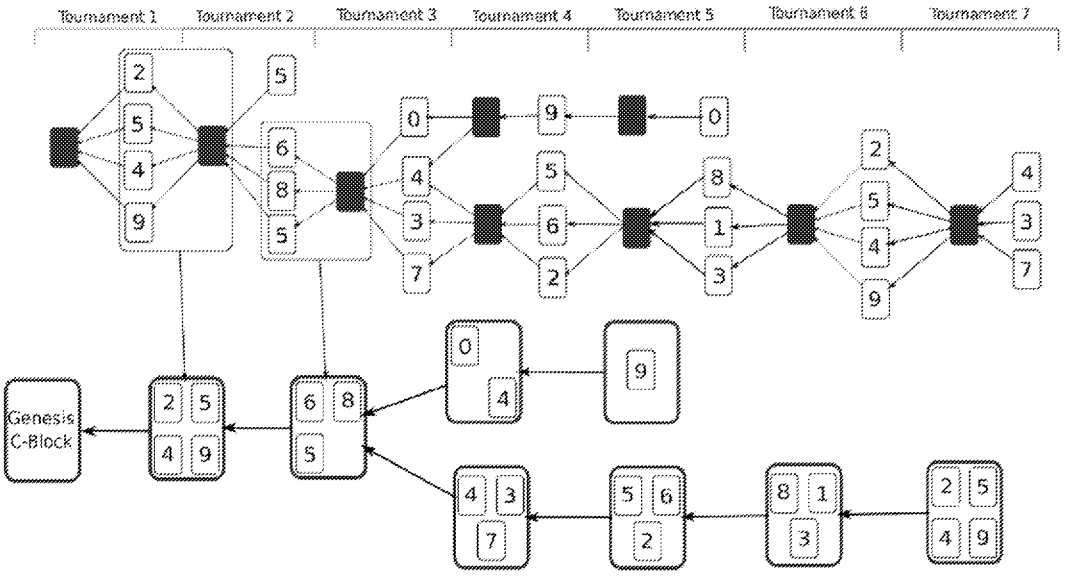


Figure 3: Progress of CDAG for a number of tournaments

(57) Abstract: This invention provides Colosseum, a Byzantine fault tolerant scalable method to achieve consensus, designed to maintain distributed ledgers over a permissioned peer-to-peer blockchain network that follows a curative approach towards achieving consensus. It assumes that adversaries are not always present in the system but efficiently detects malicious behaviour and alert the network. Colosseum follows a knockout tournament like approach over a structured ring network to have agreement on the next set of block proposers. Colosseum proposes a novel two-player game like approach which uses Proof-of-Win to eliminate nodes in each round of a tournament resulting in the selection of a subset of nodes after α ($\alpha < \log 2\ N$, where N is the number of nodes in the network) number of rounds. It allows multiple block proposers for a tournament and tries to commit a maximum of their blocks using Converging Directed Acyclic Graph (CDAG) as the ledger to increase the throughput of the transactions.

WO 2020/255158 A1

[Continued on next page]

SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) **Designated States** *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *of inventorship (Rule 4.17(iv))*

Published:

— *with international search report (Art. 21(3))*

— *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

COLOSSEUM, A SCALABLE AND BYZANTINE-FAULT TOLERANT CONSENSUS MECHANISM FOR PERMISSIONED SYSTEMS

FIELD OF INVENTION

This invention is in the field of block chain technology. Particularly, Colosseum, a method to achieve consensus that can scale, and maintain distributed ledgers over a permissioned peer-to-peer blockchain network with Byzantine fault tolerance.

BACKGROUND OF THE INVENTION

The following on the background of the invention is intended only to facilitate an understanding of the present invention.

Byzantine failures are considered the most general and most difficult class of failures among the failure modes. The so-called fail-stop failure mode occupies the simplest end of the spectrum. Whereas fail-stop failure mode simply means that the only way to fail is a node crash, detected by other nodes, Byzantine failures imply no restrictions, which means that the failed node can generate arbitrary data, pretending to be a correct one. Thus, Byzantine failures can confuse failure detection systems, which makes fault tolerance difficult. Despite the analogy, a Byzantine failure is not necessarily a security problem involving hostile human interference: it can also arise from electrical faults

In a Byzantine fault, a component such as a server can inconsistently appear both failed and functioning to failure-detection systems, presenting different symptoms to different observers. It is difficult for the other components to declare it failed and shut it out of the network, because they need to first reach a consensus regarding which component has failed in the first place. 'Byzantine fault tolerance' is the ability to defend against failures of system components with or without symptoms that prevent other components of the system from reaching an agreement among themselves, where such an agreement is needed for the correct operation of the system. Remaining correctly operational components of a Byzantine fault tolerant system will be able to continue providing the system's service as originally intended, assuming there are sufficiently many accurately operating components to maintain the service.

In US 9,807,106 B2 titled" MITIGATING BLOCKCHAIN ATTACK" discloses a computer implemented method for detecting malicious attacks presenting a threat to a blockchain associated with a blockchain data structure of a computing device. Again, it is complicated and had limited applications.

In general, permissioned blockchains use methods based on voting to reach consensus. They work on an assumption that adversaries are *always present* in the system and follows a preventive approach of having a majority vote on each batch of transactions added to the ledger. Achieving scalability for such methods is hard because of their high message complexity and thus, have a limited number of applications. Use cases like banking and healthcare blockchain requires a scalable design to maintain the consistency of the ledger and ensure its safety.

