



US011188977B2

(12) **United States Patent**
Youb et al.

(10) **Patent No.:** **US 11,188,977 B2**
(45) **Date of Patent:** **Nov. 30, 2021**

(54) **METHOD FOR CREATING COMMODITY ASSETS FROM UNREFINED COMMODITY RESERVES UTILIZING BLOCKCHAIN AND DISTRIBUTED LEDGER TECHNOLOGY**

H04L 2209/56; H04L 9/3297; H04L 9/0897; H04L 9/3247; H04L 63/10; H04L 9/0825; H04L 63/0442; G06Q 20/0658; H04W 4/80; H04W 12/001; H04W 12/08
USPC 705/1.1-912
See application file for complete search history.

(71) Applicant: **IP Oversight Corporation**
(72) Inventors: **Christopher Youb**, Edmonton (CA); **Scott Mehlman**, Merrick, NY (US); **Michael Zimits**, Massapequa Park, NY (US); **Bruce Youb**, Puerto Plato (DO)

(56) **References Cited**

U.S. PATENT DOCUMENTS

(73) Assignee: **Stichting IP-Oversight**, Lutjebroek (NL)

6,324,286	B1	11/2001	Lai et al.	
6,938,039	B1	8/2005	Bober et al.	
8,326,720	B2 *	12/2012	Spirgel	G06Q 40/00 705/35
8,449,378	B2	5/2013	Michaelson et al.	
8,453,219	B2	5/2013	Shuster et al.	
8,493,386	B2	7/2013	Burch et al.	
8,522,330	B2	8/2013	Shuster et al.	
8,523,657	B2	9/2013	Michaelson et al.	
8,572,207	B2	10/2013	Shuster et al.	
8,621,368	B2	12/2013	Shuster et al.	
8,626,641	B1 *	1/2014	Merk	G06Q 40/04 705/37
8,671,142	B2	3/2014	Shuster et al.	
8,688,525	B2	4/2014	Minde	
8,719,131	B1	5/2014	Roth et al.	
8,726,379	B1	5/2014	Stiansen et al.	
8,727,893	B2	5/2014	Lee et al.	
8,756,151	B1	6/2014	Lubling et al.	
8,756,156	B1	6/2014	Campi et al.	
8,769,643	B1	7/2014	Ben Ayed	
8,777,735	B1	7/2014	Fine et al.	
8,862,506	B2	10/2014	Smith et al.	
8,905,303	B1	12/2014	Ben Ayed	
8,909,517	B2	12/2014	Mosko et al.	
8,915,781	B2	12/2014	Fine et al.	
8,918,794	B2	12/2014	Kruglick	
8,938,534	B2	1/2015	Le et al.	
8,947,427	B2	2/2015	Shuster et al.	
8,950,004	B2	2/2015	Messinger et al.	
8,961,300	B2	2/2015	Fine	
8,961,301	B2	2/2015	Fine et al.	
8,962,964	B2	2/2015	Emmerson	
8,968,082	B2	3/2015	Fine et al.	
8,968,104	B2	3/2015	Fine et al.	
8,972,612	B2	3/2015	Le et al.	
8,974,284	B2	3/2015	Selby et al.	
8,985,442	B1	3/2015	Zhou et al.	
8,986,096	B2	3/2015	Fine et al.	
8,992,311	B2	3/2015	Fine	
8,992,312	B2	3/2015	Fine et al.	

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 794 days.

(21) Appl. No.: **15/916,128**

(22) Filed: **Mar. 8, 2018**

(65) **Prior Publication Data**

US 2019/0080392 A1 Mar. 14, 2019

Related U.S. Application Data

(60) Provisional application No. 62/468,764, filed on Mar. 8, 2017.

(51) **Int. Cl.**

G06Q 10/10	(2012.01)
G06Q 10/06	(2012.01)
G06Q 30/02	(2012.01)
G06Q 30/06	(2012.01)
G06Q 10/08	(2012.01)
G06Q 20/06	(2012.01)
G06F 21/64	(2013.01)
G06Q 40/00	(2012.01)
G06Q 20/38	(2012.01)
G06Q 20/40	(2012.01)
G06Q 50/16	(2012.01)
H04L 9/32	(2006.01)
G06F 16/18	(2019.01)
G06Q 40/04	(2012.01)
H04L 9/06	(2006.01)

(52) **U.S. Cl.**

CPC **G06Q 30/0645** (2013.01); **G06F 16/1805** (2019.01); **G06F 21/64** (2013.01); **G06Q 20/065** (2013.01); **G06Q 20/38215** (2013.01); **G06Q 20/401** (2013.01); **G06Q 30/02** (2013.01); **G06Q 40/00** (2013.01); **G06Q 40/04** (2013.01); **G06Q 50/16** (2013.01); **H04L 9/0643** (2013.01); **H04L 9/3213** (2013.01); **G06Q 2220/00** (2013.01); **H04L 2209/38** (2013.01)

(58) **Field of Classification Search**

CPC .. G06Q 30/0645; G06Q 20/065; G06Q 30/02; G06Q 40/00; G06Q 2220/00; G06Q 20/401; G06Q 20/3829; G06Q 20/389; G06Q 20/38215; G06Q 50/16; G06Q 40/04; G06F 16/1805; G06F 21/64; H04L 9/0643; H04L 2209/38; H04L 9/3213;

Primary Examiner — Jonathan P Ouellette
(74) Attorney, Agent, or Firm — Hoffberg & Associates; Steven M. Hoffberg

(57) **ABSTRACT**

A method for creating an asset-backed distributed ledger token representing a smart contract, the token being backed by a pledge of an illiquid form of a precursor or means of production of a commodity asset, comprising receiving a pledge of unrefined or pre-commodity asset, digitizing the unrefined commodity asset into fractional representations of the commodity asset using smart contracts on a distributed ledger network, and allowing account holders access to the perform transactions on the distributed ledger network to trade the fractional representations under the terms of the smart contract.

(56)

References Cited

U.S. PATENT DOCUMENTS

9,014,661	B2	4/2015	deCharms	2017/0109637	A1	4/2017	Marcu et al.
9,015,847	B1	4/2015	Kaplan et al.	2017/0109638	A1	4/2017	Marcu et al.
9,020,110	B1	4/2015	Baharav et al.	2017/0109639	A1	4/2017	Marcu et al.
9,027,007	B2	5/2015	Brackman et al.	2017/0109640	A1	4/2017	Marcu et al.
9,747,586	B1 *	8/2017	Frolov	2017/0109657	A1	4/2017	Marcu et al.
9,853,819	B2	12/2017	Truu et al.	2017/0109667	A1	4/2017	Marcu et al.
9,853,977	B1	12/2017	Laucius et al.	2017/0109668	A1	4/2017	Marcu et al.
9,855,785	B1	1/2018	Nagelberg et al.	2017/0109670	A1	4/2017	Marcu et al.
9,858,546	B2	1/2018	Rinzler et al.	2017/0109676	A1	4/2017	Marcu et al.
9,858,766	B2	1/2018	Aleksey	2017/0109735	A1	4/2017	Sheng et al.
9,858,781	B1	1/2018	Campero et al.	2017/0109744	A1	4/2017	Wilkins et al.
9,862,092	B2	1/2018	Izhikevich et al.	2017/0329996	A1	11/2017	Wilson
9,862,222	B1	1/2018	Nagelberg et al.	2017/0330143	A1	11/2017	Raghavan et al.
9,864,970	B2	1/2018	Rinzler et al.	2017/0330159	A1	11/2017	Castinado et al.
9,866,393	B1	1/2018	Rush et al.	2017/0330174	A1	11/2017	Demarinis et al.
9,866,545	B2	1/2018	Beecham	2017/0330179	A1	11/2017	Song et al.
9,870,508	B1	1/2018	Hodgson et al.	2017/0330180	A1	11/2017	Song et al.
9,870,562	B2	1/2018	Davis et al.	2017/0330181	A1	11/2017	Ortiz
9,870,591	B2	1/2018	Shah	2017/0330250	A1	11/2017	Arjomand et al.
9,872,050	B2	1/2018	Uhr et al.	2017/0331624	A1	11/2017	Samid
9,875,489	B2	1/2018	Spitz et al.	2017/0331635	A1	11/2017	Barinov et al.
9,875,510	B1	1/2018	Kasper	2017/0331810	A1	11/2017	Kurian
9,875,592	B1	1/2018	Erickson et al.	2017/0331828	A1	11/2017	Caldera et al.
9,876,646	B2	1/2018	Ebrahimi et al.	2017/0331896	A1	11/2017	Holloway et al.
9,876,775	B2	1/2018	Mossbarger	2017/0337534	A1	11/2017	Goeringer et al.
9,881,176	B2	1/2018	Goldfarb et al.	2017/0338947	A1	11/2017	Ateniese et al.
9,882,715	B2	1/2018	Alness et al.	2017/0338957	A1	11/2017	Ateniese et al.
9,882,918	B1	1/2018	Ford et al.	2017/0338963	A1	11/2017	Berg
9,885,672	B2	2/2018	Forutanpour et al.	2017/0338967	A1	11/2017	Lewison et al.
9,886,685	B2	2/2018	Keohane et al.	2017/0339503	A1	11/2017	Lyren et al.
9,887,975	B1	2/2018	Gifford et al.	2017/0344435	A1	11/2017	Davis
9,887,995	B2	2/2018	Rotter et al.	2017/0344580	A1	11/2017	King
9,888,007	B2	2/2018	Caldera et al.	2017/0344983	A1	11/2017	Muftic
9,888,026	B2	2/2018	Kaplan et al.	2017/0344987	A1	11/2017	Davis
9,889,382	B2	2/2018	Thompson et al.	2017/0344988	A1	11/2017	Cusden et al.
9,892,141	B2	2/2018	Darcy	2017/0345011	A1	11/2017	Salami et al.
9,892,400	B1	2/2018	Grassadonia	2017/0345019	A1	11/2017	Radocchia et al.
9,892,460	B1	2/2018	Winklevoss et al.	2017/0345105	A1	11/2017	Isaacson et al.
9,892,597	B1	2/2018	Trepp et al.	2017/0346637	A1	11/2017	Zhang
9,894,101	B2	2/2018	Hendrick et al.	2017/0346639	A1	11/2017	Muftic
9,894,485	B2	2/2018	Finlow-Bates	2017/0346693	A1	11/2017	Dix et al.
9,898,739	B2	2/2018	Monastyrsky et al.	2017/0346752	A1	11/2017	Krishnamurthy et al.
9,898,781	B1	2/2018	Silverman	2017/0346804	A1	11/2017	Beecham
9,898,782	B1	2/2018	Winklevoss et al.	2017/0346830	A1	11/2017	Goldfarb et al.
9,900,841	B2	2/2018	Ur et al.	2017/0346833	A1	11/2017	Zhang
9,904,544	B2	2/2018	Thomas et al.	2017/0346907	A1	11/2017	Clawsie et al.
2016/0328705	A1 *	11/2016	Sebag	2017/0347253	A1	11/2017	Clawsie et al.
2017/0011460	A1 *	1/2017	Molinari	2017/0351693	A1	12/2017	Boudville
2017/0075938	A1	3/2017	Black et al.	2017/0352012	A1	12/2017	Hearn et al.
2017/0075941	A1	3/2017	Finlow-Bates	2018/0068359	A1 *	3/2018	Preston
2017/0076109	A1	3/2017	Kaditz et al.	2018/0075421	A1 *	3/2018	Serrano
2017/0076306	A1	3/2017	Snider et al.	2018/0097836	A1	4/2018	Stolarz et al.
2017/0078097	A1	3/2017	Carter et al.	2018/0097837	A1	4/2018	Stolarz et al.
2017/0078493	A1	3/2017	Melika et al.	2018/0097838	A1	4/2018	Stolarz et al.
2017/0083907	A1	3/2017	McDonough et al.	2018/0097841	A1	4/2018	Stolarz et al.
2017/0083911	A1	3/2017	Phillips	2018/0101448	A1	4/2018	Ventura et al.
2017/0083989	A1	3/2017	Brockman et al.	2018/0101455	A1	4/2018	Ventura et al.
2017/0084118	A1	3/2017	Robinson et al.	2018/0101557	A1	4/2018	Ventura et al.
2017/0085545	A1	3/2017	Lohe et al.	2018/0101560	A1	4/2018	Christidis et al.
2017/0085555	A1	3/2017	Bisikalo et al.	2018/0101684	A1	4/2018	Murphy et al.
2017/0088397	A1	3/2017	Buckman	2018/0101701	A1	4/2018	Barinov et al.
2017/0091397	A1	3/2017	Shah	2018/0101842	A1	4/2018	Ventura et al.
2017/0091467	A1	3/2017	Pogorelik et al.	2018/0101844	A1	4/2018	Song et al.
2017/0091750	A1	3/2017	Maim	2018/0101846	A1	4/2018	Mandal et al.
2017/0091756	A1	3/2017	Stern et al.	2018/0101848	A1	4/2018	Castagna et al.
2017/0098291	A1 *	4/2017	Code	2018/0101906	A1	4/2018	McDonald et al.
2017/0103167	A1	4/2017	Shah	2018/0101914	A1	4/2018	Samuel et al.
2017/0103385	A1	4/2017	Wilson, Jr. et al.	2018/0102013	A1	4/2018	Spanos et al.
2017/0103390	A1	4/2017	Wilson, Jr. et al.	2018/0103013	A1	4/2018	Imai et al.
2017/0103391	A1	4/2017	Wilson, Jr. et al.	2018/0103020	A1	4/2018	Moiyallah, Jr. et al.
2017/0103468	A1	4/2017	Orsini et al.	2018/0103042	A1	4/2018	Castagna et al.
2017/0103472	A1	4/2017	Shah	2018/0105787	A1	4/2018	Hardin
2017/0104831	A1	4/2017	Fransen	2018/0107958	A1	4/2018	Konik et al.
2017/0109475	A1	4/2017	Kaditz et al.	2018/0108010	A1	4/2018	McHale et al.
2017/0109636	A1	4/2017	Marcu et al.	2018/0108024	A1	4/2018	Greco et al.
				2018/0108089	A1	4/2018	Jayachandran
				2018/0109507	A1	4/2018	Caldera et al.
				2018/0109516	A1	4/2018	Song et al.
				2018/0109541	A1	4/2018	Gleichauf

(56)

References Cited

U.S. PATENT DOCUMENTS

2018/0110960	A1	4/2018	Youngblood et al.	
2018/0113752	A1	4/2018	Derbakova et al.	
2018/0114045	A1	4/2018	Ebrahimi et al.	
2018/0114063	A1	4/2018	Wexler et al.	
2018/0114169	A1	4/2018	Wiig et al.	
2018/0114205	A1	4/2018	Thomas et al.	
2018/0114218	A1	4/2018	Konik et al.	
2018/0114220	A1	4/2018	Ekberg	
2018/0114261	A1	4/2018	Jayachandran	
2018/0114329	A1	4/2018	Wexler et al.	
2018/0204190	A1*	7/2018	Moy	G06Q 40/04
2018/0218176	A1*	8/2018	Voorhees	H04L 9/3213
2018/0240191	A1*	8/2018	Aronson	G06Q 20/3823