There are various researches on-going to have a scalable method, which is simple and have wide range of applications. The shortcomings of the prior art are overcome, and additional advantages are provided through the invention and method of present disclosure.

OBJECT OF THE INVENTION

The main object of the invention is to provide a Byzantine fault tolerant, scalable method to achieve consensus, designed to maintain distributed ledgers over a permissioned peer-to-peer blockchain network.

SUMMARY OF THE INVENTION:

This invention provides Colosseum, a Byzantine fault tolerant scalable method to achieve consensus, designed to maintain distributed ledgers over a permissioned peer-to-peer blockchain network that follows a curative approach towards achieving consensus. It assumes that adversaries are not always present in the system but efficiently detects malicious behaviour and alert the network. Colosseum follows a knockout tournament like approach over a structured ring network to have agreement on the next set of block proposers. Colosseum proposes a novel two-player game like approach which uses Proof-of-Win to eliminate nodes in each round of a tournament resulting in the selection of a subset of nodes after α (α < log 2 N, where N is the number of nodes in the network) number of rounds. It allows multiple block proposers for a tournament and tries to commit

a maximum of their blocks using Converging Directed Acyclic Graph (CDAG) as the ledger.

BRIEF DESCRIPTION OF THE DRAWINGS AND FIGURES

It is to be expressly understood, that each of the following figures is provided for the purpose of illustration and description only and is not intended as a definition of the limits of the present disclosure.

FIGURE:1 proof of win for a match

FIGURE:2 depicts message flow of a match in Colosseum A, B are the players, V_{AB} is the validator, and K_A, K_B are the Keepers.

FIGURE:3 depicts the progress of CDAG for the number of tournaments.

FIGURE:4 throughput for α=3

FIGURE:5 Throughput for α=4

DETAILED DESCRIPTION

The foregoing has broadly outlined the features and technical advantages of the present disclosure in order that the detailed description that follows may be better understood. Additional features and advantages of the disclosure will be described herein after which form the subject of the claims of the disclosure. It should be appreciated by those skilled in the art that the conception and specific aspect disclosed may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present disclosure. It should also be realized by those skilled in the art that such equivalent constructions do not depart from the spirit and scope of the disclosure as set forth in the appended claims. The novel features which are believed to be characteristic of the disclosure, both as to its organization and method of operation, together with further objects and advantages will be better understood from the following description when considered in connection with the accompanying figures. The novel features and characteristic of the disclosure are set forth in the appended claims. The embodiments of the disclosure itself, however, as well as a preferred mode of use, further objectives and

advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the following details of the accompanying drawings. One or more embodiments are described in the later part, by way of example, with reference to the accompanying drawings.

All publications, patents and patent applications cited herein, whether supra or infra, are hereby incorporated by reference in their entirety to the same extent as if each individual publication, patent and patent application is specifically and individually indicated to be incorporated by reference.

Unless defined otherwise, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which the invention pertains.

The terms "an embodiment", "embodiment", "embodiments", "the embodiment", "the embodiments", "one or more embodiments", "some embodiments", and "one embodiment" mean "one or more (but not all) embodiments of the invention(s)" unless expressly specified otherwise.

The terms "including", "comprising", “having” and variations thereof mean "including but not limited to", unless expressly specified otherwise. The enumerated listing of items does not imply that any or all of the items are mutually exclusive, unless expressly specified otherwise. The terms "a", "an" and "the" mean "one or more", unless expressly specified otherwise.

The present invention describes a method Colosseum, a scalable consensus technique designed for permissioned blockchains. At the core, it uses synchronous knockout approach over a Distributed Hash Table (DHT) [1] to agree on the next set of block proposers. To propose a block, participants play matches of a novel two-player game to eliminate others and qualify. It introduces Proof-of-Win (PoWin) as a verifiable, tamper-proof certificate for the result of matches and use it to store the state of nodes over the DHT.

Permissioned blockchains generally use methods based on voting to reach consensus. They work on an assumption that adversaries *are always present* in the system and follows a preventive approach of having a majority vote on each batch of transactions added to the

ledger. Achieving scalability for such methods is hard because of their high message complexity and thus, have a limited number of applications. Use cases like banking and healthcare blockchain requires a scalable design to maintain the consistency of the ledger and ensure its safety.

Colosseum introduces Converging Directed Acyclic Graph (CDAG) as an alternative to the blockchain and blockDAG structures to commit multiple blocks of transactions simultaneously to the ledger and converge them at each step providing finality to the blocks.

Moreover, Colosseum uses bucketing of transactions to prevent simultaneous proposers from generating conflicting blocks and introduces a distributed time barrier to grant enough time for the blocks to disseminate in the network.

A tournament in Colosseum consists of log 2 N (N is the number of participants) asynchronous rounds. Nodes play matches in pair of two for each round and accept the winners of α (α < log 2 N) number of rounds as the proposers for that tournament. The elected proposers can then propose a block for the blockchain and propagate them in the network along with their PoWin.