* cited by examiner




Fig. 1

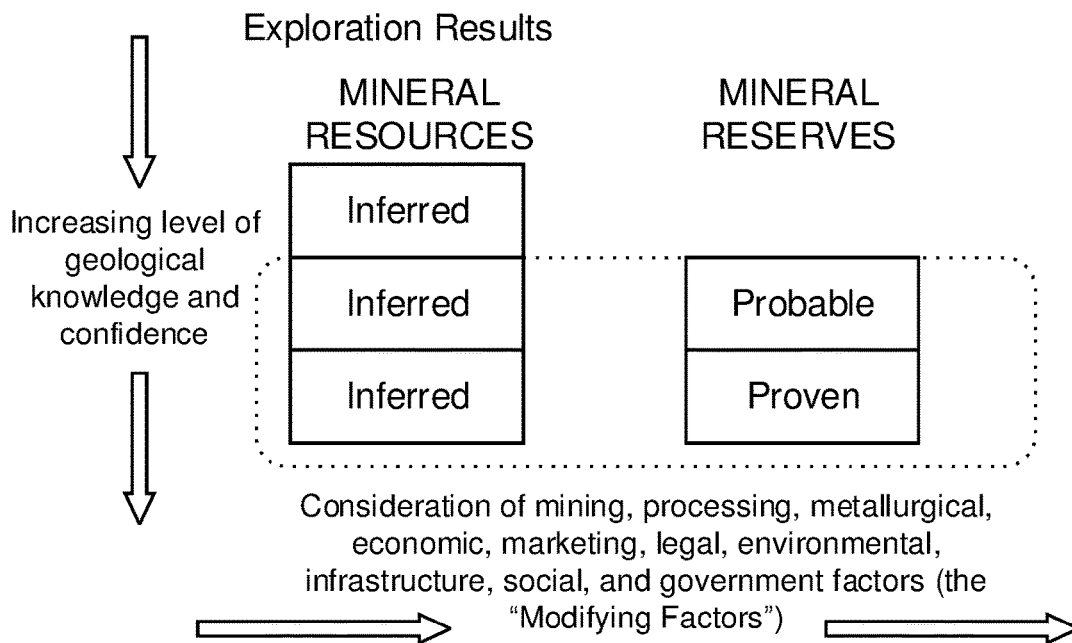


Fig. 3

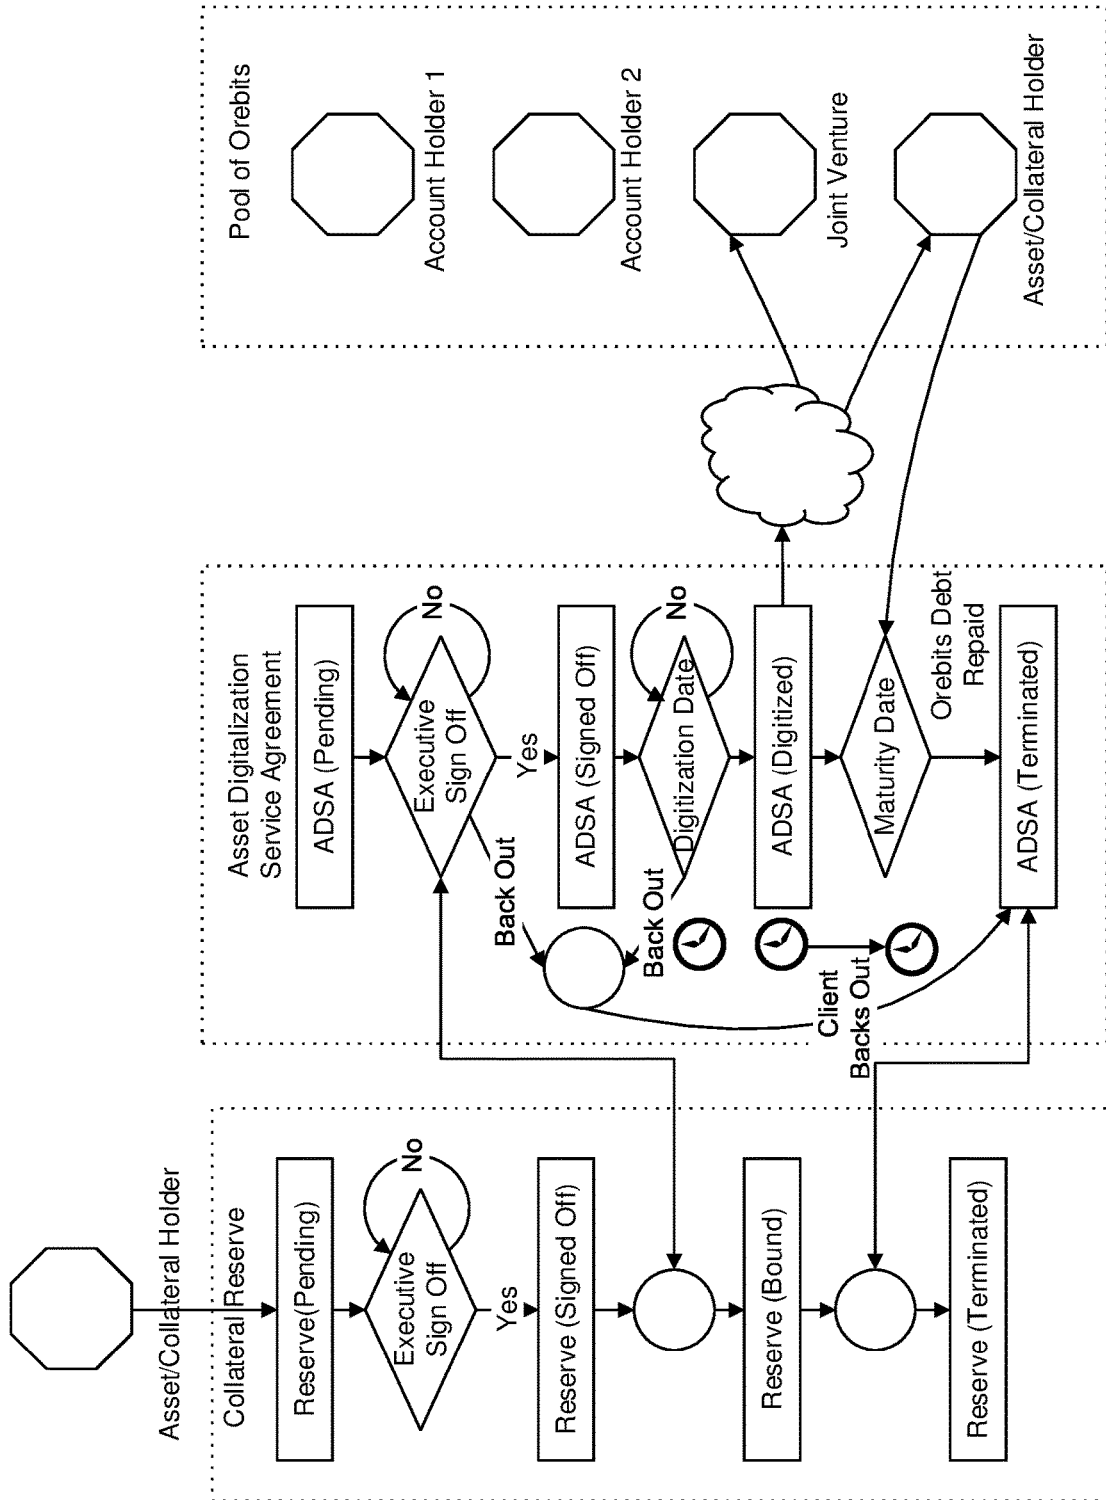


Fig. 2

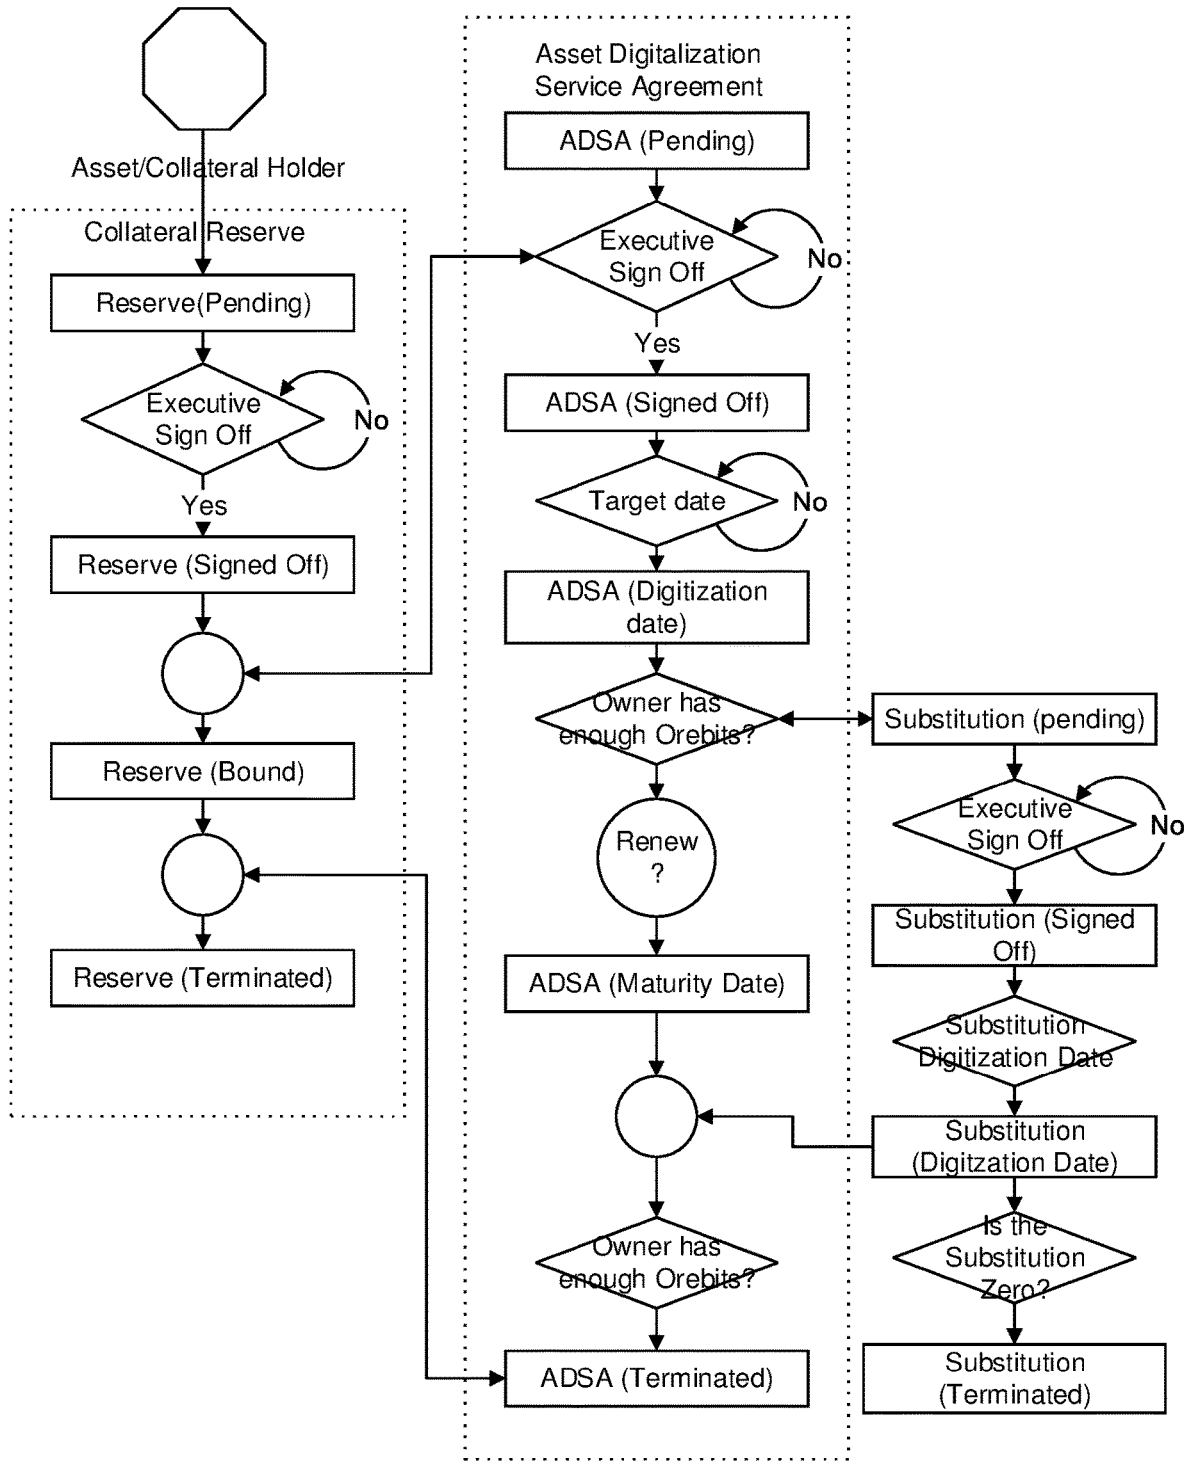


Fig. 4

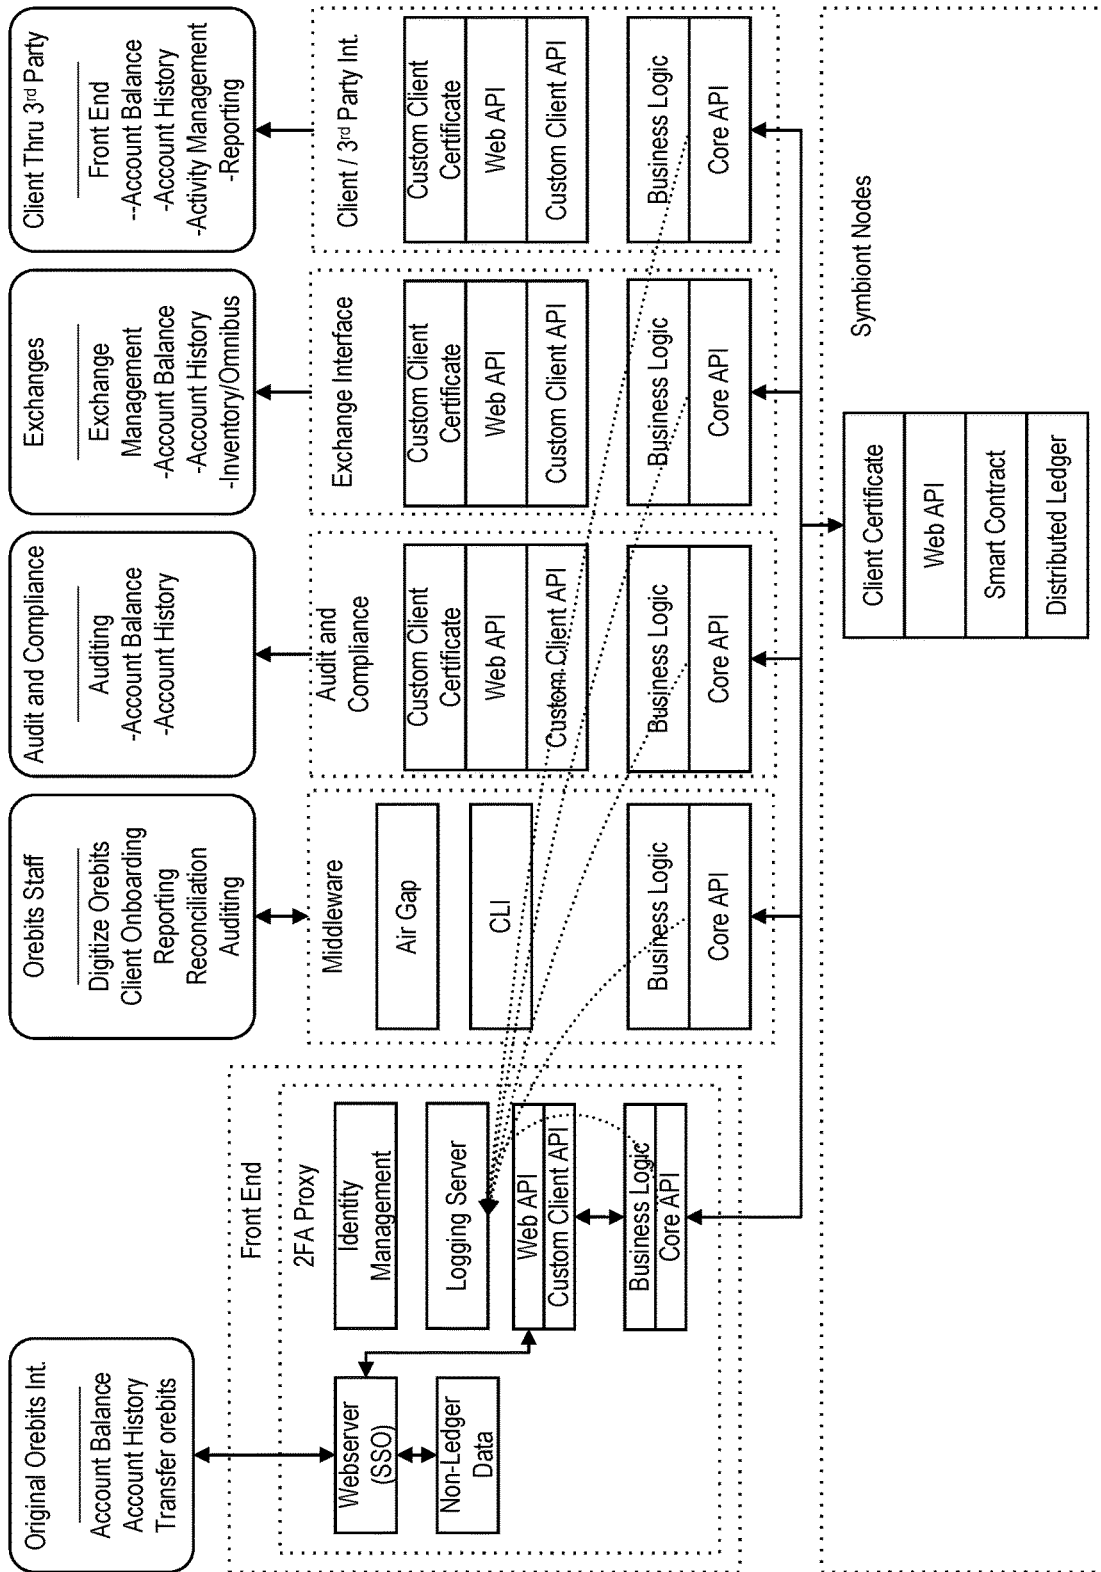


Fig. 5

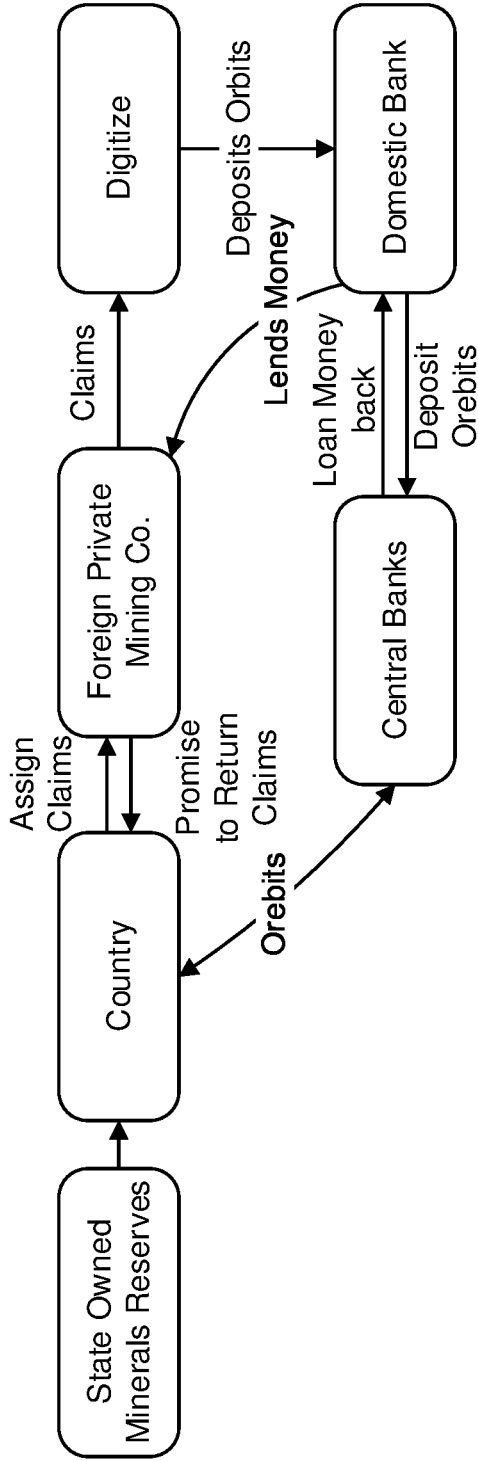


Fig. 6

METHOD FOR CREATING COMMODITY ASSETS FROM UNREFINED COMMODITY RESERVES UTILIZING BLOCKCHAIN AND DISTRIBUTED LEDGER TECHNOLOGY

CROSS REFERENCE TO RELATED APPLICATION

The present application claims benefit of priority from U.S. Provisional Patent Application Ser. No. 62/468,764, filed Mar. 8, 2017, the entirety of which is incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates to the field of securitized transactions and smart contracts, and encompasses systems and methods for conducting transactions.

BACKGROUND OF THE INVENTION

Each reference cited herein is expressly incorporated herein by reference in its entirety, for all purposes.

In the current marketplace, a commodity asset owner can go to a lender and securitize the commodity assets thereby gaining liquidity. The problem with this current model is that it requires a liquid commodity, and when securitized, the commodity may be restricted from beneficial use. Further, the value of the commodity may be deeply discounted, and ongoing interest charges are accrued.

Frolov et al., U.S. Pat. No. 9,747,586, discloses a system and method for issuance of electronic currency substantiated by a reserve of assets. The reserve is a commodity or asset that is actively traded.

Miner, US 20150332256, discloses a system and method for converting cryptocurrency to virtual assets whose value is substantiated by reserve of assets. The reserve is, for example, book entries for fiat currencies, which are actively traded.

Doney, US 20170213289, expressly incorporated herein by reference in its entirety, describes creation of collateralized portfolios, as a collection of income-producing assets, generated through transactions that exchange estimated asset value for liquid instruments in the portfolio. Transaction elasticity is provided by liquid instruments (reserve funds and portfolio-owned shares) held in reserve in the portfolio's reservoir which provides a market smoothing function to adapt to changes in asset demand and risk. Each portfolio's reservoir is collectively owned by the shareholders; continuously replenishing itself with income generated by assets in the portfolio. Shares can be represented by digital tokens, traded as digital currency such as cryptocurrency, and monetized with the convenience of cash through a network of exchanges and payment gateways.

Vieira et al., US20180047111, expressly incorporated herein by reference in its entirety, describes enhanced organizational transparency using a linked activity chain in a ledger, employing a block chain.

So-called "Smart Contracts" are legal obligations tied to a computer protocol intended to digitally facilitate, verify, or enforce the negotiation or performance of the contracts. Smart contracts allow the performance of credible transactions without third parties. These transactions are trackable and may be irreversible. See, en.wikipedia.org/wiki/Smart_contract.

The phrase "smart contracts" was coined by computer scientist Nick Szabo in 1996. See:

"Nick Szabo—Smart Contracts: Building Blocks for Digital Markets". www.fon.hum.uva.nl, Extropy #16;

Szabo, Nick (1997-09-01). "Formalizing and Securing Relationships on Public Networks". First Monday. 2 (9);

Tapscott, Don; Tapscott, Alex (May 2016). *The Blockchain Revolution: How the Technology Behind Bitcoin is Changing Money, Business, and the World*. pp. 72, 83, 101, 127. ISBN 978-0670069972.

A smart contract is a set of promises, specified in digital form, including protocols within which the parties perform on these promises. Recent implementations of smart contracts are based on blockchains, though this is not an intrinsic requirement. Building on this base, some recent interpretations of "smart contract" are mostly used more specifically in the sense of general purpose computation that takes place on a blockchain or distributed ledger. In this interpretation, used for example by the Ethereum Foundation or IBM, a smart contract is not necessarily related to the classical concept of a contract, but can be any kind of computer program. See:

Buterin, Vitalik. "Ethereum Whitepaper", Github;

Cachin, Christian. "Architecture of the Hyperledger Blockchain Fabric", ibm.com.

Byzantine fault tolerant algorithms allowed digital security through decentralization to form smart contracts. Additionally, the programming languages with various degrees of Turing-completeness as a built-in feature of some blockchains make the creation of custom sophisticated logic possible. See:

"Smart contracts: Turing completeness & reality";

"Dumb Contracts and Smart Scripts—We Use Cash". weuse.cash.

Notable examples of implementation of smart contracts are Decentralized cryptocurrency protocols are smart contracts with decentralized security, encryption, and limited trusted parties that fit Szabo's definition of a digital agreement with observability, verifiability, privity, and enforceability. See:\

"How Do Ethereum Smart Contracts Work?—CoinDesk". CoinDesk. Retrieved 2017-10-27;

"Bitcoin as a Smart Contract Platform". Richard Gendal Brown. 2015-03-30;

"Blockchain: Forget Bitcoin, here comes the real thing". Idealog. 2016-03-29;

"What are Smart Contracts" (PDF). chainfrog. Aug. 3, 2017.

Bitcoin also provides a Turing-incomplete Script language that allows the creation of custom smart contracts on top of Bitcoin like multisignature accounts, payment channels, escrows, time locks, atomic cross-chain trading, oracles, or multi-party lottery with no operator. See:

Rosa, Davide De. "The Bitcoin Script language (pt. 1)". davidederosa.com;

bitcoinbook: Mastering Bitcoin 2nd Edition—Programming the Open Blockchain—Chapter 7, Mastering Bitcoin, 2017-05-30;

"Smart contracts and bitcoin", medium.com/@maraoz/smart-contracts-and-bitcoin-a5d61011d9b1;

"Contract—Bitcoin Wiki". en.bitcoin.it;

"What is a Bitcoin Merklized Abstract Syntax Tree (MAST)?" Bitcoin Tech Talk. 2017-10-12;

"Smart Contracts on Bitcoin Blockchain" (PDF). Sep. 4, 2015;

Andrychowicz, Marcin; Dziembowski, Stefan; Malinowski, Daniel;

Mazurek, Łukasz (2013). "Secure Multiparty Computations on Bitcoin";

Atzei, Nicola; Bartoletti, Massimo; Cimoli, Tiziana; Lande, Stefano; Zunino, Roberto (2018), "SoK: unraveling

Bitcoin smart contracts”, 7th International Conference on Principles of Security and Trust (POST), European Joint Conferences on Theory and Practice of Software.

Ethereum implements a nearly Turing-complete language on its blockchain, a prominent smart contract framework. 5
See:

Atzei, Nicola; Bartoletti, Massimo; Cimoli, Tiziana (2017), “A survey of attacks on Ethereum smart contracts”, 6th International Conference on Principles of Security and Trust (POST), European Joint Conferences on Theory and Practice of Software; 10

“Vitalik Buterin on Tweeter (verified)”. 18 Apr. 2017.

RootStock (RSK) is a smart contract platform that is connected to the Bitcoin blockchain through sidechain technology. RSK is compatible with smart contracts created for Ethereum. See: 15

“RSK—Rootstock Open-Source Smart Contract Bitcoin Technology?”;

“Digrate Express rating report on Project Rootstock”;

Thomas Bocek (15 Sep. 2017). Digital Marketplaces Unleashed. Springer-Verlag GmbH. p. 169-184. ISBN 978-3-662-49274-1; 20

“A Solution for the Problems of Translation and Transparency in Smart Contracts”;

“Trust in Smart Contracts is a Process, As Well”;

“Scripting smart contracts for distributed ledger technology”.

See also:

“Namecoin”. Cointelegraph. 23 May 2015. Automated Transactions;

“Ripple discontinues smart contract platform Codius”. Bitcoin Magazine. Jun. 24, 2015;

“Automated Transactions Specification”;

“Qora and Burst Now Able to Make Cross-Chain Transactions”. May 22, 2015.

Smart contract infrastructure can be implemented by replicated asset registries and contract execution using cryptographic hash chains and *Byzantine* fault tolerant replication. See:

Nick Szabo (2005). “Secure Property Titles with Owner Authority”;

Jörg F. Wittenberger (2002). “Askemos a distributed settlement”;

“Proceedings of International Conference on Advances in Infrastructure for e-Business, e-Education, e-Science, and e-Medicine on the Internet”;

Martin Möbius (2009). “Erstellung eines Archivierungskonzepts für die Speicherung rückverfolgbarer Datenbestände im Askemos-System”;

Tom-Steve Watzke (2010). “Entwicklung einer Datenbankschnittstelle als Grundlage für Shop-Systeme unter dem Betriebssystem Askemos”.

RA Markus Heinker (2007). “Beweiswürdigung elektronischer Dokumente im Zivilprozess unter vergleichender Betrachtung von qualifizierten elektronischen Signaturen nach dem Signaturgesetz und dem Askemos-Verfahren”;

Hal Hodson (20 Nov. 2013). “Bitcoin moves beyond mere money”. New Scientist.

Smart contracts have advantages over equivalent conventional financial instruments, including minimizing counterparty risk, reducing settlement times, and increased transparency. UBS proposed “smart bonds” that use the bitcoin blockchain in which payment streams could hypothetically be fully automated, creating a self-paying instrument. See: 65

“Blockchain Technology: Preparing for Change”, Accenture;

Ross, Rory (2015-09-12). “Smart Money: Blockchains Are the Future of the Internet”, Newsweek;

Wigan, David (2015-06-11). “Bitcoin technology will disrupt derivatives, says banker”, IFR Asia.

See also,

Buterin, Vitalik. “A next-generation smart contract and decentralized application platform.” *white paper* (2014).

Delmolino, Kevin, Mitchell Arnett, Ahmed Kosba, Andrew Miller, and Elaine Shi. “Step by step towards creating a safe smart contract: Lessons and insights from a cryptocurrency lab.” In *International Conference on Financial Cryptography and Data Security*, pp. 79-94. Springer, Berlin, Heidelberg, 2016.

Clack, Christopher D., Vikram A. Bakshi, and Lee Braine. “Smart contract templates: foundations, design landscape and research directions.” *arXiv preprint arXiv:1608.00771* (2016).

Swan, Melanie. “Blockchain temporality: smart contract time speciifiability with blocktime.” In *International symposium on rules and rule markup languages for the semantic web*, pp. 184-196. Springer, Cham, 2016.

Vukolić, Marko. “The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication.” In *International Workshop on Open Problems in Network Security*, pp. 112-125. Springer, Cham, 2015.

Yasin, Affan, and Lin Liu. “An online identity and smart contract management system.” In *Computer Software and Applications Conference (COMPSAC)*, 2016 *IEEE 40th Annual*, vol. 2, pp. 192-198. IEEE, 2016. 30

Clack, Christopher D., Vikram A. Bakshi, and Lee Braine. “Smart Contract Templates: essential requirements and design options.” *arXiv preprint arXiv:1612.04496* (2016).

Norta, Alex. “Creation of smart-contracting collaborations for decentralized autonomous organizations.” In *International Conference on Business Informatics Research*, pp. 3-17. Springer, Cham, 2015.

Christidis, Konstantinos, and Michael Devetsikiotis. “Blockchains and smart contracts for the internet of things.” *IEEE Access* 4 (2016): 2292-2303.

Beck, Roman, Jacob Stenum Czepluch, Nikolaj Lolllike, and Simon Malone. “Blockchain—the Gateway to Trust-Free Cryptographic Transactions.” In *ECIS*, p. ResearchPaper153. 2016.

Frantz, Christopher K., and Mariusz Nowostawski. “From institutions to code: Towards automated generation of smart contracts.” In *Foundations and Applications of Self*Systems, IEEE International Workshops on*, pp. 210-215. IEEE, 2016.

Idelberger, Florian, Guido Governatori, Régis Riveret, and Giovanni Sartor. “Evaluation of logic-based smart contracts for blockchain systems.” In *International Symposium on Rules and Rule Markup Languages for the Semantic Web*, pp. 167-183. Springer, Cham, 2016.

Weber, Ingo, Xiwei Xu, Régis Riveret, Guido Governatori, Alexander Ponomarev, and Jan Mendling. “Untrusted business process monitoring and execution using blockchain.” In *International Conference on Business Process Management*, pp. 329-347. Springer, Cham, 2016.

Xu, Xiwei, Cesare Pautasso, Liming Zhu, Vincent Gramoli, Alexander Ponomarev, An Binh Tran, and Shiping Chen. “The blockchain as a software connector.” In *Software Architecture (WICSA)*, 2016 *13th Working IEEE/IFIP Conference on*, pp. 182-191. IEEE, 2016.

Kosba, Ahmed, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. “Hawk: The blockchain model of cryptography and privacy-preserving smart con-

tracts.” In *Security and Privacy (SP)*, 2016 *IEEE Symposium on*, pp. 839-858. IEEE, 2016.

Bogner, Andreas, Mathieu Chanson, and Arne Meeuw. “A decentralised sharing app running a smart contract on the ethereum blockchain.” In *Proceedings of the 6th International Conference on the Internet of Things*, pp. 177-178. ACM, 2016.

Brown, Richard Gendal, James Carlyle, Ian Grigg, and Mike Hearn. “Corda: An Introduction.” *R3 CEV*, August (2016).

Huckle, Steve, Rituparna Bhattacharya, Martin White, and Natalia Beloff. “Internet of things, blockchain and shared economy applications.” *Procedia Computer Science* 98 (2016): 461-466.

Luu, Loi, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. “Making smart contracts smarter.” In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 254-269. ACM, 2016.

Ammous, Saifedean Hisham. “Blockchain Technology: What is it good for?.” *Browser Download This Paper* (2016).

Al-Bassam, Mustafa. “SCPki: a smart contract-based PKI and identity system.” In *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, pp. 35-40. ACM, 2017.

Glaser, Florian. “Pervasive decentralisation of digital infrastructures: a framework for blockchain enabled system and use case analysis.” (2017).

Porru, Simone, Andrea Pinna, Michele Marchesi, and Roberto Tonelli. “Blockchain-oriented software engineering: challenges and new directions.” In *Proceedings of the 39th International Conference on Software Engineering Companion*, pp. 169-171. IEEE Press, 2017.

Zhang, Fan, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. “Town crier: An authenticated data feed for smart contracts.” In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 270-282. ACM, 2016.

Bartoletti, Massimo, and Livio Pompianu. “An empirical analysis of smart contracts: platforms, applications, and design patterns.” In *International Conference on Financial Cryptography and Data Security*, pp. 494-509. Springer, Cham, 2017.

Maskell, Brian. “Just-in-time manufacturing.” *Industrial Management & Data Systems* 87, no. 9/10 (1987): 17-20.

Bahga, Arshdeep, and Vijay K. Madiseti. “Blockchain platform for industrial Internet of Things.” *Journal of Software Engineering and Applications* 9, no. 10 (2016): 533.

Korpela, Kari, Jukka Hallikas, and Tomi Dahlberg. “Digital supply chain transformation toward blockchain integration.” In *proceedings of the 50th Hawaii international conference on system sciences*. 2017.

Bhargavan, Karthikeyan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova et al. “Formal verification of smart contracts: Short paper.” In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security*, pp. 91-96. ACM, 2016.

Vukolić, Marko. “Rethinking permissioned blockchains.” In *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, pp. 3-7. ACM, 2017.

Levy, Karen E C. “Book-smart, not street-smart: blockchain-based smart contracts and the social workings of law.” *Engaging Science, Technology, and Society* 3 (2017): 1-15.

Marino, Bill, and Ari Juels. “Setting standards for altering and undoing smart contracts.” In *International Symposium*

on Rules and Rule Markup Languages for the Semantic Web, pp. 151-166. Springer, Chain, 2016.

Morrison, Alan, and Subhankar Sinha. “Blockchain and smart contract automation: Blockchains defined.” (2016).

Atzei, Nicola, Massimo Bartoletti, and Tiziana Cimoli. “A survey of attacks on Ethereum smart contracts (SoK).” In *International Conference on Principles of Security and Trust*, pp. 164-186. Springer, Berlin, Heidelberg, 2017.

Peters, Gareth W., and Efstathios Panayi. “Understanding modern banking ledgers through blockchain technologies: Future of transaction processing and smart contracts on the internet of money.” In *Banking Beyond Banks and Money*, pp. 239-278. Springer, Chain, 2016.

Dai, Patrick, Neil Mahi, Jordan Earls, and Alex Norta. “Smart-contract value-transfer protocols on a distributed mobile application platform” tum.org/uploads/files/cf6d69348ca50dd985b60425ccf282f3. (2017).

Omohundro, Steve. “Cryptocurrencies, smart contracts, and artificial intelligence.” *AI matters* 1, no. 2 (2014): 19-21.

Foroglou, George, and Anna-Lali Tsilidou. “Further applications of the blockchain.” In *12th Student Conference on Managerial Science and Technology*. 2015.

Kolvart, Merit, Margus Poola, and Addi Rull. “Smart contracts.” In *The Future of Law and etechnologies*, pp. 133-147. Springer, Chain, 2016.

Hull, Richard, Vishal S. Batra, Yi-Min Chen, Alin Deutsch, Fenno F. Terry Heath III, and Victor Vianu. “Towards a shared ledger business collaboration language based on data-aware processes.” In *International Conference on Service-Oriented Computing*, pp. 18-36. Springer, Chain, 2016.

Szabo, Nick. “Formalizing and securing relationships on public networks.” *First Monday* 2, no. 9 (1997).

English, Matthew, Soren Auer, and John Domingue. “Block chain technologies & the semantic web: A framework for symbiotic development.” In *Computer Science Conference for University of Bonn Students*, J. Lehmann, H. Thakkar, L. Halilaj, and R. Asmat, Eds, pp. 47-61. 2016.

Miller, Mark S., Chip Morningstar, and Bill Frantz. “Capability-based financial instruments.” In *International Conference on Financial Cryptography*, pp. 349-378. Springer, Berlin, Heidelberg, 2000.

Zheng, Zibin, Shaoan Xie, Hong-Ning Dai, and Huaimin Wang. “Blockchain challenges and opportunities: A survey.” *Work Pap.-2016* (2016).

Hirai, Yoichi. “Defining the ethereum virtual machine for interactive theorem provers.” In *International Conference on Financial Cryptography and Data Security*, pp. 520-535. Springer, Chain, 2017.

Xu, Xiwei, Ingo Weber, Mark Staples, Liming Zhu, Jan Bosch, Len Bass, Cesare Pautasso, and Paul Rimba. “A taxonomy of blockchain-based systems for architecture design.” In *Software Architecture (ICSA), 2017 IEEE International Conference on*, pp. 243-252. IEEE, 2017.

Chesebro, Russell. *A contract that manages itself the time has arrived*. Defense Acquisition Univ Ft Belvoir VA, 2015.

Savelyev, Alexander. “Contract law 2.0: Smart contracts as the beginning of the end of classic contract law.” *Information & Communications Technology Law* 26, no. 2 (2017): 116-134.