Colosseum elects multiple block proposers in each tournament to propose blocks for the ledger. It divides the transactions into multiple buckets based on their hashes to enable the proposers to generate non-conflicting blocks from different buckets. It also introduces Converging Directed Acyclic Graph (CDAG) as the structure for the ledger to support the addition of multiple simultaneous blocks and helps them converge at the next step providing finality to the blocks. CDAG is a hybrid structure of blockchain [2] and blockDAG [3], designed to achieve the serialization of blockchain (to progress as a chain) and the parallel nature of blockDAG (to enable addition of multiple concurrent blocks to the ledger). Moreover, Colosseum also introduces a Proof of Elapsed Time [4] based distributed time barrier to divide the timeline into slots. The barrier makes the nodes to individually wait for a particular time after the start of a tournament to prevent the nodes from moving ahead of others. It helps CDAG to converge better at each step as blocks get adequate time to disseminate in the network.

Accordingly, the principal embodiment of this invention is a scalable and Byzantine fault tolerant method to reach consensus on the next set of block proposers in permissioned blockchain systems comprising the steps of;

a. forming pair of nodes and generating their respective block proposals,
b. validating the game proposals from the involved nodes of step (a) by Validator, and generating proof of win (PoWin) certificate as result, and
c. storing the PoWin from step (b) by Keeper and verifying that the nodes are forming authorized pairs.

Another embodiment of this invention is the Pairing of nodes in step (a) is not absolute, pairing is on the fly to avoid targeted attacks by malicious opponents.

Yet another embodiment of this invention is for pairing in step (a), the nodes send TCP requests to other nodes on the ring to pair and provide the PoWin of the previous round as a token of their eligibility.

In yet another embodiment of this invention, in step (a) of this method after pairing, the confirmation nodes generate their block proposal by signing common string S_{ABtr} with their (digital signature) private key p_k to generate *proposed string* (*PS*).

A further embodiment of this invention, in step (b) of this method *Validator* compares the *PS* of both the nodes with its signed S_{ABtr} and the node with the value closer to that of the validator earns the PoWin.

In still another embodiment of the present invention the validator is a random third party in the network uniquely identifiable for a match.

In another embodiment of this invention the validation process is delegated to prevent the nodes from leaving the method in-between.

Yet another embodiment of this invention, in step (c) of the method the 'Keeper' stores the PoWin as verifiable and tamper-proof state of users over a Distributed Hash Table (DHT).

In an embodiment of this invention the DHT restricts the nodes (already lost or not eligible) from pairing unauthorized matches; as the verification process which is

independent of the nodes will identify any wrongful attempt and notify with verifiable proofs.

A further embodiment of this invention is the 'Keepers' are different for every pair to evade any malicious or attack by adversary.

In a further embodiment of this method in step (c) for verification the keeper will;

a. query the previous round keepers (multiple replicas) of the nodes included in a PoWin by keepers (all replicas) of a match,
b. ensure both the nodes pairing in r_{th} round must have won earlier $(r-1)_{th}$ round,
c. querying keeper collecting the PoWin for round *r* as a proof for a valid query (i.e., node has played the r_{th} round), and the queried keeper reply with the PoWin of round *r-1* to confirm the previous victory of the nodes.
d. In the absence of proof with the queried keeper and query is valid, it broadcasts a negative vote against the pair and queries the previous keeper using the same proof.

wherein, the match is considered as a *foul* when a majority (> 2/3rd of keepers) of negative votes from valid keepers are against it.

In a preferred embodiment of the present invention, a new structure to store blocks in a distributed ledger across peer-to-peer networks proposed by the block proposers selected using the method is Converging Directed Acyclic Graph (CDAG).

Yet another embodiment of this invention, the structure progresses by adding (Converging block) *C-Block* (Fig.3) to the structure of blockchain to give the ability to commit multiple blocks simultaneously and by bucketing of transactions to minimize the conflicts.

Still another embodiment, the *C-Block* is the converging point in CDAG allowing the current block proposers to refer multiple blocks with a single link and enables different proposers to have the same view about the ledger.

As an embodiment of this invention, the C-Block is characterised by following set of rules:

a. belong to different buckets, to ensure that a C-Block does not contain conflicting blocks due to the presence of the same transactions;

b. have non-conflicting transactions, to avoid any double spending transactions in a C-Block, and
c. refer to the same previous C-Block, providing the serial nature to CDAG.

As an embodiment to the structure, the number of qualifying proposers is inversely proportional to α.

In a further embodiment to the structure, for larger values of α the saturation point shifts to a higher number of nodes, allowing CDAG to *scale* with the number of nodes.

The invention is now illustrated by various examples and accompanying drawings, which are not meant to limit the scope of the invention in any manner. All embodiments that may be obvious to a skilled person would fall within the scope of the present invention.

EXAMPLES

EXAMPLE 1

1. KNOCKOUT TOURNAMENTS

Tournaments in Colosseum are the basis for the selection of block proposers. Users participate in these tournaments to qualify and propose blocks for the ledger. A tournament consists of multiple asynchronous rounds. Users play matches of a two-player game for each round to win and move forward and receive PoWin as the certificate of their win. Nodes with valid PoWin advance to the next round and the winners of $\alpha < \log(N)$ (N is the number of participants) number of rounds are eligible to propose a block in Colosseum.
The two-player game plays a vital role in maintaining the integrity of the system. It is a process to select a random winner amongst two participants and rely on random nodes in the network for validation and verification. The randomness of the game comes from the cryptography involved to generate the game proposals and Ids of the third-party nodes on the ring. The dependency on random nodes in the network prevents a player to defraud others and ensures fairness among the players.

1.1 Proof-of-Win

PoWin is a cryptographically secured certificate designed to maintain the state of users on a peer-to-peer network. It is tamper-proof as it contains the digitally signed game proposals of both the nodes involved in a match, as shown in Figure 1. PoWin is a single

certificate generated as the result of matches in Colosseum and gets distributed over the network. Winners of a round use their PoWin as a ticket to play the next round, whereas it works as a restriction for the losers and prevents them from moving ahead.

PoWin proves the occurrence of a match. A node on receiving conflicting PoWins of a user alert others with verifiable proofs (multiple PoWins for the same round or PoWin for a higher round of a player already lost in previous rounds).

Since the cryptography involved does not allow to generate fake PoWin certificates, it becomes evident that the user itself has performed such a task. Peers discard blocks proposed by these nodes and individually block them for a certain number of tournaments on multiple violations. Moreover, nodes are blocked on the basis of proofs available with the peers and therefore does not require a consensus between nodes.

1.2 Two Player Game

Colosseum uses a two-player game to eliminate nodes participating in the tournaments and select the final set of block proposers. Each node in the network performs *three* roles during a tournament.

- A *Player* forms a pair with an opponent and generate their game proposals.
- A *Validator* validates the game proposal from the involved players and generate PoWin certificate as a result, and
- A *Keeper,* responsible for storing the PoWin and verifying that the players are playing authorized matches.

1.2.1. Finding Competitors: Pairing for the tournaments is not absolute in Colosseum. Participants find their opponents on the fly to avoid targeted attacks by malicious opponents. They send TCP requests to other nodes on the ring to play and provide the PoWin of the previous round as a token of their eligibility. Eligible nodes respond with their PoWin while others send a negative response.

1.2.2. Generating game proposals: After pairing confirmation nodes generate their *Game Proposal*. Both the players digitally sign a common string S_{ABtr} with their private key p_k to generate *proposed string* (*PS*), refer equation (2). Validator compares *PS* of both the opponents with its signed S_{ABtr} and the player with the value closer to that of the validator wins.

$$S_{ABtr} = \text{concat}\ (I_A, I_B, t, r) \quad (1)$$

$$PS = \text{sign}\ (p_k, S_{ABtr}) \quad (2)$$

Here, I_A and I_B are the Ids of the players (I_A lesser than I_B), and t and r is the tournament and round number respectively. Digital signatures used to generate the proposals must be deterministic to ensure safety of the technique.

Moreover, the Public Key infrastructure should not be compromised, i.e., the players should not have multiple private keys to generate various game proposals which can result in a probabilistic advantage. It also prevents the validators from using different keys that can affect the outcome of a match.

1.2.3. Game Validation: Colosseum delegates the validation and result computation of matches to a random third-party on the network known as *validator*.

Validator is a randomly selected participant on the ring network uniquely identifiable for a match. Its responsibility is to validate the game proposals sent by the players involved, produce PoWin and forward it to the contestants and their keepers, as shown in Figure 2.

Delegating the validation process prevents the competitors from leaving matches in-between, foreseeing the result. Selecting a random node to validate ensures that the game is fair unless the validator is biased. Moreover, the outcome of a match is verifiable; therefore, a malicious validator can at most drop a match if an ally is losing. In such a case, nodes not receiving results from the validator find other eligible contenders to play after a particular time.

$$V_{AB} = \text{hash}\,(S_{ABtr}) \qquad (3)$$

Validator Id V_{AB} depends on S_{ABtr}, as shown in equation (3). Dynamic pairing at run time influence S_{ABtr} and therefore prevent targeted attacks on the validator. Using a common string S_{ABtr} also restricts a player to select a validator of choice and hence provides a fair chance to both the players.

1.2.4. Storing Result: Storing PoWin with the *keeper* is also an important design decision. *Keeper* manages the verifiable and tamper-proof state of users over a DHT. It restricts the participants (already lost or not eligible) from playing unauthorized matches; as the verification process which is independent of the players will identify any wrongful attempt by the players and notify it to others with verifiable proofs.

$$K_A = \text{hash}\,(\text{concat}\,(I_A, t, r)) \qquad (4)$$

DHT root for the hash K_A is considered as the keeper of a match, as shown in equation (4). To allow independent state verification (i.e., verifying the state of a player irrespective of

its competitor) K_A only depends on the Id of the player and the match information. Also, the *keepers* are different for every match because the same *keeper* can become a liability if malicious or can get attacked by some adversary.

Furthermore, a malicious or faulty *keeper* can result in loss of data and match verification can get affected. Thus, a PoWin is stored at multiple *keepers* evenly distributed over the ring to handle faults. Nodes can query any number of replicas to verify the state of a node for a match.

1.2.5. State Verification in Tournament Tree

Colosseum verifies all the matches of a tournament to make sure that only winning participants of a round proceed towards higher rounds. Its goal is to double-check the entire tournament tree in a top-down approach and helps to ensure that nodes play honestly and identify any unauthorized match played in a tournament.

To verify matches, keepers (all replicas) of a match query one or more previous round keepers (multiple replicas) of the players included in a PoWin. It is important because both the nodes playing r_{th} round must have won earlier $(r-1)_{th}$ round. The querying keeper provides the PoWin for round *r* as a proof for a valid query (i.e., node has played the r_{th} round), and the queried keeper reply with the PoWin of round *r-1* to confirm the previous victory of the players. In case, the proof is not available with the queried keeper and query is valid, it broadcasts a negative vote against the match and queries the previous keeper using the same proof. A match is considered as a *foul* when a majority (> 2/3rd of keepers) of negative votes from valid keepers are against it.