See, U.S. Pat. Nos. 6,324,286; 6,938,039; 9,014,661; 9,135,787; 9,298,806; 9,300,467; 9,331,856; 9,338,148; 9,351,124; 9,397,985; 9,413,735; 9,436,923; 9,436,935; 9,480,188; 9,507,984; 9,509,690; 9,513,627; 9,558,524; 9,563,873; 9,569,771; 9,608,829; 9,635,000; 9,641,338; 9,641,342; 9,665,734; 9,667,427; 9,667,600; 9,679,276; 9,702,582; 9,703,986; 9,705,682; 9,705,851; 9,710,808;

US 11,188,977 B2

7

9,716,595; 9,722,790; 9,743,272; 9,747,586; 9,749,140;
 9,749,297; 9,749,766; 9,754,131; 9,760,574; 9,760,827;
 9,767,520; 9,773,099; 9,774,578; 9,785,369; 9,792,101;
 9,794,074; 9,805,381; 9,807,106; 9,818,092; 9,818,116;
 9,820,120; 9,824,031; 9,824,408; 9,825,931; 9,832,026;
 9,836,908; 9,847,997; 9,848,271; 9,849,364; 9,852,427;
 9,853,819; 9,853,977; 9,855,785; 9,858,781; 9,862,222;
 9,866,545; 9,870,508; 9,870,562; 9,870,591; 9,875,510;
 9,875,592; 9,876,646; 9,876,775; 9,881,176; 9,882,918;
 9,888,007; 9,892,141; 9,892,460; 9,894,485; 9,898,782;
 9,904,544; 9,906,513; 9,910,969; 9,912,659; D759073;
 20050203815; 20140344015; 20140368601; 20150067143;
 20150081566; 20150127940; 20150170112; 20150206106;
 20150244690; 20150262137; 20150262138; 20150262139;
 20150262140; 20150262141; 20150262168; 20150262171;
 20150262172; 20150262176; 20150269624; 20150278820;
 20150278887; 20150294425; 20150310476; 20150324764;
 20150332256; 20150332283; 20150348169; 20150356524;
 20150356555; 20150379510; 20160005032; 20160012424;
 20160012465; 20160027229; 20160028552; 20160055236;
 20160071108; 20160072800; 20160092988; 20160098723;
 20160098730; 20160117471; 20160123620; 20160134593;
 20160140653; 20160170996; 20160170998; 20160171514;
 20160180338; 20160191243; 20160192166; 20160203448;
 20160203522; 20160203572; 20160203575; 20160210626;
 20160210710; 20160212109; 20160212146; 20160217436;
 20160217532; 20160218879; 20160224803; 20160224949;
 20160234026; 20160253663; 20160254910; 20160259923;
 20160260091; 20160260169; 20160261411; 20160267472;
 20160267474; 20160267558; 20160267566; 20160267601;
 20160267605; 20160269182; 20160269402; 20160275461;
 20160283920; 20160283939; 20160283941; 20160284033;
 20160292396; 20160292672; 20160292680; 20160294783;
 20160300200; 20160300223; 20160300234; 20160300252;
 20160306982; 20160307197; 20160321316; 20160321434;
 20160321435; 20160321629; 20160321654; 20160321675;
 20160321676; 20160321751; 20160321752; 20160321769;
 20160323109; 20160327294; 20160328713; 20160330027;
 20160330034; 20160335533; 20160335609; 20160342958;
 20160342959; 20160342976; 20160342977; 20160342978;
 20160342980; 20160342981; 20160342982; 20160342983;
 20160342984; 20160342985; 20160342986; 20160342987;
 20160342988; 20160342989; 20160342994; 20160350728;
 20160350749; 20160357550; 20160358158; 20160358165;
 20160358169; 20160358184; 20160358186; 20160358187;
 20160358253; 20160358267; 20160359637; 20160364700;
 20160364787; 20160365978; 20160371771; 20160379212;
 20160379213; 20160379256; 20160379298; 20160379312;
 20160379330; 20170004563; 20170004578; 20170004588;
 20170005804; 20170011053; 20170011392; 20170011460;
 20170012780; 20170012943; 20170013047; 20170017936;
 20170017954; 20170017955; 20170017958; 20170019496;
 20170024817; 20170024818; 20170028622; 20170031676;
 20170031874; 20170033932; 20170034197; 20170039330;
 20170041148; 20170042068; 20170046526; 20170046638;
 20170046651; 20170046652; 20170046664; 20170046670;
 20170046680; 20170046689; 20170046693; 20170046694;
 20170046698; 20170046709; 20170046799; 20170046806;
 20170048209; 20170048216; 20170048234; 20170048235;
 20170052676; 20170053036; 20170053131; 20170054611;
 20170061396; 20170070778; 20170075877; 20170075938;
 20170075941; 20170076109; 20170076306; 20170078097;
 20170078493; 20170083907; 20170083911; 20170083989;
 20170084118; 20170085545; 20170085555; 20170088397;
 20170091397; 20170091467; 20170091750; 20170091756;
 20170098291; 20170103167; 20170103385; 20170103390;
 20170103391; 20170103468; 20170103472; 20170104831;

8

20170109475; 20170109636; 20170109637; 20170109638;
 20170109639; 20170109640; 20170109657; 20170109667;
 20170109668; 20170109670; 20170109676; 20170109735;
 20170109744; 20170109748; 20170109814; 20170109955;
 20170111175; 20170111385; 20170115976; 20170116693;
 20170118301; 20170124535; 20170124556; 20170124647;
 20170126702; 20170131988; 20170132393; 20170132615;
 20170132619; 20170132620; 20170132621; 20170132625;
 20170132626; 20170132630; 20170132635; 20170132636;
 20170134161; 20170134162; 20170134280; 20170134326;
 20170134375; 20170134937; 20170140145; 20170140371;
 20170140375; 20170140394; 20170140408; 20170147808;
 20170147975; 20170148016; 20170148021; 20170149560;
 20170149795; 20170149796; 20170149819; 20170150939;
 20170154331; 20170155515; 20170161439; 20170161517;
 20170161652; 20170161697; 20170161733; 20170161734;
 20170161762; 20170161829; 20170161833; 20170163733;
 20170169125; 20170169363; 20170169473; 20170169800;
 20170173262; 20170177855; 20170177898; 20170178127;
 20170178128; 20170178236; 20170178237; 20170178263;
 20170178417; 20170180128; 20170180134; 20170180211;
 20170185692; 20170185981; 20170185998; 20170186057;
 20170187535; 20170188168; 20170191688; 20170192994;
 20170193464; 20170193619; 20170195299; 20170195336;
 20170195397; 20170195747; 20170198931; 20170199671;
 20170200137; 20170200147; 20170201385; 20170205102;
 20170206382; 20170206522; 20170206523; 20170206532;
 20170206603; 20170206604; 20170207917; 20170212781;
 20170213198; 20170213209; 20170213210; 20170213221;
 20170213287; 20170213289; 20170214522; 20170214675;
 20170214699; 20170214701; 20170220781; 20170220815;
 20170220998; 20170221021; 20170221029; 20170221032;
 20170221052; 20170222814; 20170228371; 20170228557;
 20170228705; 20170228706; 20170228731; 20170228734;
 20170228822; 20170228841; 20170230189; 20170230285;
 20170230345; 20170230349; 20170230353; 20170230375;
 20170230791; 20170232300; 20170234709; 20170235970;
 20170236094; 20170236102; 20170236103; 20170236104;
 20170236120; 20170236121; 20170236123; 20170236143;
 20170236176; 20170236196; 20170237553; 20170237554;
 20170237569; 20170237570; 20170237700; 20170238072;
 20170242475; 20170243025; 20170243177; 20170243179;
 20170243193; 20170243208; 20170243209; 20170243212;
 20170243213; 20170243214; 20170243215; 20170243216;
 20170243217; 20170243222; 20170243239; 20170244324;
 20170243284; 20170243286; 20170243287; 20170244707;
 20170244720; 20170244721; 20170244757; 20170244909;
 20170249482; 20170249623; 20170250796; 20170250815;
 20170250972; 20170251025; 20170255912; 20170256000;
 20170256001; 20170256003; 20170256951; 20170257358;
 20170262778; 20170262862; 20170262902; 20170264428;
 20170265789; 20170270492; 20170270493; 20170270527;
 20170277909; 20170278080; 20170278186; 20170279620;
 20170279774; 20170279783; 20170279818; 20170285720;
 20170286717; 20170286880; 20170286951; 20170287068;
 20170287090; 20170289111; 20170289124; 20170289134;
 20170291295; 20170293503; 20170293669; 20170293898;
 20170293912; 20170295021; 20170295023; 20170295157;
 20170295180; 20170295232; 20170300627; 20170300872;
 20170300875; 20170300876; 20170300877; 20170300898;
 20170300905; 20170300910; 20170300928; 20170300946;
 20170300978; 20170301031; 20170301033; 20170301047;
 20170302450; 20170302460; 20170307387; 20170308070;
 20170308893; 20170308920; 20170308928; 20170309117;
 20170310484; 20170310653; 20170310747; 20170316162;
 20170316390; 20170316391; 20170316409; 20170316410;
 20170316487; 20170316497; 20170317833; 20170317834;

20170317997; 20170318008; 20170323294; 20170323392;
 20170324738; 20170329922; 20170329980; 20170329996;
 20170330143; 20170330159; 20170330174; 20170330179;
 20170330180; 20170330181; 20170330250; 20170331624;
 20170331635; 20170331810; 20170331828; 20170331896;
 20170337534; 20170338947; 20170338957; 20170338963;
 20170338967; 20170339503; 20170344435; 20170344580;
 20170344983; 20170344987; 20170344988; 20170345011;
 20170345019; 20170345105; 20170346637; 20170346639;
 20170346693; 20170346752; 20170346804; 20170346830;
 20170346833; 20170346907; 20170347253; 20170351693;
 20170352012; 20170352027; 20170352031; 20170352033;
 20170352116; 20170352185; 20170352219; 20170353309;
 20170353311; 20170353320; 20170357966; 20170357970;
 20170358041; 20170359288; 20170359316; 20170359374;
 20170359408; 20170364450; 20170364552; 20170364637;
 20170364655; 20170364698; 20170364699; 20170364700;
 20170364701; 20170364702; 20170364825; 20170364860;
 20170364871; 20170364900; 20170364908; 20170364934;
 20170364936; 20170366348; 20170366353; 20170366399;
 20170366416; 20170366516; 20170366547; 20170372278;
 20170372300; 20170372308; 20170372391; 20170372392;
 20170372431; 20170373849; 20170374049; 20180000367;
 20180001184; 20180005186; 20180005318; 20180005326;
 20180005489; 20180005492; 20180006826; 20180006829;
 20180006831; 20180006990; 20180007131; 20180012195;
 20180012262; 20180012311; 20180013567; 20180013815;
 20180015838; 20180017791; 20180018590; 20180018655;
 20180018695; 20180018723; 20180018738; 20180019867;
 20180019872; 20180019873; 20180019879; 20180019921;
 20180019984; 20180019993; 20180020324; 20180025135;
 20180025140; 20180025166; 20180025181; 20180025272;
 20180025365; 20180025368; 20180025388; 20180025401;
 20180025435; 20180025442; 20180026979; 20180032273;
 20180032383; 20180032696; 20180032759; 20180034634;
 20180034636; 20180034642; 20180034804; 20180039397;
 20180039512; 20180039667; 20180039785; 20180039786;
 20180039942; 20180039982; 20180039993; 20180040007;
 20180040040; 20180040041; 20180040062; 20180041072;
 20180041338; 20180041345; 20180041445; 20180041446;
 20180041486; 20180041487; 20180041571; 20180043386;
 20180046766; 20180046956; 20180046992; 20180047111;
 20180048461; 20180048463; 20180048469; 20180048485;
 20180048738; 20180052462; 20180052813; 20180052926;
 20180052927; 20180052970; 20180053158; 20180053160;
 20180053161; 20180053182; 20180054436; 20180054491;
 20180060496; 20180060596; 20180060600; 20180060771;
 20180060835; 20180060836; 20180060860; 20180060927;
 20180061237; 20180062831; 20180062835; 20180062848;
 20180063099; 20180063139; 20180063189; 20180063238;
 20180068091; 20180068097; 20180068130; 20180068271;
 20180068282; 20180068359; 20180069798; and
 20180069899.

Liquidity, the ability to efficiently convert asset value to cash on demand, is a key characteristic of optimal markets. Likewise, high transaction costs, inability to liquidate an asset on demand, and discounting of an asset in order to facilitate a transaction represent market inefficiencies. Many investment opportunities, such as emerging technologies or real estate projects in the developing markets, offer significant earning potential but suffer from a lack of liquidity. Despite the potential for return, earning potential for these asset classes may remain dormant. Asset liquidity may be limited to due to lack of information, individual asset risk, uncertain market conditions, large transaction sizes, and irregular or infrequent payouts.

A particular class of assets involve resources capable of generating or producing commodities at a predictable cost (i.e., assets for which there exists a robust market with high efficiency), but which themselves have limited liquidity because of the cost, time, and risk of bringing the asset to commodity status. These assets are thus ultimately valuable, but suffer significant discounts beyond the predicted cost of commoditization. For example, mineral deposits which produce commodity minerals fluctuate based on the value of the mineral, but suffer discounts far in excess of the cost of extraction. Further, because the value is intrinsic to the deposit, the use of the asset for securitization depends on the extractability of the resource, and not its actual extraction, permitting leverage of the capital investment independent of the mineral business. Indeed, because the market for commodity minerals may fluctuate or suffer cyclic pricing, it may be inefficient to extract mineral deposits under all circumstances, yet liquidity of the underlying real estate or leases is a significant risk on owning or lending based on these deposits.

Currently, a company which owns a source or means for production of a commodity may borrow at interest from a lender secured by the assets, or sells equity which intrinsically shares in the profits of production, or offers some hybrid security. While an equity holder may have liquidity in selling the shares, the value of the shares is heavily weighted toward the profits made from operations, and negatively impacted by the size of the capital investment involved. Meanwhile, a debt lender (e.g., bondholder) typically has more limited liquidity, and even if the security interest is sufficient to ensure eventual payment, the default risk for the borrower remains a critical factor in liquidity.

The Ethereum White Paper, Vitalik Buterin, "Ethereum White Paper A Next Generation Smart Contract & Decentralized Application Platform" describes the Ethereum platform.

Because currency is a first-to-file application, where the order of transactions is often of critical importance, decentralized currencies require a solution to decentralized consensus. The main roadblock that all pre-Bitcoin currency protocols faced is the fact that, while there had been plenty of research on creating secure *Byzantine*-fault-tolerant multiparty consensus systems for many years, all of the protocols described were solving only half of the problem. The protocols assumed that all participants in the system were known, and produced security margins of the form "if N parties participate, then the system can tolerate up to N/4 malicious actors". The problem is, however, that in an anonymous setting such security margins are vulnerable to sybil attacks, where a single attacker creates thousands of simulated nodes on a server or botnet and uses these nodes to unilaterally secure a majority share.

The innovation provided by Satoshi Nakamoto is the idea of combining a very simple decentralized consensus protocol, based on nodes combining transactions into a "block" every ten minutes creating an ever-growing blockchain, with proof of work as a mechanism through which nodes gain the right to participate in the system. While nodes with a large amount of computational power do have proportionately greater influence, coming up with more computational power than the entire network combined is much harder than simulating a million nodes. Despite the Bitcoin blockchain model's crudeness and simplicity, it has proven to be good enough, and would over the next five years become the bedrock of over two hundred currencies and protocols around the world.

11

Bitcoin as a State Transition System

From a technical standpoint, the Bitcoin ledger can be thought of as a state transition system, where there is a “state” consisting of the ownership status of all existing bitcoins and a “state transition function” that takes a state and a transaction and outputs a new state which is the result. In a standard banking system, for example, the state is a balance sheet, a transaction is a request to move \$X from A to B, and the state transition function reduces the value in A’s account by \$X and increases the value in B’s account by \$X. If A’s account has less than \$X in the first place, the state transition function returns an error. Hence, one can formally define:

APPLY(S,TX)→S’ or ERROR

In the banking system defined above: APPLY({Alice: \$50, Bob: \$50}, “send \$20 from Alice to Bob”)={Alice: \$30, Bob: \$70}

But: APPLY({Alice: \$50, Bob: \$50}, “send \$70 from Alice to Bob”) = ERROR

The “state” in Bitcoin is the collection of all coins (technically, “unspent transaction outputs” or UTXO) that have been minted and not yet spent, with each UTXO having a denomination and an owner (defined by a 20-byte address which is essentially a cryptographic public key[1]). A transaction contains one or more inputs, with each input containing a reference to an existing UTXO and a cryptographic signature produced by the private key associated with the owner’s address, and one or more outputs, with each output containing a new UTXO to be added to the state.

The state transition function APPLY(S,TX)→S’ can be defined roughly as follows:

1. For each input in TX:
 - i. If the referenced UTXO is not in S, return an error.
 - ii. If the provided signature does not match the owner of the UTXO, return an error.
2. If the sum of the denominations of all input UTXO is less than the sum of the denominations of all output UTXO, return an error.

3. Return S with all input UTXO removed and all output UTXO added.

The first half of the first step prevents transaction senders from spending coins that do not exist, the second half of the first step prevents transaction senders from spending other people’s coins, and the second step enforces conservation of value. In order to use this for payment, the protocol is as follows. Suppose Alice wants to send 11.7 BTC to Bob. First, Alice will look for a set of available UTXO that she owns that totals up to at least 11.7 BTC. Realistically, Alice will not be able to get exactly 11.7 BTC; say that the smallest she can get is 6+4+2=12. She then creates a transaction with those three inputs and two outputs. The first output will be 11.7 BTC with Bob’s address as its owner, and the second output will be the remaining 0.3 BTC “change”, with the owner being Alice herself.

Mining

If we had access to a trustworthy centralized service, this system would be trivial to implement; it could simply be coded exactly as described. However, with Bitcoin we are trying to build a decentralized currency system, so we will need to combine the state transition system with a consensus system in order to ensure that everyone agrees on the order of transactions. Bitcoin’s decentralized consensus process requires nodes in the network to continuously attempt to produce packages of transactions called “blocks”. The network is intended to produce roughly one block every ten minutes, with each block containing a timestamp, a nonce, a reference to (i.e., hash of) the previous block and a list of

12

all of the transactions that have taken place since the previous block. Over time, this creates a persistent, ever-growing, “blockchain” that constantly updates to represent the latest state of the Bitcoin ledger.

The algorithm for checking if a block is valid, expressed in this paradigm, is as follows:

1. Check if the previous block referenced by the block exists and is valid
2. Check that the timestamp of the block is greater than that of the previous block[2] and less than 2 hours into the future.
3. Check that the proof of work on the block is valid.
4. Let S[0] be the state at the end of the previous block.
5. Suppose TX is the block’s transaction list with n transactions. For all i in 0 . . . n-1, setS[i+1]=APPLY(S[i], TX[i]) If any application returns an error, exit and return false.
6. Return true, and register S[n] as the state at the end of this block

Essentially, each transaction in the block must provide a state transition that is valid. Note that the state is not encoded in the block in any way; it is purely an abstraction to be remembered by the validating node and can only be (securely) computed for any block by starting from the genesis state and sequentially applying every transaction in every block. Additionally, note that the order in which the miner includes transactions into the block matters; if there are two transactions A and B in a block such that B spends a UTXO created by A, then the block will be valid if A comes before B but not otherwise.

The interesting part of the block validation algorithm is the concept of “proof of work”: the condition is that the SHA256 hash of every block, treated as a 256-bit number, must be less than a dynamically adjusted target, which as of the time of this writing is approximately 2190. The purpose of this is to make block creation computationally “hard”, thereby preventing sybil attackers from remaking the entire blockchain in their favor.

Because SHA256 is designed to be a completely unpredictable pseudorandom function, the only way to create a valid block is simply trial and error, repeatedly incrementing the nonce and seeing if the new hash matches. At the current target of **2192**, this means an average of 264 tries; in general, the target is recalibrated by the network every 2016 blocks so that on average a new block is produced by some node in the network every ten minutes. In order to compensate miners for this computational work, the miner of every block is entitled to include a transaction giving themselves 25 BTC out of nowhere. Additionally, if any transaction has a higher total denomination in its inputs than in its outputs, the difference also goes to the miner as a “transaction fee”. Incidentally, this is also the only mechanism by which BTC are issued; the genesis state contained no coins at all.

In order to better understand the purpose of mining, let us examine what happens in the event of a malicious attacker. Since Bitcoin’s underlying cryptography is known to be secure, the attacker will target the one part of the Bitcoin system that is not protected by cryptography directly: the order of transactions. The attacker’s strategy is simple:

1. Send 100 BTC to a merchant in exchange for some product (preferably a rapid-delivery digital good)
2. Wait for the delivery of the product
3. Produce another transaction sending the same 100 BTC to himself
4. Try to convince the network that his transaction to himself was the one that came first.

Once step (1) has taken place, after a few minutes some miner will include the transaction in a block, say block number 270000. After about one hour, five more blocks will have been added to the chain after that block, with each of those blocks indirectly pointing to the transaction and thus “confirming” it. At this point, the merchant will accept the payment as finalized and deliver the product; since we are assuming this is a digital good, delivery is instant. Now, the attacker creates another transaction sending the 100 BTC to himself. If the attacker simply releases it into the wild, the transaction will not be processed; miners will attempt to run APPLY(S,TX) and notice that TX consumes a UTXO which is no longer in the state. So instead, the attacker creates a “fork” of the blockchain, starting by mining another version of block 270000 pointing to the same block 269999 as a parent but with the new transaction in place of the old one. Because the block data is different, this requires redoing the proof of work. Furthermore, the attacker’s new version of block 270000 has a different hash, so the original blocks 270001 to 270005 do not “point” to it; thus, the original chain and the attacker’s new chain are completely separate. The rule is that in a fork the longest blockchain (i.e., the one backed by the largest quantity of proof of work) is taken to be the truth, and so legitimate miners will work on the 270005 chain while the attacker alone is working on the 270000 chain. In order for the attacker to make his blockchain the longest, he would need to have more computational power than the rest of the network combined in order to catch up (hence, “51% attack”).

Merkle Trees

Left: it suffices to present only a small number of nodes in a Merkle tree to give a proof of the validity of a branch.

Right: any attempt to change any part of the Merkle tree will eventually lead to an inconsistency somewhere up the chain.

An important scalability feature of Bitcoin is that the block is stored in a multi-level data structure. The “hash” of a block is actually only the hash of the block header, a roughly 200-byte piece of data that contains the timestamp, nonce, previous block hash and the root hash of a data structure called the Merkle tree storing all transactions in the block.

A Merkle tree is a type of binary tree, composed of a set of nodes with a large number of leaf nodes at the bottom of the tree containing the underlying data, a set of intermediate nodes where each node is the hash of its two children, and finally a single root node, also formed from the hash of its two children, representing the “top” of the tree. The purpose of the Merkle tree is to allow the data in a block to be delivered piecemeal: a node can download only the header of a block from one source, the small part of the tree relevant to them from another source, and still be assured that all of the data is correct. The reason why this works is that hashes propagate upward: if a malicious user attempts to swap in a fake transaction into the bottom of a Merkle tree, this change will cause a change in the node above, and then a change in the node above that, finally changing the root of the tree and therefore the hash of the block, causing the protocol to register it as a completely different block (almost certainly with an invalid proof of work).

The Merkle tree protocol is arguably essential to long-term sustainability. A “full node” in the Bitcoin network, one that stores and processes the entirety of every block, takes up about 15 GB of disk space in the Bitcoin network as of April 2014, and is growing by over a gigabyte per month. Currently, this is viable for some desktop computers and not phones, and later on in the future only businesses and

hobbyists will be able to participate. A protocol known as “simplified payment verification” (SPV) allows for another class of nodes to exist, called “light nodes”, which download the block headers, verify the proof of work on the block headers, and then download only the “branches” associated with transactions that are relevant to them. This allows light nodes to determine with a strong guarantee of security what the status of any Bitcoin transaction, and their current balance, is while downloading only a very small portion of the entire blockchain.

Alternative Blockchain Applications

The idea of taking the underlying blockchain idea and applying it to other concepts also has a long history. In 2005, Nick Szabo came out with the concept of “secure property titles with owner authority”, a document describing how “new advances in replicated database technology” will allow for a blockchain-based system for storing a registry of who owns what land, creating an elaborate framework including concepts such as homesteading, adverse possession and Georgian land tax. However, there was unfortunately no effective replicated database system available at the time, and so the protocol was never implemented in practice. After 2009, however, once Bitcoin’s decentralized consensus was developed a number of alternative applications rapidly began to emerge:

Namecoin—created in 2010, Namecoin is best described as a decentralized name registration database. In decentralized protocols like Tor, Bitcoin and BitMessage, there needs to be some way of identifying accounts so that other people can interact with them, but in all existing solutions the only kind of identifier available is a pseudorandom hash like 1LW79wp5ZBqaHW1jL5TCiBCrHqYtHagUWy. Ideally, one would like to be able to have an account with a name like “george”. However, the problem is that if one person can create an account named “george” then someone else can use the same process to register “george” for themselves as well and impersonate them. The only solution is a first-to-file paradigm, where the first registrant succeeds and the second fails—a problem perfectly suited for the Bitcoin consensus protocol. Namecoin is the oldest, and most successful, implementation of a name registration system using such an idea.

Colored coins—the purpose of colored coins is to serve as a protocol to allow people to create their own digital currencies—or, in the important trivial case of a currency with one unit, digital tokens, on the Bitcoin blockchain. In the colored coins protocol, one “issues” a new currency by publicly assigning a color to a specific Bitcoin UTXO, and the protocol recursively defines the color of other UTXO to be the same as the color of the inputs that the transaction creating them spent (some special rules apply in the case of mixed-color inputs). This allows users to maintain wallets containing only UTXO of a specific color and send them around much like regular bitcoins, backtracking through the blockchain to determine the color of any UTXO that they receive.

Metacoins—the idea behind a metacoin is to have a protocol that lives on top of Bitcoin, using Bitcoin transactions to store metacoin transactions but having a different state transition function, APPLY’. Because the metacoin protocol cannot prevent invalid metacoin transactions from appearing in the Bitcoin blockchain, a rule is added that if APPLY’(S,TX) returns an error, the protocol defaults to APPLY’(S,TX)=S. This pro-

vides an easy mechanism for creating an arbitrary cryptocurrency protocol, potentially with advanced features that cannot be implemented inside of Bitcoin itself, but with a very low development cost since the complexities of mining and networking are already 5 handled by the Bitcoin protocol.

Thus, in general, there are two approaches toward building a consensus protocol: building an independent network, and building a protocol on top of Bitcoin. The former approach, while reasonably successful in the case of applications like Namecoin, is difficult to implement; each individual implementation needs to bootstrap an independent blockchain, as well as building and testing all of the necessary state transition and networking code. Additionally, we predict that the set of applications for decentralized consensus technology will follow a power law distribution where the vast majority of applications would be too small to warrant their own blockchain, and we note that there exist large classes of decentralized applications, particularly decentralized autonomous organizations, that need to interact with each other. 20

The Bitcoin-based approach, on the other hand, has the flaw that it does not inherit the simplified payment verification features of Bitcoin. SPV works for Bitcoin because it can use blockchain depth as a proxy for validity; at some point, once the ancestors of a transaction go far enough back, it is safe to say that they were legitimately part of the state. Blockchain-based meta-protocols, on the other hand, cannot force the blockchain not to include transactions that are not valid within the context of their own protocols. Hence, a fully secure SPV meta-protocol implementation would need to backward scan all the way to the beginning of the Bitcoin blockchain to determine whether or not certain transactions are valid. Currently, all “light” implementations of Bitcoin-based meta-protocols rely on a trusted server to provide the data, arguably a highly suboptimal result especially when one of the primary purposes of a cryptocurrency is to eliminate the need for trust. 35

Scripting

Even without any extensions, the Bitcoin protocol actually does facilitate a weak version of a concept of “smart contracts”. UTXO in Bitcoin can be owned not just by a public key, but also by a more complicated script expressed in a simple stack-based programming language. In this paradigm, a transaction spending that UTXO must provide data that satisfies the script. Indeed, even the basic public key ownership mechanism is implemented via a script: the script takes an elliptic curve signature as input, verifies it against the transaction and the address that owns the UTXO, and returns 1 if the verification is successful and 0 otherwise. Other, more complicated, scripts exist for various additional use cases. For example, one can construct a script that requires signatures from two out of a given three private keys to validate (“multisig”), a setup useful for corporate accounts, secure savings accounts and some merchant escrow situations. Scripts can also be used to pay bounties for solutions to computational problems, and one can even construct a script that says something like “this Bitcoin UTXO is yours if you can provide an SPV proof that you sent a Dogecoin transaction of this denomination to me”, essentially allowing decentralized cross-cryptocurrency exchange. 50

However, the scripting language as implemented in Bitcoin has several important limitations:

Lack of Turing-completeness—that is to say, while there is a large subset of computation that the Bitcoin scripting language supports, it does not nearly support every-

thing. The main category that is missing is loops. This is done to avoid infinite loops during transaction verification; theoretically it is a surmountable obstacle for script programmers, since any loop can be simulated by simply repeating the underlying code many times with an if statement, but it does lead to scripts that are very space-inefficient. For example, implementing an alternative elliptic curve signature algorithm would likely require 256 repeated multiplication rounds all individually included in the code.

Value-blindness—there is no way for a UTXO script to provide fine-grained control over the amount that can be withdrawn. For example, one powerful use case of an oracle contract would be a hedging contract, where A and B put in \$1000 worth of BTC and after 30 days the script sends \$1000 worth of BTC to A and the rest to B. This would require an oracle to determine the value of 1 BTC in USD, but even then it is a massive improvement in terms of trust and infrastructure requirement over the fully centralized solutions that are available now. However, because UTXO are all-or-nothing, the only way to achieve this is through the very inefficient hack of having many UTXO of varying denominations (e.g., one UTXO of 2k for every k up to 30) and having the oracle pick which UTXO to send to A and which to B.

Lack of state—UTXO can either be spent or unspent; there is no opportunity for multi-stage contracts or scripts which keep any other internal state beyond that. This makes it hard to make multi-stage options contracts, decentralized exchange offers or two-stage cryptographic commitment protocols (necessary for secure computational bounties). It also means that UTXO can only be used to build simple, one-off contracts and not more complex “stateful” contracts such as decentralized organizations, and makes meta-protocols difficult to implement. Binary state combined with value-blindness also mean that another important application, withdrawal limits, is impossible.

Blockchain-blindness—UTXO are blind to blockchain data such as the nonce and previous hash. This severely limits applications in gambling, and several other categories, by depriving the scripting language of a potentially valuable source of randomness.

Thus, we see three approaches to building advanced applications on top of cryptocurrency: building a new blockchain, using scripting on top of Bitcoin, and building a meta-protocol on top of Bitcoin. Building a new blockchain allows for unlimited freedom in building a feature set, but at the cost of development time and bootstrapping effort. Using scripting is easy to implement and standardize, but is very limited in its capabilities, and meta-protocols, while easy, suffer from faults in scalability. With Ethereum, we intend to build a generalized framework that can provide the advantages of all three paradigms at the same time. 55

Ethereum

The intent of Ethereum is to merge together and improve upon the concepts of scripting, altcoins and on-chain meta-protocols, and allow developers to create arbitrary consensus-based applications that have the scalability, standardization, feature-completeness, ease of development and interoperability offered by these different paradigms all at the same time. Ethereum does this by building what is essentially the ultimate abstract foundational layer: a blockchain with a built-in Turing-complete programming language, allowing anyone to write smart contracts and decentralized applications where they can create their own

arbitrary rules for ownership, transaction formats and state transition functions. A bare-bones version of Namecoin can be written in two lines of code, and other protocols like currencies and reputation systems can be built in under twenty. Smart contracts, cryptographic “boxes” that contain value and only unlock it if certain conditions are met, can also be built on top of our platform, with vastly more power than that offered by Bitcoin scripting because of the added powers of Turing-completeness, value-awareness, block-chain-awareness and state.

Ethereum Accounts

In Ethereum, the state is made up of objects called “accounts”, with each account having a 20-byte address and state transitions being direct transfers of value and information between accounts. An Ethereum account contains four fields:

The nonce, a counter used to make sure each transaction can only be processed once

The account’s current ether balance

The account’s contract code, if present

The account’s storage (empty by default)

“Ether” is the main internal crypto-fuel of Ethereum, and is used to pay transaction fees. In general, there are two types of accounts: externally owned accounts, controlled by private keys, and contract accounts, controlled by their contract code. An externally owned account has no code, and one can send messages from an externally owned account by creating and signing a transaction; in a contract account, every time the contract account receives a message its code activates, allowing it to read and write to internal storage and send other messages or create contracts in turn.

Messages and Transactions

“Messages” in Ethereum are somewhat similar to “transactions” in Bitcoin, but with three important differences. First, an Ethereum message can be created either by an external entity or a contract, whereas a Bitcoin transaction can only be created externally. Second, there is an explicit option for Ethereum messages to contain data. Finally, the recipient of an Ethereum message, if it is a contract account, has the option to return a response; this means that Ethereum messages also encompass the concept of functions.

The term “transaction” is used in Ethereum to refer to the signed data package that stores a message to be sent from an externally owned account. Transactions contain the recipient of the message, a signature identifying the sender, the amount of ether and the data to send, as well as two values called STARTGAS and GASPRICE. In order to prevent exponential blowup and infinite loops in code, each transaction is required to set a limit to how many computational steps of code execution it can spawn, including both the initial message and any additional messages that get spawned during execution. STARTGAS is this limit, and GASPRICE is the fee to pay to the miner per computational step. If transaction execution “runs out of gas”, all state changes revert—except for the payment of the fees, and if transaction execution halts with some gas remaining then the remaining portion of the fees is refunded to the sender. There is also a separate transaction type, and corresponding message type, for creating a contract; the address of a contract is calculated based on the hash of the account nonce and transaction data.

An important consequence of the message mechanism is the “first class citizen” property of Ethereum—the idea that contracts have equivalent powers to external accounts, including the ability to send message and create other contracts. This allows contracts to simultaneously serve many different roles: for example, one might have a member

of a decentralized organization (a contract) be an escrow account (another contract) between an paranoid individual employing custom quantum-proof Lamport signatures (a third contract) and a co-signing entity which itself uses an account with five keys for security (a fourth contract). The strength of the Ethereum platform is that the decentralized organization and the escrow contract do not need to care about what kind of account each party to the contract is.

Ethereum State Transition Function

The Ethereum state transition function, $APPLY(S, TX) \rightarrow S'$ can be defined as follows:

1. Check if the transaction is well-formed (i.e., has the right number of values), the signature is valid, and the nonce matches the nonce in the sender’s account. If not, return an error.

2. Calculate the transaction fee as $STARTGAS * GASPRICE$, and determine the sending address from the signature. Subtract the fee from the sender’s account balance and increment the sender’s nonce. If there is not enough balance to spend, return an error.

3. Initialize $GAS = STARTGAS$, and take off a certain quantity of gas per byte to pay for the bytes in the transaction.

4. Transfer the transaction value from the sender’s account to the receiving account. If the receiving account does not yet exist, create it. If the receiving account is a contract, run the contract’s code either to completion or until the execution runs out of gas.

5. If the value transfer failed because the sender did not have enough money, or the code execution ran out of gas, revert all state changes except the payment of the fees, and add the fees to the miner’s account.

6. Otherwise, refund the fees for all remaining gas to the sender, and send the fees paid for gas consumed to the miner.

For example, suppose that the contract’s code is:

```
if !contract.storage[msg.data[0]]:
    contract.storage[msg.data[0]] = msg.data[1]
```

Note that in reality the contract code is written in the low-level EVM code; this example is written in Serpent, our high-level language, for clarity, and can be compiled down to EVM code. Suppose that the contract’s storage starts off empty, and a transaction is sent with 10 ether value, 2000 gas, 0.001 ether gasprice, and two data fields: [2, ‘CHARLIE’][3]. The process for the state transition function in this case is as follows:

1. Check that the transaction is valid and well formed.

2. Check that the transaction sender has at least $2000 * 0.001 = 2$ ether. If it is, then subtract 2 ether from the sender’s account.