EXAMPLE 2

CDAG (Converging Block directed acyclic graph)

A directed acyclic graph (**DAG**), is a finite directed graph with no directed cycles. That is, it consists of finitely many vertices and edges (also called *arcs*), with each edge directed from one vertex to another, such that there is no way to start at any vertex *v* and follow a consistently-directed sequence of edges that eventually loops back to *v* again. Equivalently, a DAG is a directed graph that has a topological ordering, a sequence of the vertices such that every edge is directed from earlier to later in the sequence.

DAGs can model many different kinds of information. For example, a spreadsheet can be modeled as a DAG, with a vertex for each cell and an edge whenever the formula in one cell uses the value from another; a topological ordering of this DAG can be used to update all cell values when the spreadsheet is changed.

Converging Directed Acyclic Graph (CDAG) is introduced as a new structure to maintain distributed ledgers across peer-to-peer networks. It shares the same trust model as the traditional blockchain (blocks of transactions linked to each other progressing as a modification resistant ledger) and is designed to provide higher throughput. It adds *C-Block* as a new construct to the structure of blockchain to give the ability to commit multiple blocks simultaneously and use bucketing of transactions to minimize the conflicts. The structure progresses as a chain of C-Blocks as shown in (**Figure 3**). C-Block provides finality to the locks added in CDAG and serialize the structure to support smart contracts.

EXAMPLE 3

C-Block

C-Block (Converging block) is the converging point in CDAG. It is proposed as an extra layer between two layers of blocks providing a point of convergence to the blocks of the previous slot and a single point of reference for all the blocks of the current slot, shown in Figure 3. As multiple blocks are proposed in each slot, a C-Block refers multiple non-conflicting blocks and contains the list of hashes of all the previous blocks it refers. A block consistent with a C-Block is bound to be consistent with all the blocks included in it. It allows the current block proposers to refer multiple blocks with a single link and enables different proposers to have the same view about the ledger.

Creation of C-Block is also the responsibility of the block proposers. While proposing a block, a proposer needs to make it consistent with a C-Block and includes a pointer to it. To do so a proposer can either generate a new C-Block or use an existing C-Block created and propagated by other selected proposers of the same tournament. In case of a conflict the heaviest chain rule is used to select the C-Block with greater probability to survive.

Generating a valid C-Block is a trivial process and obeys a specific set of rules. Proposers need to make sure that all the blocks included in a C-Block must:

1) belong to different buckets, to ensure that a C-Block does not contain conflicting blocks due to the presence of the same transactions;

2) have non-conflicting transactions, to avoid any double spending transactions in a C-Block, and

3) refer to the same previous C-Block, providing the serial nature to CDAG.

Blocks of a tournament with different previous C-Blocks are similar to a fork in the traditional blockchain. In figure 3 the block of bucket 9 points to a different C-Block than that of buckets 5, 6, 2 at tournament 4, i.e., the previous blocks of both the C-Blocks are different, creating a fork. Thus, the structure progresses as a chain where a C-Block can be viewed as a large block encapsulating all the blocks included in its list.

RESULTS

To evaluate Colosseum, we deploy it on Google Cloud Platform using Google Kubernetes Engine having n1-highmem-16 instances, each with 16 cores and 104 GB memory. Multiple Docker containers each running as a user of the structured layer run on these instances with a cap of 25 Mbps on their bandwidth. Users proposed 1 MByte blocks each with 2500-3000 transactions. We ran each experiment for at least 30 tournaments, and graphs show an average over all of them. We ran the experiments for three different configurations varying the block size and the time slot of the barrier, as shown in table 1.

Table 1: Configurations to test Colosseum

Configuration	Block Size	Time slot
Config 1	1 MB	20s
Config 2	0.75 MB	15s
Config 3	0.50 MB	10s

We vary the number of nodes in network for multiple values of α and all the three configurations as mentioned in table 1. Both the figures 4, and 5 show that for a fixed value of α the throughput of the system increases initially with the number of nodes and then starts to decrease. The initial increase is because the number of blocks for a tournament of Colosseum is directly proportional to the number of nodes. More number of nodes result in a higher degree of parallelism for CDAG and thus, increases the throughput. After a certain point, the throughput starts decreasing for all the different configurations in 4 and 5, due to the generation of a large number of simultaneous blocks. Bandwidth constraints and network congestion do not provide them enough time to reach everyone in the network. Also, it is visible from both the graphs that the network with

smaller time slots saturates faster in comparison to that of a higher time slot even though the size of the block also increases in the same ratio.

Figure 4 and 5 also shows that for a fixed number of nodes, increasing the value of α results in lower throughput since the number of qualifying proposers is inversely proportional to α. However, for larger values of α the saturation point shifts to a higher number of nodes. It allows CDAG to *scale* with the number of nodes. The value of α can be manipulated to achieve the optimum throughput for a given configuration similar to setting the hardness of the cryptographic puzzle in Bitcoin.

Finally, the language used in the specification has been principally selected for readability and instructional purposes, and it may not have been selected to delineate or circumscribe the inventive subject matter. It is therefore intended that the scope of the invention be limited not by this detailed description, but rather by any claims that issue on an application based here on. Accordingly, the disclosure of the embodiments of the invention is intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the following claims.