3. Initialize $gas = 2000$; assuming the transaction is 170 bytes long and the byte-fee is 5, subtract 850 so that there is 1150 gas left.

4. Subtract 10 more ether from the sender’s account, and add it to the contract’s account.

5. Run the code. In this case, this is simple: it checks if the contract’s storage at index 2 is used, notices that it is not, and so it sets the storage at index 2 to the value CHARLIE., Suppose this takes 187 gas, so the remaining amount of gas is $1150 - 187 = 963$

6. Add $963 * 0.001 = 0.963$ ether back to the sender’s account, and return the resulting state.

If there was no contract at the receiving end of the transaction, then the total transaction fee would simply be equal to the provided GASPRICE multiplied by the length of the transaction in bytes, and the data sent alongside the transaction would be irrelevant. Additionally, note that contract-initiated messages can assign a gas limit to the com-

putation that they spawn, and if the sub-computation runs out of gas it gets reverted only to the point of the message call. Hence, just like transactions, contracts can secure their limited computational resources by setting strict limits on the sub-computations that they spawn.

Code Execution

The code in Ethereum contracts is written in a low-level, stack-based bytecode language, referred to as “Ethereum virtual machine code” or “EVM code”. The code consists of a series of bytes, where each byte represents an operation. In general, code execution is an infinite loop that consists of repeatedly carrying out the operation at the current program counter (which begins at zero) and then incrementing the program counter by one, until the end of the code is reached or an error or STOP or RETURN instruction is detected. The operations have access to three types of space in which to store data:

The stack, a last-in-first-out container to which 32-byte values can be pushed and popped

Memory, an infinitely expandable byte array

The contract’s long-term storage, a key/value store where keys and values are both 32 bytes. Unlike stack and memory, which reset after computation ends, storage persists for the long term.

The code can also access the value, sender and data of the incoming message, as well as block header data, and the code can also return a byte array of data as an output.

The formal execution model of EVM code is surprisingly simple. While the Ethereum virtual machine is running, its full computational state can be defined by the tuple (block_state, transaction, message, code, memory, stack, pc, gas), where block_state is the global state containing all accounts and includes balances and storage. Every round of execution, the current instruction is found by taking the pc-th byte of code, and each instruction has its own definition in terms of how it affects the tuple. For example, ADD pops two items off the stack and pushes their sum, reduces gas by 1 and increments pc by 1, and SSTORE pushes the top two items off the stack and inserts the second item into the contract’s storage at the index specified by the first item, as well as reducing gas by up to 200 and incrementing pc by 1. Although there are many ways to optimize Ethereum via just-in-time compilation, a basic implementation of Ethereum can be done in a few hundred lines of code.

Blockchain and Mining

The Ethereum blockchain is in many ways similar to the Bitcoin blockchain, although it does have some differences. The main difference between Ethereum and Bitcoin with regard to the blockchain architecture is that, unlike Bitcoin, Ethereum blocks contain a copy of both the transaction list and the most recent state. Aside from that, two other values, the block number and the difficulty, are also stored in the block. The block validation algorithm in Ethereum is as follows:

1. Check if the previous block referenced exists and is valid.
2. Check that the timestamp of the block is greater than that of the referenced previous block and less than 15 minutes into the future
3. Check that the block number, difficulty, transaction root, uncle root and gas limit (various low-level Ethereum-specific concepts) are valid.
4. Check that the proof of work on the block is valid.
5. Let S[0] be the STATE_ROOT of the previous block.
6. Let TX be the block’s transaction list, with n transactions. For all i in $0 \dots n-1$, set $S[i+1]=APPLY(S[i],TX[i])$.

If any application returns an error, or if the total gas consumed in the block up until this point exceeds the GASLIMIT, return an error.

7. Let S_FINAL be S[n], but adding the block reward paid to the miner.

8. Check if S_FINAL is the same as the STATE_ROOT. If it is, the block is valid; otherwise, it is not valid.

The approach may seem highly inefficient at first glance, because it needs to store the entire state with each block, but in reality efficiency should be comparable to that of Bitcoin. The reason is that the state is stored in the tree structure, and after every block only a small part of the tree needs to be changed. Thus, in general, between two adjacent blocks the vast majority of the tree should be the same, and therefore the data can be stored once and referenced twice using pointers (i.e., hashes of subtrees). A special kind of tree known as a “Patricia tree” is used to accomplish this, including a modification to the Merkle tree concept that allows for nodes to be inserted and deleted, and not just changed, efficiently. Additionally, because all of the state information is part of the last block, there is no need to store the entire blockchain history—a strategy which, if it could be applied to Bitcoin, can be calculated to provide 5-20x savings in space.

Applications

In general, there are three types of applications on top of Ethereum. The first category is financial applications, providing users with more powerful ways of managing and entering into contracts using their money. This includes sub-currencies, financial derivatives, hedging contracts, savings wallets, wills, and ultimately even some classes of full-scale employment contracts. The second category is semi-financial applications, where money is involved but there is also a heavy non-monetary side to what is being done; a perfect example is self-enforcing bounties for solutions to computational problems. Finally, there are applications such as online voting and decentralized governance that are not financial at all.

Token Systems

On-blockchain token systems have many applications ranging from sub-currencies representing assets such as USD or gold to company stocks, individual tokens representing smart property, secure unforgeable coupons, and even token systems with no ties to conventional value at all, used as point systems for incentivization. Token systems are surprisingly easy to implement in Ethereum. The key point to understand is that all a currency, or token system, fundamentally is a database with one operation: subtract X units from A and give X units to B, with the proviso that (1) X had at least X units before the transaction and (2) the transaction is approved by A. All that it takes to implement a token system is to implement this logic into a contract.

The basic code for implementing a token system in Serpent looks as follows:

```

55 from =msg.sender to =msg.data[0] value=msg.data[1]
   if contract.storage[from]>=value:
       contract.storage[from]=contract.storage[from]-value
       contract.storage[to]=contract.storage[to]+value

```

This is essentially a literal implementation of the “banking system” state transition function described further above in this document. A few extra lines of code need to be added to provide for the initial step of distributing the currency units in the first place and a few other edge cases, and ideally a function would be added to let other contracts query for the balance of an address. But that’s all there is to it. Theoretically, Ethereum-based token systems acting as sub-currencies can potentially include another important feature that

on-chain Bitcoin-based meta-currencies lack: the ability to pay transaction fees directly in that currency. The way this would be implemented is that the contract would maintain an ether balance with which it would refund ether used to pay fees to the sender, and it would refill this balance by collecting the internal currency units that it takes in fees and reselling them in a constant running auction. Users would thus need to “activate” their accounts with ether, but once the ether is there it would be reusable because the contract would refund it each time.

Financial Derivatives and Stable-Value Currencies

Financial derivatives are the most common application of a “smart contract”, and one of the simplest to implement in code. The main challenge in implementing financial contracts is that the majority of them require reference to an external price ticker; for example, a very desirable application is a smart contract that hedges against the volatility of ether (or another cryptocurrency) with respect to the US dollar, but doing this requires the contract to know what the value of ETH/USD is. The simplest way to do this is through a “data feed” contract maintained by a specific party (e.g., NASDAQ) designed so that that party has the ability to update the contract as needed, and providing an interface that allows other contracts to send a message to that contract and get back a response that provides the price.

Given that critical ingredient, the hedging contract would look as follows:

1. Wait for party A to input 1000 ether.
2. Wait for party B to input 1000 ether.
3. Record the USD value of 1000 ether, calculated by querying the data feed contract, in storage, say this is \$x.
4. After 30 days, allow A or B to “ping” the contract in order to send \$x worth of ether (calculated by querying the data feed contract again to get the new price) to A and the rest to B.

Such a contract would have significant potential in cryptocurrency. One of the main problems cited about cryptocurrency is the fact that it’s volatile; although many users and merchants may want the security and convenience of dealing with cryptographic assets, they many not wish to face that prospect of losing 23% of the value of their funds in a single day. Up until now, the most commonly proposed solution has been issuer-backed assets; the idea is that an issuer creates a sub-currency in which they have the right to issue and revoke units, and provide one unit of the currency to anyone who provides them (offline) with one unit of a specified underlying asset (e.g., gold, USD). The issuer then promises to provide one unit of the underlying asset to anyone who sends back one unit of the crypto-asset. This mechanism allows any non-cryptographic asset to be “uplifted” into a cryptographic asset, provided that the issuer can be trusted.

In practice, however, issuers are not always trustworthy, and in some cases the banking infrastructure is too weak, or too hostile, for such services to exist. Financial derivatives provide an alternative. Here, instead of a single issuer providing the funds to back up an asset, a decentralized market of speculators, betting that the price of a cryptographic reference asset will go up, plays that role. Unlike issuers, speculators have no option to default on their side of the bargain because the hedging contract holds their funds in escrow. Note that this approach is not fully decentralized, because a trusted source is still needed to provide the price ticker, although arguably even still this is a massive improvement in terms of reducing infrastructure requirements (unlike being an issuer, issuing a price feed requires

no licenses and can likely be categorized as free speech) and reducing the potential for fraud.

Identity and Reputation Systems

The earliest alternative cryptocurrency of all, Namecoin, attempted to use a Bitcoin-like blockchain to provide a name registration system, where users can register their names in a public database alongside other data. The major cited use case is for a DNS system, mapping domain names like “bitcoin.org” (or, in Namecoin’s case, “bitcoin.bit”) to an IP address. Other use cases include email authentication and potentially more advanced reputation systems. Here is the basic contract to provide a Namecoin-like name registration system on Ethereum:

```
if !contract.storage[tx.data[0]]:
    contract.storage[tx.data[0]]=tx.data[1]
```

The contract is very simple; all it is a database inside the Ethereum network that can be added to, but not modified or removed from. Anyone can register a name with some value, and that registration then sticks forever. A more sophisticated name registration contract will also have a “function clause” allowing other contracts to query it, as well as a mechanism for the “owner” (i.e., the first registerer) of a name to change the data or transfer ownership. One can even add reputation and web-of-trust functionality on top.

Decentralized File Storage

Over the past few years, there have emerged a number of popular online file storage startups, the most prominent being Dropbox, seeking to allow users to upload a backup of their hard drive and have the service store the backup and allow the user to access it in exchange for a monthly fee. However, at this point the file storage market is at times relatively inefficient; a cursory look at various existing solutions shows that, particularly at the “uncanny valley” 20-200 GB level at which neither free quotas nor enterprise-level discounts kick in, monthly prices for mainstream file storage costs are such that you are paying for more than the cost of the entire hard drive in a single month. Ethereum contracts can allow for the development of a decentralized file storage ecosystem, where individual users can earn small quantities of money by renting out their own hard drives and unused space can be used to further drive down the costs of file storage.

The key underpinning piece of such a device would be what we have termed the “decentralized Dropbox contract”. This contract works as follows. First, one splits the desired data up into blocks, encrypting each block for privacy, and builds a Merkle tree out of it. One then makes a contract with the rule that, every N blocks, the contract would pick a random index in the Merkle tree (using the previous block hash, accessible from contract code, as a source of randomness), and give X ether to the first entity to supply a transaction with a simplified payment verification-like proof of ownership of the block at that particular index in the tree. When a user wants to re-download their file, they can use a micropayment channel protocol (e.g., pay 1 szabo per 32 kilobytes) to recover the file; the most fee-efficient approach is for the payer not to publish the transaction until the end, instead replacing the transaction with a slightly more lucrative one with the same nonce after every 32 kilobytes.

An important feature of the protocol is that, although it may seem like one is trusting many random nodes not to decide to forget the file, one can reduce that risk down to near-zero by splitting the file into many pieces via secret sharing, and watching the contracts to see each piece is still in some node’s possession. If a contract is still paying out money, that provides a cryptographic proof that someone out there is still storing the file.

Decentralized Autonomous Organizations

The general concept of a “decentralized organization” is that of a virtual entity that has a certain set of members or shareholders which, perhaps with a 67% majority, have the right to spend the entity’s funds and modify its code. The members would collectively decide on how the organization should allocate its funds. Methods for allocating a DAO’s funds could range from bounties, salaries to even more exotic mechanisms such as an internal currency to reward work. This essentially replicates the legal trappings of a traditional company or nonprofit but using only cryptographic blockchain technology for enforcement. So far much of the talk around DAOs has been around the “capitalist” model of a “decentralized autonomous corporation” (DAC) with dividend-receiving shareholders and tradable shares; an alternative, perhaps described as a “decentralized autonomous community”, would have all members have an equal share in the decision making and require 67% of existing members to agree to add or remove a member. The requirement that one person can only have one membership would then need to be enforced collectively by the group.

A general outline for how to code a DO is as follows. The simplest design is simply a piece of self-modifying code that changes if two thirds of members agree on a change. Although code is theoretically immutable, one can easily get around this and have de-facto mutability by having chunks of the code in separate contracts, and having the address of which contracts to call stored in the modifiable storage. In a simple implementation of such a DAO contract, there would be three transaction types, distinguished by the data provided in the transaction:

[0,i,K,V] to register a proposal with index i to change the address at storage index K to value V

[0,i] to register a vote in favor of proposal i

[2,i] to finalize proposal i if enough votes have been made

The contract would then have clauses for each of these. It would maintain a record of all open storage changes, along with a list of who voted for them. It would also have a list of all members. When any storage change gets to two thirds of members voting for it, a finalizing transaction could execute the change. A more sophisticated skeleton would also have built-in voting ability for features like sending a transaction, adding members and removing members, and may even provide for Liquid Democracy-style vote delegation (i.e., anyone can assign someone to vote for them, and assignment is transitive so if A assigns B and B assigns C then C determines A’s vote). This design would allow the DO to grow organically as a decentralized community, allowing people to eventually delegate the task of filtering out who is a member to specialists, although unlike in the “current system” specialists can easily pop in and out of existence over time as individual community members change their alignments.

An alternative model is for a decentralized corporation, where any account can have zero or more shares, and two thirds of the shares are required to make a decision. A complete skeleton would involve asset management functionality, the ability to make an offer to buy or sell shares, and the ability to accept offers (preferably with an order-matching mechanism inside the contract). Delegation would also exist Liquid Democracy-style, generalizing the concept of a “board of directors”.

In the future, more advanced mechanisms for organizational governance may be implemented; it is at this point that a decentralized organization (DO) can start to be described as a decentralized autonomous organization (DAO). The difference between a DO and a DAO is fuzzy, but the general

dividing line is whether the governance is generally carried out via a political-like process or an “automatic” process; a good intuitive test is the “no common language” criterion: can the organization still function if no two members spoke the same language? Clearly, a simple traditional shareholder-style corporation would fail, whereas something like the Bitcoin protocol would be much more likely to succeed. Robin Hanson’s futarchy, a mechanism for organizational governance via prediction markets, is a good example of what truly “autonomous” governance might look like. Note that one should not necessarily assume that all DAOs are superior to all DOs; automation is simply a paradigm that is likely to have very large benefits in certain particular places and may not be practical in others, and many semi-DAOs are also likely to exist.

Further Applications

1. Savings wallets. Suppose that Alice wants to keep her funds safe, but is worried that she will lose or someone will hack her private key. She puts ether into a contract with Bob, a bank, as follows:

Alice alone can withdraw a maximum of 1% of the funds per day.

Bob alone can withdraw a maximum of 1% of the funds per day, but Alice has the ability to make a transaction with her key shutting off this ability.

Alice and Bob together can withdraw anything.

Normally, 1% per day is enough for Alice, and if Alice wants to withdraw more she can contact Bob for help. If Alice’s key gets hacked, she runs to Bob to move the funds to a new contract. If she loses her key, Bob will get the funds out eventually. If Bob turns out to be malicious, then she can turn off his ability to withdraw.

2. Crop insurance. One can easily make a financial derivatives contract but using a data feed of the weather instead of any price index. If a farmer in Iowa purchases a derivative that pays out inversely based on the precipitation in Iowa, then if there is a drought, the farmer will automatically receive money and if there is enough rain the farmer will be happy because their crops would do well.

3. A decentralized data feed. For financial contracts for difference, it may actually be possible to decentralize the data feed via a protocol called “SchellingCoin”. SchellingCoin basically works as follows: N parties all put into the system the value of a given datum (e.g., the ETH/USD price), the values are sorted, and everyone between the 25th and 75th percentile gets one token as a reward. Everyone has the incentive to provide the answer that everyone else will provide, and the only value that a large number of players can realistically agree on is the obvious default: the truth. This creates a decentralized protocol that can theoretically provide any number of values, including the ETH/USD price, the temperature in Berlin or even the result of a particular hard computation.

4. Smart multi-signature escrow. Bitcoin allows multisignature transaction contracts where, for example, three out of a given five keys can spend the funds. Ethereum allows for more granularity; for example, four out of five can spend everything, three out of five can spend up to 10% per day, and two out of five can spend up to 0.5% per day. Additionally, Ethereum multisig is asynchronous—two parties can register their signatures on the blockchain at different times and the last signature will automatically send the transaction.

5. Cloud computing. The EVM technology can also be used to create a verifiable computing environment, allowing users to ask others to carry out computations and then optionally ask for proofs that computations at certain ran-

domly selected checkpoints were done correctly. This allows for the creation of a cloud computing market where any user can participate with their desktop, laptop or specialized server, and spot-checking together with security deposits can be used to ensure that the system is trustworthy (i.e., nodes cannot profitably cheat). Although such a system may not be suitable for all tasks; tasks that require a high level of inter-process communication, for example, cannot easily be done on a large cloud of nodes. Other tasks, however, are much easier to parallelize; projects like SETI@home, folding@home and genetic algorithms can easily be implemented on top of such a platform.

6. Peer-to-peer gambling. Any number of peer-to-peer gambling protocols, such as Frank Stajano and Richard Clayton's Cyberdice, can be implemented on the Ethereum blockchain. The simplest gambling protocol is actually simply a contract for difference on the next block hash, and more advanced protocols can be built up from there, creating gambling services with near-zero fees that have no ability to cheat.

7. Prediction markets. Provided an oracle or SchellingCoin, prediction markets are also easy to implement, and prediction markets together with SchellingCoin may prove to be the first mainstream application of futarchy as a governance protocol for decentralized organizations.

8. On-chain decentralized marketplaces, using the identity and reputation system as a base.

Miscellanea And Concerns

Modified GHOST Implementation

The "Greedy Heaviest Observed Subtree" (GHOST) protocol is an innovation first introduced by Yonatan Sompolinsky and Aviv Zohar in December 2013. Note however that a recent paper [31] argues that, while uncle blocks do provide block rewards to miners, they do not contribute towards the difficulty of the main chain. Therefore, Ethereum does not actually apply the GHOST protocol. The motivation behind GHOST is that blockchains with fast confirmation times currently suffer from reduced security due to a high stale rate—because blocks take a certain time to propagate through the network, if miner A mines a block and then miner B happens to mine another block before miner A's block propagates to B, miner B's block will end up wasted and will not contribute to network security. Furthermore, there is a centralization issue: if miner A is a mining pool with 30% hashpower and B has 10% hashpower, A will have a risk of producing a stale block 70% of the time (since the other 30% of the time A produced the last block and so will get mining data immediately) whereas B will have a risk of producing a stale block 90% of the time. Thus, if the block interval is short enough for the stale rate to be high, A will be substantially more efficient simply by virtue of its size. With these two effects combined, blockchains which produce blocks quickly are very likely to lead to one mining pool having a large enough percentage of the network hashpower to have de facto control over the mining process.

As described by Sompolinsky and Zohar, GHOST solves the first issue of network security loss by including stale blocks in the calculation of which chain is the "longest"; that is to say, not just the parent and further ancestors of a block, but also the stale children of the block's ancestors (in Ethereum jargon, "uncles") are added to the calculation of which block has the largest total proof of work backing it. To solve the second issue of centralization bias, we go beyond the protocol described by Sompolinsky and Zohar, and also allow stales to be registered into the main chain to receive a block reward: a stale block receives 93.75% of its base

reward, and the nephew that includes the stale block receives the remaining 6.25%. Transaction fees, however, are not awarded to uncles.

Ethereum implements a simplified version of GHOST which only goes down five levels. Specifically, a stale block can only be included as an uncle by the 2nd to 5th generation child of its parent, and not any block with a more distant relation (e.g., 6th generation child of a parent, or 3rd generation child of a grandparent). This was done for several reasons. First, unlimited GHOST would include too many complications into the calculation of which uncles for a given block are valid. Second, unlimited GHOST with compensation as used in Ethereum removes the incentive for a miner to mine on the main chain and not the chain of a public attacker. Finally, calculations show that five-level GHOST with incentivization is over 95% efficient even with a 15s block time, and miners with 25% hashpower show centralization gains of less than 3%.

Fees

Because every transaction published into the blockchain imposes on the network the cost of needing to download and verify it, there is a need for some regulatory mechanism, typically involving transaction fees, to prevent abuse. The default approach, used in Bitcoin, is to have purely voluntary fees, relying on miners to act as the gatekeepers and set dynamic minimums. This approach has been received very favorably in the Bitcoin community particularly because it is "market-based", allowing supply and demand between miners and transaction senders determine the price. The problem with this line of reasoning is, however, that transaction processing is not a market; although it is intuitively attractive to construe transaction processing as a service that the miner is offering to the sender, in reality every transaction that a miner includes will need to be processed by every node in the network, so the vast majority of the cost of transaction processing is borne by third parties and not the miner that is making the decision of whether or not to include it. Hence, tragedy-of-the-commons problems are very likely to occur.

However, as it turns out this flaw in the market-based mechanism, when given a particular inaccurate simplifying assumption, magically cancels itself out. The argument is as follows. Suppose that:

1. A transaction leads to k operations, offering the reward kR to any miner that includes it where R is set by the sender and k and R are (roughly) visible to the miner beforehand.
2. An operation has a processing cost of C to any node (i.e., all nodes have equal efficiency)
3. There are N mining nodes, each with exactly equal processing power (i.e., $1/N$ of total)
4. No non-mining full nodes exist.

A miner would be willing to process a transaction if the expected reward is greater than the cost. Thus, the expected reward is kR/N since the miner has a $1/N$ chance of processing the next block, and the processing cost for the miner is simply kC . Hence, miners will include transactions where $kR/N > kC$, or $R > NC$. Note that R is the per-operation fee provided by the sender, and is thus a lower bound on the benefit that the sender derives from the transaction, and NC is the cost to the entire network together of processing an operation. Hence, miners have the incentive to include only those transactions for which the total utilitarian benefit exceeds the cost.

However, there are several important deviations from those assumptions in reality:

1. The miner does pay a higher cost to process the transaction than the other verifying nodes, since the extra

verification time delays block propagation and thus increases the chance the block will become a stale.

2. There do exist non-mining full nodes.

3. The mining power distribution may end up radically inegalitarian in practice.

4. Speculators, political enemies and crazies whose utility function includes causing harm to the network do exist, and they can cleverly set up contracts whose cost is much lower than the cost paid by other verifying nodes.

Point 1 above provides a tendency for the miner to include fewer transactions, and point 2 increases NC; hence, these two effects at least partially cancel each other out. Points 3 and 4 are the major issue; to solve them we simply institute a floating cap: no block can have more operations than `BLK_LIMIT_FACTOR` times the long-term exponential moving average. Specifically:

```
blk.oplimit=floor((blk.parent.oplimit*(EMAFAC-
1)+floor(parent.opcount*BLK_LIMIT_FACTOR))/EMA_
FACTOR)
```

`BLK_LIMIT_FACTOR` and `EMAFAC` are constants that will be set to 65536 and 1.5 for the time being, but will likely be changed after further analysis.

Computation And Turing-Completeness

An important note is that the Ethereum virtual machine is Turing-complete; this means that EVM code can encode any computation that can be conceivably carried out, including infinite loops. EVM code allows looping in two ways. First, there is a JUMP instruction that allows the program to jump back to a previous spot in the code, and a JUMPI instruction to do conditional jumping, allowing for statements like while `x<27: x=x*2`. Second, contracts can call other contracts, potentially allowing for looping through recursion. This naturally leads to a problem: can malicious users essentially shut miners and full nodes down by forcing them to enter into an infinite loop? The issue arises because of a problem in computer science known as the halting problem: there is no way to tell, in the general case, whether or not a given program will ever halt.

As described in the state transition section, our solution works by requiring a transaction to set a maximum number of computational steps that it is allowed to take, and if execution takes longer computation is reverted but fees are still paid. Messages work in the same way. To show the motivation behind our solution, consider the following examples:

An attacker creates a contract which runs an infinite loop, and then sends a transaction activating that loop to the miner. The miner will process the transaction, running the infinite loop, and wait for it to run out of gas. Even though the execution runs out of gas and stops halfway through, the transaction is still valid and the miner still claims the fee from the attacker for each computational step.

An attacker creates a very long infinite loop with the intent of forcing the miner to keep computing for such a long time that by the time computation finishes a few more blocks will have come out and it will not be possible for the miner to include the transaction to claim the fee. However, the attacker will be required to submit a value for `STARTGAS` limiting the number of computational steps that execution can take, so the miner will know ahead of time that the computation will take an excessively large number of steps.

An attacker sees a contract with code of some form like `send (A,contract.storage[A]); contract.storage[A]=0`, and sends a transaction with just enough gas to run the first step but not the second (i.e., making a withdrawal

but not letting the balance go down). The contract author does not need to worry about protecting against such attacks, because if execution stops halfway through the changes get reverted.

5 A financial contract works by taking the median of nine proprietary data feeds in order to minimize risk. An attacker takes over one of the data feeds, which is designed to be modifiable via the variable-address-call mechanism described in the section on DAOs, and converts it to run an infinite loop, thereby attempting to force any attempts to claim funds from the financial contract to run out of gas. However, the financial contract can set a gas limit on the message to prevent this problem.

The alternative to Turing-completeness is Turing-incompleteness, where JUMP and JUMPI do not exist and only one copy of each contract is allowed to exist in the call stack at any given time. With this system, the fee system described and the uncertainties around the effectiveness of our solution might not be necessary, as the cost of executing a contract would be bounded above by its size. Additionally, Turing-incompleteness is not even that big a limitation; out of all the contract examples we have conceived internally, so far only one required a loop, and even that loop could be removed by making 26 repetitions of a one-line piece of code. Given the serious implications of Turing-completeness, and the limited benefit, why not simply have a Turing-incomplete language? In reality, however, Turing-incompleteness is far from a neat solution to the problem. To see why, consider the following contracts:

```
C0: call(C1); call(C1);
C1: call(C2); call(C2);
C2: call(C3); call(C3);
...
C49: call(C50); call(C50);
C50: (run one step of a program and record the change in
storage)
```

Now, send a transaction to A. Thus, in 51 transactions, we have a contract that takes up 250 computational steps. Miners could try to detect such logic bombs ahead of time by maintaining a value alongside each contract specifying the maximum number of computational steps that it can take, and calculating this for contracts calling other contracts recursively, but that would require miners to forbid contracts that create other contracts (since the creation and execution of all 50 contracts above could easily be rolled into a single contract). Another problematic point is that the address field of a message is a variable, so in general it may not even be possible to tell which other contracts a given contract will call ahead of time. Hence, all in all, we have a surprising conclusion: Turing-completeness is surprisingly easy to manage, and the lack of Turing-completeness is equally surprisingly difficult to manage unless the exact same controls are in place—but in that case why not just let the protocol be Turing-complete?

Currency And Issuance

The Ethereum network includes its own built-in currency, ether, which serves the dual purpose of providing a primary liquidity layer to allow for efficient exchange between various types of digital assets and, more importantly, of providing a mechanism for paying transaction fees. For convenience and to avoid future argument (see the current `mBTC/uBTC/satoshi` debate in Bitcoin), the denominations will be pre-labelled:

```
1: wei
10^12: szabo
10^15: finney
10^18: ether
```

This should be taken as an expanded version of the concept of “dollars” and “cents” or “BTC” and “satoshi”. In the near future, we expect “ether” to be used for ordinary transactions, “finney” for microtransactions and “szabo” and “wei” for technical discussions around fees and protocol implementation.

The issuance model will be as follows:

Ether will be released in a currency sale at the price of 1337-2000 ether per BTC, a mechanism intended to fund the Ethereum organization and pay for development that has been used with success by a number of other cryptographic platforms. Earlier buyers will benefit from larger discounts. The BTC received from the sale will be used entirely to pay salaries and bounties to developers, researchers and projects in the cryptocurrency ecosystem.

0.099× the total amount sold will be allocated to early contributors who participated in development before BTC funding or certainty of funding was available, and another 0.099× will be allocated to long-term research projects.

0.26× the total amount sold will be allocated to miners per year forever after that point.

Issuance Breakdown

The permanent linear supply growth model reduces the risk of what some see as excessive wealth concentration in Bitcoin, and gives individuals living in present and future eras a fair chance to acquire currency units, while at the same time discouraging depreciation of ether because the “supply growth rate” as a percentage still tends to zero over time. We also theorize that because coins are always lost over time due to carelessness, death, etc, and coin loss can be modeled as a percentage of the total supply per year, that the total currency supply in circulation will in fact eventually stabilize at a value equal to the annual issuance divided by the loss rate (e.g., at a loss rate of 1%, once the supply reaches 26× then 0.26× will be mined and 0.26× lost every year, creating an equilibrium).

Group At launch After 1 year After 5 years

Currency units 1.198X1.458X2.498X

Purchasers 83.5% 68.6% 40.0%

Early contributor distribution 8.26% 6.79% 3.96%

Long-term endowment 8.26% 6.79% 3.96%

Miners 0% 17.8% 52.0%

Despite the linear currency issuance, just like with Bitcoin over time the supply growth rate nevertheless tends to zero.

Mining Centralization

The Bitcoin mining algorithm basically works by having miners compute SHA256 on slightly modified versions of the block header millions of times over and over again, until eventually one node comes up with a version whose hash is less than the target (currently around 2190). However, this mining algorithm is vulnerable to two forms of centralization. First, the mining ecosystem has come to be dominated by ASICs (application-specific integrated circuits), computer chips designed for, and therefore thousands of times more efficient at, the specific task of Bitcoin mining. This means that Bitcoin mining is no longer a highly decentralized and egalitarian pursuit, requiring millions of dollars of capital to effectively participate in. Second, most Bitcoin miners do not actually perform block validation locally; instead, they rely on a centralized mining pool to provide the block headers. This problem is arguably worse: as of the time of this writing, the top two mining pools indirectly control roughly 50% of processing power in the Bitcoin

network, although this is mitigated by the fact that miners can switch to other mining pools if a pool or coalition attempts a 51% attack.

The current intent at Ethereum is to use a mining algorithm based on randomly generating a unique hash function for every 1000 nonces, using a sufficiently broad range of computation to remove the benefit of specialized hardware. Such a strategy will certainly not reduce the gain of centralization to zero, but it does not need to. Note that each individual user, on their private laptop or desktop, can perform a certain quantity of mining activity almost for free, paying only electricity costs, but after the point of 100% CPU utilization of their computer additional mining will require them to pay for both electricity and hardware. ASIC mining companies need to pay for electricity and hardware starting from the first hash. Hence, if the centralization gain can be kept to below this ratio, $(E+H)/E$, then even if ASICs are made there will still be room for ordinary miners.

Additionally, we intend to design the mining algorithm so that mining requires access to the entire blockchain, forcing miners to store the entire blockchain and at least be capable of verifying every transaction. This removes the need for centralized mining pools; although mining pools can still serve the legitimate role of evening out the randomness of reward distribution, this function can be served equally well by peer-to-peer pools with no central control. It additionally helps fight centralization, by increasing the number of full nodes in the network so that the network remains reasonably decentralized even if most ordinary users prefer light clients.

Scalability

One common concern about Ethereum is the issue of scalability. Like Bitcoin, Ethereum suffers from the flaw that every transaction needs to be processed by every node in the network. With Bitcoin, the size of the current blockchain rests at about 20 GB, growing by about 1 MB per hour. If the Bitcoin network were to process Visa’s 2000 transactions per second, it would grow by 1 MB per three seconds (1 GB per hour, 8 TB per year). Ethereum is likely to suffer a similar growth pattern, worsened by the fact that there will be many applications on top of the Ethereum blockchain instead of just a currency as is the case with Bitcoin, but ameliorated by the fact that Ethereum full nodes need to store just the state instead of the entire blockchain history.

The problem with such a large blockchain size is centralization risk. If the blockchain size increases to, say, 100 TB, then the likely scenario would be that only a very small number of large businesses would run full nodes, with all regular users using light SPV nodes. In such a situation, there arises the potential concern that the full nodes could band together and all agree to cheat in some profitable fashion (e.g., change the block reward, give themselves BTC). Light nodes would have no way of detecting this immediately. Of course, at least one honest full node would likely exist, and after a few hours information about the fraud would trickle out through channels like Reddit, but at that point it would be too late: it would be up to the ordinary users to organize an effort to blacklist the given blocks, a massive and likely infeasible coordination problem on a similar scale as that of pulling off a successful 51% attack. In the case of Bitcoin, this is currently a problem, but there exists a blockchain modification suggested by Peter Todd which will alleviate this issue.

In the near term, Ethereum will use two additional strategies to cope with this problem. First, because of the blockchain-based mining algorithms, at least every miner will be forced to be a full node, creating a lower bound on the number of full nodes. Second and more importantly,

however, we will include an intermediate state tree root in the blockchain after processing each transaction. Even if block validation is centralized, as long as one honest verifying node exists, the centralization problem can be circumvented via a verification protocol. If a miner publishes an invalid block, that block must either be badly formatted, or the state $S[n]$ is incorrect. Since $S[0]$ is known to be correct, there must be some first state $S[i]$ that is incorrect where $S[i-1]$ is correct. The verifying node would provide the index i , along with a “proof of invalidity” consisting of the subset of Patricia tree nodes needing to process $\text{APPLY}(S[i-1], \text{TX}[i]) \rightarrow S[i]$. Nodes would be able to use those nodes to run that part of the computation, and see that the $S[i]$ generated does not match the $S[i]$ provided.

Another, more sophisticated, attack would involve the malicious miners publishing incomplete blocks, so the full information does not even exist to determine whether or not blocks are valid. The solution to this is a challenge-response protocol: verification nodes issue “challenges” in the form of target transaction indices, and upon receiving a node a light node treats the block as untrusted until another node, whether the miner or another verifier, provides a subset of Patricia nodes as a proof of validity.

Putting It All Together: Decentralized Applications

The contract mechanism described above allows anyone to build what is essentially a command line application run on a virtual machine that is executed by consensus across the entire network, allowing it to modify a globally accessible state as its “hard drive”. However, for most people, the command line interface that is the transaction sending mechanism is not sufficiently user-friendly to make decentralization an attractive mainstream alternative. To this end, a complete “decentralized application” should consist of both low-level business-logic components, whether implemented entirely on Ethereum, using a combination of Ethereum and other systems (e.g., a P2P messaging layer, one of which is currently planned to be put into the Ethereum clients) or other systems entirely, and high-level graphical user interface components. The Ethereum client’s design is to serve as a web browser, but include support for a “eth” Javascript API object, which specialized web pages viewed in the client will be able to use to interact with the Ethereum blockchain. From the point of view of the “traditional” web, these web pages are entirely static content, since the blockchain and other decentralized protocols will serve as a complete replacement for the server for the purpose of handling user-initiated requests. Eventually, decentralized protocols, hopefully themselves in some fashion using Ethereum, may be used to store the web pages themselves.

CONCLUSION

The Ethereum protocol was originally conceived as an upgraded version of a cryptocurrency, providing advanced features such as on-blockchain escrow, withdrawal limits and financial contracts, gambling markets and the like via a highly generalized programming language. The Ethereum protocol would not “support” any of the applications directly, but the existence of a Turing-complete programming language means that arbitrary contracts can theoretically be created for any transaction type or application. What is more interesting about Ethereum, however, is that the Ethereum protocol moves far beyond just currency. Protocols and decentralized applications around decentralized file storage, decentralized computation and decentralized prediction markets, among dozens of other such concepts, have

the potential to substantially increase the efficiency of the computational industry, and provide a massive boost to other peer-to-peer protocols by adding for the first time an economic layer. Finally, there is also a substantial array of applications that have nothing to do with money at all.