The scope of the present invention includes any novel features or combination of features disclosed herein. The applicant hereby gives notice that new claims may be formulated to such features or combination of features during prosecution of this application or of any such further applications derived therefrom. In particular, with reference to the appended claims, features from dependent claims may be combined with those of the independent claims and features from respective independent claims may be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

REFERENCES

[1] Wikipedia.org. 2018. Distributed Hash Table. Retrieved 2018-17-07 from https://en.wikipedia.org/wiki/Distributed_hash_table

[2] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf

[3] Alexandra Tran. 2018. An Introduction to the BlockDAG Paradigm. Retrieved 2018-11-26 from https://blog.daglabs.com/an-introduction-to-the-blockdag-paradigm-50027f44facb

[4] Hyperledger.org. 2018. PoET-Proof of Elapsed Time. Retrieved 2018-11-11 from https://sawtooth.hyperledger.org/docs/core/releases/latest/architecture/poet.html

The claim:

1. A scalable and Byzantine fault tolerant method to reach consensus on the next set of block proposers in permissioned blockchain systems comprising the steps of;
    a. forming pair of nodes and generating their respective block proposals,
    b. validating the game proposals from the involved nodes of step (a) by Validator, and generating proof of win (PoWin) certificate as result, and
    c. storing the PoWin from step (b) by Keeper and verifying that the nodes are forming authorized pairs.

2. The method as claimed in claim 1, wherein step (a) Pairing of nodes is not absolute, pairing is on the fly to avoid targeted attacks by malicious opponents.

3. The method as claimed in claim 2, wherein step (a) for pairing, the nodes send TCP requests to other nodes on the ring to pair and provide the PoWin of the previous round as a token of their eligibility.

4. The method as claimed in claim 1, wherein step (a) after pairing, the confirmation nodes generate their block proposal by signing common string S_{ABtr} with their (digital signature) private key p_k to generate *proposed string* (*PS*).

5. The method as claimed in claim 1, wherein step (b) Validator compares the *PS* of both the nodes with its signed S_{ABtr} and the node with the value closer to that of the validator earns the PoWin.

6. The method as claimed in claim 5, wherein the validator is a random third party in the network uniquely identifiable for a match.

7. The method as claimed in claim 5, wherein the validation process is delegated to prevent the nodes from leaving the method in-between.

8. The method as claimed in claim 1, wherein step (c) Keeper stores the PoWin as verifiable and tamper-proof state of users over a Distributed Hash Table (DHT).

9. The method as claimed in claim 8, wherein said DHT restricts the nodes (already lost or not eligible) from pairing unauthorized matches; as the verification process which is independent of the nodes will identify any wrongful attempt and notify with verifiable proofs.

10. The method as claimed in claim 8, wherein said keepers are different for every pair to evade any malicious or attack by adversary.

11. The method as claimed in claim 1, wherein step (c) for verifying the keeper will
    a. query the previous round keepers (multiple replicas) of the nodes included in a PoWin by keepers (all replicas) of a match,
    b. ensure both the nodes pairing in r_{th} round must have won earlier $(r-1)_{th}$ round,
    c. querying keeper collecting the PoWin for round *r* as a proof for a valid query (i.e., node has played the r_{th} round), and the queried keeper reply with the PoWin of round *r-1* to confirm the previous victory of the nodes.
    d. In the absence of proof with the queried keeper and query is valid, it broadcasts a negative vote against the pair and queries the previous keeper using the same proof.

    wherein, the match is considered as a *foul* when a majority (> 2/3rd of keepers) of negative votes from valid keepers are against it.

12. A new structure to store blocks in a distributed ledger across peer-to-peer networks proposed by the block proposers selected using the method of claim 1, is Converging Directed Acyclic Graph (CDAG).

13. The structure as claimed in claim 12, wherein said structure progresses by adding (Converging block) *C-Block* (Fig.3) to the structure of blockchain to give the ability to commit multiple blocks simultaneously and by bucketing of transactions to minimize the conflicts.

14. The structure as claimed in claim 13, wherein said *C-Block* is the converging point in CDAG allowing the current block proposers to refer multiple blocks with a single link and enables different proposers to have the same view about the ledger.

15. The structure as claimed in claim 13, wherein said C-Block is characterised to follow the following set of rules:
    a. belong to different buckets, to ensure that a C-Block does not contain conflicting blocks due to the presence of the same transactions;
    b. have non-conflicting transactions, to avoid any double spending transactions in a C-Block, and
    c. refer to the same previous C-Block, providing the serial nature to CDAG.

16. The structure as claimed in claim 13, wherein the number of qualifying proposers is inversely proportional to α.

17. The structure as claimed in claim 13, wherein for larger values of α the saturation point shifts to a higher number of nodes, allowing CDAG to *scale* with the number of nodes.

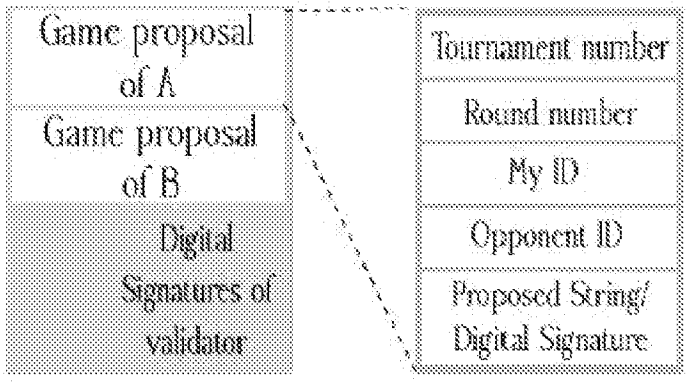


Figure 1: Proof-of-Win for a match

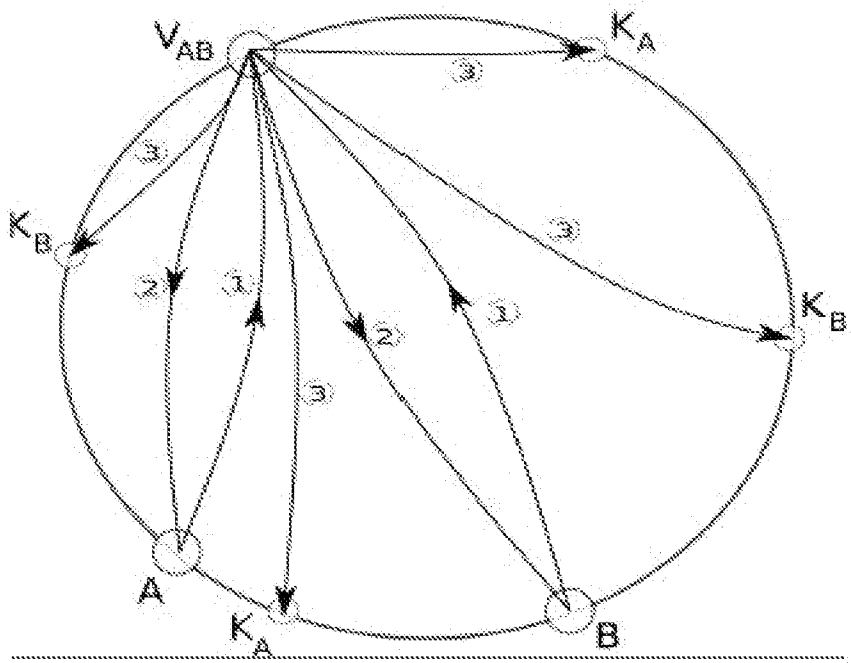


Figure 2: Message flow of a match in Colosseum. A,B are the players, V_{AB} is the validator, and K_A, K_B are the keepers

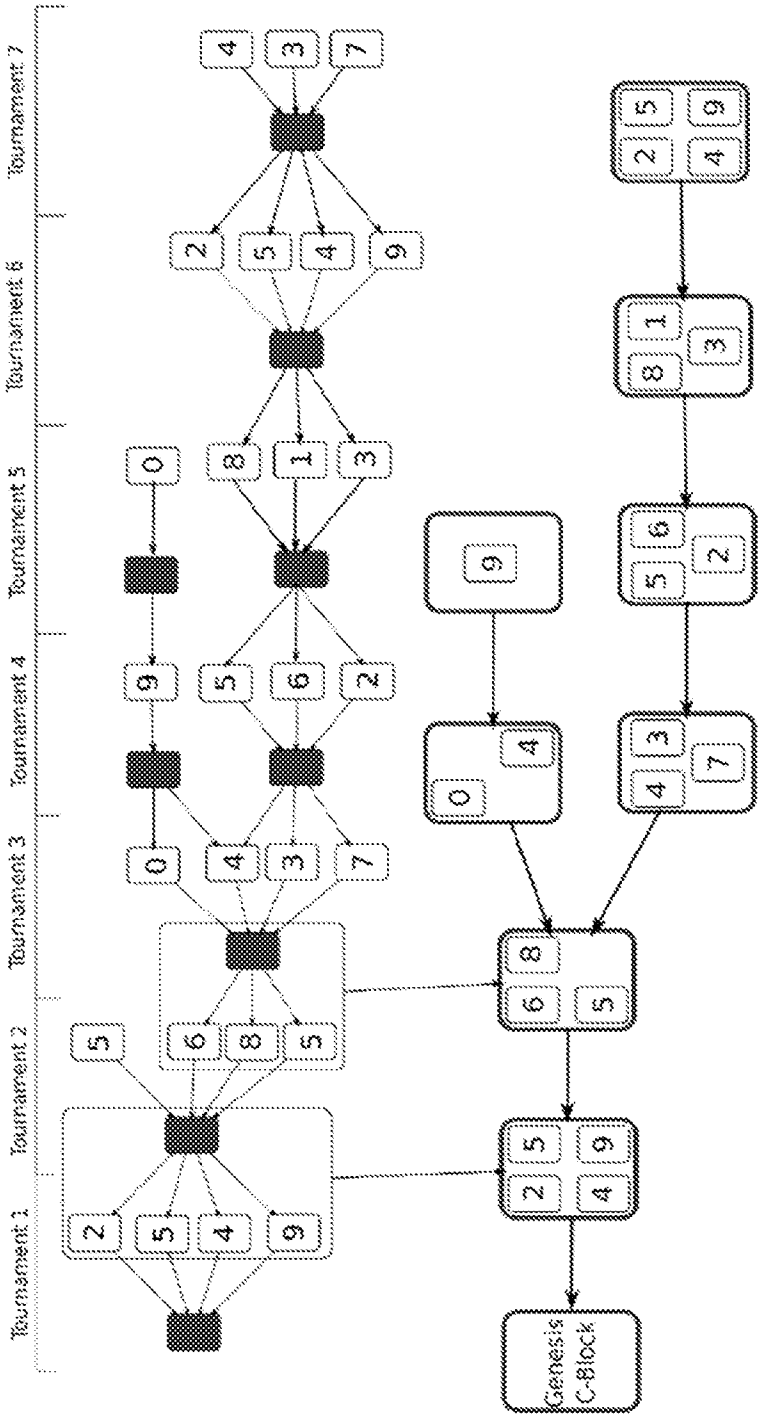


Figure 3: Progress of CDAG for a number of tournaments

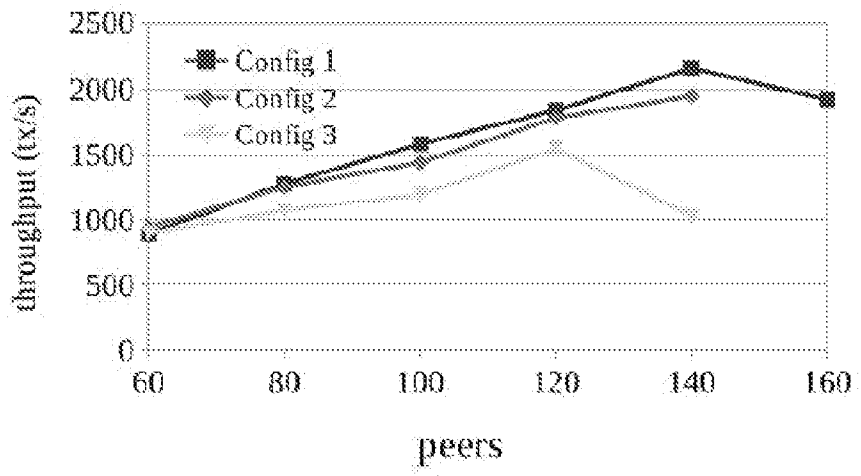


Figure 4: Throughput for $\alpha = 3$

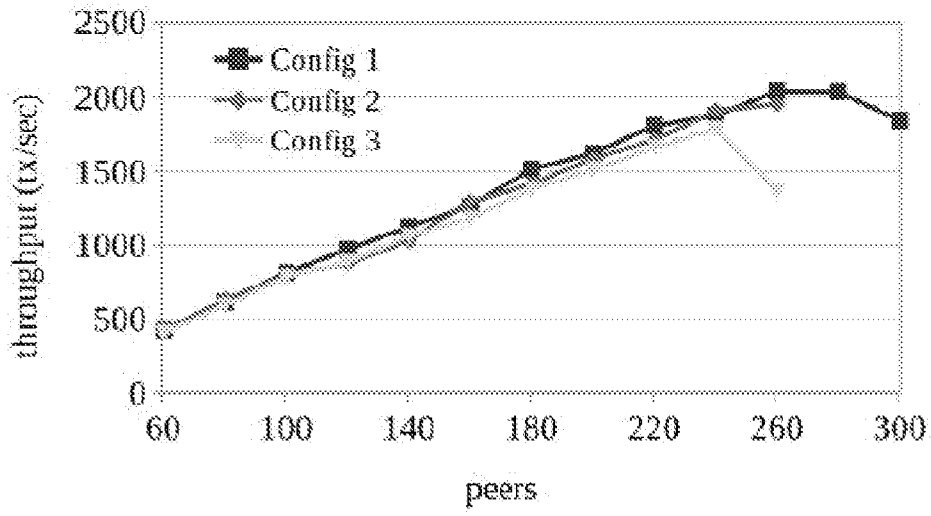


Figure 5: Throughput for $\alpha = 4$

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IN2020/050501

A. CLASSIFICATION OF SUBJECT MATTER
G06F21/10,G06F11/18 Version=2020.01

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F, H04L, G06Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Databases: TotalPatent One, IPO Internal Database
Keywords : byzantinefault,tolerance,scalable,DAG,blockchain,ledger,verify

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US20180341930A1 (ORACLE INT CORP), 29.11.2018 (29 November 2018) [see the whole document]	1-17
Y	Bencic, Federico M. Et Al, "Distributed Ledger Technology: Blockchain Compared to Directed Acyclic Graph", 2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS) 23.07.2018 (23 July 2018) [see the whole document]	1-17
Y	US20180005186A1 (CLAUSE INC), 04.01.2018 (04 January 2018) [see paras. 67, 117-118, 173, 181, 231-240, 331; figure 20]	3, 4, 7, 9

☐ Further documents are listed in the continuation of Box C. ☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"D" document cited by the applicant in the international application
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
23-09-2020	23-09-2020
Name and mailing address of the ISA/ Indian Patent Office Plot No.32, Sector 14,Dwarka,New Delhi-110075 Facsimile No.	Authorized officer Rakesh Kumar Telephone No. +91-1125300200

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IN2020/050501

Citation	Pub.Date	Family	Pub.Date
US 20180341930 A1	29-11-2018	CN 110800005 A	14-02-2020
		EP 3631722 A1	08-04-2020
		WO 2018217511 A1	29-11-2018
		JP 2020521252 A	16-07-2020
US 20180005186 A1	04-01-2018	WO 2018006072 A1	04-01-2018