The concept of an arbitrary state transition function as implemented by the Ethereum protocol provides for a platform with unique potential; rather than being a closed-ended, single-purpose protocol intended for a specific array of applications in data storage, gambling or finance, Ethereum is open-ended by design, and we believe that it is extremely well-suited to serving as a foundational layer for a very large number of both financial and non-financial protocols in the years to come.

A review of Ethereum and its vulnerabilities are discussed in Atzei, Nicola, Massimo Bartoletti, and Tiziana Cimoli. “A survey of attacks on ethereum smart contracts (sok).” In International conference on principles of security and trust, pp. 164-186. Springer, Berlin, Heidelberg, 2017.

Smart contracts are computer programs that can be correctly executed by a network of mutually distrusting nodes, without the need of an external trusted authority. Since smart contracts handle and transfer assets of considerable value, besides their correct execution it is also crucial that their implementation is secure against attacks which aim at stealing or tampering the assets. We study this problem in Ethereum, the most well-known and used framework for smart contracts so far. We analyse the security vulnerabilities of Ethereum smart contracts, providing a taxonomy of common programming pitfalls which may lead to vulnerabilities. We show a series of attacks which exploit these vulnerabilities, allowing an adversary to steal money or cause other damage.

Ethereum [23] is a decentralized virtual machine, which runs programs—called contracts—upon request of users. Contracts are written in a Turing-complete bytecode language, called EVM bytecode [47]. Roughly, a contract is a set of functions, each one defined by a sequence of bytecode instructions. A remarkable feature of contracts is that they can transfer ether (a cryptocurrency similar to Bitcoin [37]) to/from users and to other contracts.

Users send transactions to the Ethereum network in order to: (i) create new contracts; (ii) invoke functions of a contract; (iii) transfer ether to contracts or to other users. All the transactions are recorded on a public, append-only data structure, called blockchain. The sequence of transactions on the blockchain determines the state of each contract, and the balance of each user.

Since contracts have an economic value, it is crucial to guarantee that their execution is performed correctly. To this purpose, Ethereum does not rely on a trusted central authority: rather, each transaction is processed by a large network of mutually untrusted peers—called miners. Potential conflicts in the execution of contracts (due e.g., to failures or attacks) are resolved through a consensus protocol based on “proof-of-work” puzzles. Ideally, the execution of contracts is correct whenever the adversary does not control the majority of the computational power of the network.

The security of the consensus protocol relies on the assumption that honest miners are rational, i.e., that it is more convenient for a miner to follow the protocol than to try to attack it. To make this assumption hold, miners receive some economic incentives for performing the (time-consuming) computations required by the protocol. Part of these incentives is given by the execution fees paid by users upon each transaction. These fees bound the execution steps of a

transaction, so preventing from DoS attacks where users try to overwhelm the network with time-consuming computations.

Programming smart contracts. We illustrate contracts through a small example (A Wallet, in FIG. 1), which implements a personal wallet associated to an owner. Rather than programming it directly as EVM bytecode, we use Solidity, a Javascript-like programming language which compiles into EVM bytecode¹. Intuitively, the contract can receive ether from other users, and its owner can send (part of) that ether to other users via the function `pay`. The hashtable `outflow` records all the addresses² to which it sends money, and associates to each of them the total transferred amount. All the ether received is held by the contract. Its amount is automatically recorded in `balance`: this is a special variable, which cannot be altered by the programmer.

Contracts are composed by fields and functions. A user can invoke a function by sending a suitable transaction to the Ethereum nodes. The transaction must include the execution fee (for the miners), and may include a transfer of ether from the caller to the contract. Solidity also features exceptions, but with a peculiar behavior. When an exception is thrown, it cannot be caught: the execution stops, the fee is lost, and all the side effects—including transfers of ether—are reverted. Were the gas returned to callers in case of exceptions, an adversary could mount a DoS attack by repeatedly invoking a function which just throws an exception.

Each function invocation is ideally executed by all miners in the Ethereum network. Miners are incentivized to do such work by the execution fees paid by the users which invoke functions. Besides being used as incentives, execution fees also protect against denial-of-service attacks, where an adversary tries to slow down the network by requesting time-consuming Execution fees are defined in terms of gas and gas price, and their product represents the cost paid by the user to execute code. More specifically, the transaction which triggers the invocation specifies the gas limit up to which the user is willing to pay, and the price per unit of gas. Roughly, the higher is the price per unit, the higher is the chance that miners will choose to execute the transaction. Each EVM operation consumes a certain amount of gas [47], and the overall fee depends on the whole sequence of operations executed by miners.

Miners execute a transaction until its normal termination, unless an exception is thrown. If the transaction terminates successfully, the remaining gas is returned to the caller, otherwise all the gas allocated for the transaction is lost. If a computation consumes all the allocated gas, it terminates with an “out-of-gas” exception—hence the caller loses all the gas³. An adversary wishing to attempt a denial-of-service attack (e.g., by invoking a time-consuming function) should allocate a large amount of gas, and pay the corresponding ether. If the adversary chooses a gas price consistently with the market, miners will execute the transaction, but the attack will be too expensive; otherwise, if the price is too low, miners will not execute the transaction.

The mining process. Miners group the transactions sent by users into blocks, and try to append them to the blockchain in order to collect the associated fees. Only those blocks which satisfy a given set of conditions, which altogether are called validity, can be appended to the blockchain. In particular, one of these conditions requires to solve a moderately hard “proof-of-work” puzzle, which depends on the previous block and on the transactions in the new block. The difficulty of the puzzle is dynamically updated so that the average mining rate is 1 block every 12 s.

When a miner solves the puzzle and broadcasts a new valid block to the network, the other miners discard their attempts, update their local copy of the blockchain by appending the new block, and start “mining” on top of it. The miner who solves the puzzle is rewarded with the fees of the transactions in the new block (and also with some fresh ether).

It may happen that two (or more) miners solve the puzzle almost simultaneously. In this case, the blockchain forks in two (or more) branches, with the new blocks pointing to the same parent block. The consensus protocol prescribes miners to extend the longest branch. Hence, even though both branches can transiently continue to exist, eventually the fork will be resolved for the longest branch. Only the transactions therein will be part of the blockchain, while those in the shortest branch will be discarded. The reward mechanism, inspired to the GHOST protocol [43], assigns the full fees to the miners of the blocks in the longest branch, and a portion of the fees to those who mined the roots of the discarded branch⁴. E.g., assume that blocks A and B have the same parent, and that a miner appends a new block on top of A. The miner can donate part of its reward to the miner of the “uncle block” B, in order to increase the weight of its branch in the fork resolution process. Systems with low mining rate—like e.g., Bitcoin (1 block/10 min)—have a small probability of forks, hence typically they do not reward discarded blocks.

Compiling Solidity into EVM bytecode. Although contracts are rendered as sets of functions in Solidity, the EVM bytecode has no support for functions. Therefore, the Solidity compiler translates contracts so that their first part implements a function dispatching mechanism. More specifically, each function is uniquely identified by a signature, based on its name and type parameters. Upon function invocation, this signature is passed as input to the called contract: if it matches some function, the execution jumps to the corresponding code, otherwise it jumps to the fallback function. This is a special function with no name and no arguments, which can be arbitrarily programmed. The fallback function is executed also when the contract is passed an empty signature: this happens e.g., when sending ether to the contract.

Solidity features three different constructs to invoke a contract from another contract, which also allow to send ether. All these constructs are compiled using the same bytecode instruction. The result is that the same behavior can be implemented in several ways.

Ether lost in transfer. When sending ether, one has to specify the recipient address, which takes the form of a sequence of 160 bits. Addresses are sequences of 160 bits which uniquely identify contracts and users. However, many of these addresses are orphan, i.e., they are not associated to any user or contract. If some ether is sent to an orphan address, it is lost forever (note that there is no way to detect whether an address is orphan). Since lost ether cannot be recovered, programmers have to manually ensure the correctness of the recipient addresses.

Unpredictable state. The state of a contract is determined by the value of its fields and balance. In general, when a user sends a transaction to the network in order to invoke some contract, he cannot be sure that the transaction will be run in the same state the contract was at the time of sending that transaction. This may happen because, in the meanwhile, other transactions have changed the contract state. Even if the user was fast enough to be the first to send a transaction, it is not guaranteed that such transaction will be the first to be run. Indeed, when miners group transactions into blocks,

they are not required to preserve any order; they could also choose not to include some transactions.

Generating randomness. The execution of EVM bytecode is deterministic: in the absence of misbehavior, all miners executing a transaction will have the same results. Hence, to simulate non-deterministic choices, many contracts (e.g., lotteries, games, etc.) generate pseudo-random numbers, where the initialization seed is chosen uniquely for all miners.

Time constraints. A wide range of applications use time constraints in order to determine which actions are permitted (or mandatory) in the current state. Typically, time constraints are implemented by using block timestamps, which are agreed upon by all the miners. Contracts can retrieve the timestamp in which the block was mined; all the transactions within a block share the same timestamp.

1. Announcement of imminent hard fork for EIP150 gas cost changes. blog.ethereum.org/2016/10/13/announcement-imminent-hard-fork-eip150-gas-cost-changes/
2. Bitcointalk: Hi! My name is Rubixi. bitcointalk.org/index.php?topic=1400536.60
3. Ethereum Classic. ethereumclassic.github.io/
4. The ethereum network is currently undergoing a dos attack. blog.ethereum.org/2016/09/22/ethereum-network-currently-undergoing-dos-attack/
5. Ethereum reddit page. www.reddit.com/r/ethereum
6. Ethereum Wiki: Contract security techniques and tips. github.com/ethereum/wiki/wiki/Safety
7. Explaining eip 150. www.reddit.com/r/ethereum/comments/56f6we/explaining_eip_150/
8. GovernMental main page. governmental.github.io/GovernMental/
9. Hacking, Distribute: Scanning live Ethereum contracts for the “unchecked-send” bug. hackingdistributed.com/2016/06/16/scanning-live-ethereum-contracts-for-bugs/
10. King of the Ether Throne: Post mortem investigation. www.kingoftheether.com/postmortem.html
11. MAKER DART: a random number generating game for Ethereum. github.com/makerdao/maker-darts
12. RANDAO: a DAO working as RNG of Ethereum. github.com/randao/randao
13. Solidity: security considerations. solidity.readthedocs.io/en/develop/index.html
14. Understanding the DAO attack. www.coindesk.com/understanding-dao-hack-journalists/
15. Anderson, L., Holz, R., Ponomarev, A., Rimba, P., Weber, I.: New kids on the block: an analysis of modern blockchains. CoRR, abs/1606.06530 (2016)
16. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. In: IEEE S&P, pp. 443-458 (2014)
17. Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Gollamudi, A., Gonthier, G., Kobeissi, N., Rastogi, A., Sibut-Pinote, T., Swamy, N., Zanella-Beguelin, S.: Formal verification of smart contracts. In: PLAS (2016)
18. Bishop, M.: A taxonomy of Unix system and network vulnerabilities. Technical Report, CSE-95-10, Department of Computer Science, University of California at Davis (1995)
19. Bishop, M.: Vulnerabilities analysis. In: Proceedings of the Recent Advances in Intrusion Detection, pp. 125-136 (1999)
20. Boneh, D., Naor, M.: Timed commitments. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 236-254. Springer, Heidelberg (2000). doi: 10.1007/3-540-44598-6_15CrossRef

21. Bonneau, J., Clark, J., Goldfeder, S.: On Bitcoin as a public randomness source. IACR Cryptology ePrint Archive 2015, 1015 (2015)
22. Brown, R. G., Carlyle, J., Grigg, I., Hearn, M.: Corda: an introduction (2016). r3cev.com/s/corda-introductory-whitepaper-final.pdf
23. Buterin, V.: Ethereum: a next generation smart contract and decentralized application platform (2013). github.com/ethereum/wiki/wiki/White-Paper
24. Churyumov, A.: Byteball: a decentralized system for transfer of value (2016). byteball.org/Byteball.pdf
25. Clack, C. D., Bakshi, V. A., Braine, L.: Smart contract templates: foundations, design landscape and research directions. CoRR abs/1608.00771 (2016)
26. Delmolino, K., Arnett, M., Kosba, A. M. A., Shi, E.: Step by step towards creating a safe smart contract: lessons and insights from a crypto currency lab (2016)
27. Etherscripeter. etherscripeter.com
28. Eyal, I., Sirer, E.: Majority is not enough: bitcoin mining is vulnerable. In: Financial Cryptography and Data Security, pp. 436-454 (2014)
29. Frantz, C. K., Nowostawski, M.: From institutions to code: towards automated generation of smart contracts. In: Workshop on Engineering Collective Adaptive Systems (eCAS) (2016)
30. Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: analysis and applications. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 281-310. Springer, Heidelberg (2015). doi: 10.1007/978-3-662-46803-6_10
31. Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., Capkun, S.: On the security and performance of proof of work blockchains. In: ACM CCS, pp. 3-16 (2016)
32. Hirai, Y.: Formal verification of Deed contract in Ethereum name service. yoichihirai.com/deed.pdf
33. Landwehr, C. E., Bull, A. R., McDermott, J. P., Choi, W. S.: A taxonomy of computer program security flaws. ACM Comput. Surv. 26(3), 211-254 (1994) CrossRef
34. Luu, L., Chu, D. H., Olickel, H., Saxena, P., Hobor, A.: Making smart contracts smarter. In: ACM CCS (2016). eprint.iacr.org/2016/633
35. Luu, L., Teutsch, J., Kulkarni, R., Saxena, P.: Demystifying incentives in the consensus computer. In: ACM CCS, pp. 706-719 (2015)
36. Marino, B., Juels, A.: Setting standards for altering and undoing smart contracts. In: RuleML, pp. 151-166 (2016)
37. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008). bitcoin.org/bitcoin.pdf
38. Nipkow, T., Wenzel, M., Paulson, L. C. (eds.): Isabelle/HOL: A Proof Assistant for Higherorder. LNCS, vol. 2283. Springer, Heidelberg (2002)zbMATH
39. Nomura Research Institute: Survey on blockchain technologies and related services. www.meti.go.jp/english/press/2016/pdf/0531_01f.pdf
40. Pierrot, C., Wesolowski, B.: Malleability of the blockchain’s entropy. IACR Cryptology ePrint Archive 2016, 370 (2016)
41. Piessens, F.: A taxonomy of causes of software vulnerabilities in internet software. In: International Symposium on Software Reliability Engineering, pp. 47-52 (2002)
42. Popejoy, S.: The Pact smart contract language (2016). kadana.io/pact
43. Sompolinsky, Y., Zohar, A.: Secure high-rate transaction processing in bitcoin. In: Böhme, R., Okamoto, T. (eds.)

- FC 2015. LNCS, vol. 8975, pp. 507-527. Springer, Heidelberg (2015). doi: 10.1007/978-3-662-47854-7_32CrossRef
44. Swamy, N., Hritcu, C., Keller, C., Rastogi, A., Delignat-Lavaud, A., Forest, S., Bhargavan, K., Fournet, C., Strub, P., Kohlweiss, M., Zinzindohoue, J. K., Béguelin, S. Z.: Dependent types and multi-monadic effects in F*. In: POPL (2016)
45. Szabo, N.: Formalizing and securing relationships on public networks. *First Monday* 2(9) (1997). firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548
46. UK Government Chief Scientific Adviser: Distributed ledger technology: beyond block chain. www.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf
47. Wood, G.: Ethereum: a secure decentralised generalised transaction ledger (2014). gavwood.com/paper.pdf
48. Wüst, K., Gervais, A.: Ethereum Eclipse Attacks. Technical report, ETH-Zurich (2016)

SUMMARY OF THE INVENTION

The present technology provides various solutions addressing these issues. It converts an illiquid asset which represents demonstrable and recognized wealth, into a liquid resource, typically without requiring the illiquid asset itself to be monetized. A contingent interest in real property rights (e.g., a security interest) permits, upon conclusion of a smart contract, a counterparty to the asset owner, to force either compliance with the terms of the smart contract, e.g., repayment of the tokens, or have available legal remedies available under the contingent interest. In a preferred embodiment, the smart contract requires return of the same number of tokens as were issued to the asset owner, in order to release the contingent interest. This permits a distributed ledger to be a complete and accurate accounting means for the tokens, from creation to extinction by completion of the cycle. Of course, more complex terms are possible, such as a fractional or surplus return ratio. In a preferred embodiment, a substitution is possible, which serves to secure release of the contingent interest, absent return of all tokens, for example in case of unavailability or pricing disparity.

In general, a borrower seeks to monetize a capital investment in a productive resource that has a predictable cost of production of a commodity, and other predictable (and/or insurable) risks. A buyer, lender, or issuer seeks to securitize the investment by creating a salable and tradable token, guaranteed by the value of the security (security interest), and subject to a "smart contract" which limits unpredictable human-factor risks, so that the value of the loan can be immediately off-loaded into public markets. Finally, public markets acquire the tokens, which have advantages over cash, in terms of transferability, security, and in some cases (i.e., unstable political regimes), asset backing. The tokens are ultimately backed by a security interest in the productive means, and may have a discount in excess of the cost of extraction or production, at least at issuance. The tokens, once issued, have a fundamental value, tied to the security interest and smart contract that governs it. Therefore, in contrast to a fiat currency, the tokens are less correlated in value with political instability, inflation, devaluation, foreign trade balance, and the like. On the other hand, the token would typically have a value positively correlated with the value of the underlying security asset, and the commodity to which it pertains.

The tokens may be tied to a single smart contract, or issued backed by a portfolio of secured assets. A typical

smart contract will have a predefined term, such as 10 years. At the end of the term, the borrower must acquire the tokens at market price, or provide a permitted alternate. In a preferred embodiment, the smart contract authorizes a right of substitution, wherein the borrower may release the security interest by depositing with a reliable repository, a predetermined amount of the commodity, and at termination, the holder(s) of tokens can offer their tokens to the original source, presumably at or near the underlying commodity value, as of the expiration of the contract. In another scenario, the token holders receive the commodity itself, and not its value, though in most cases, the commodity is represented by a liquid right in a defined currency. The fixed expiration date permits deposit of commodities futures contracts as the substituted assets, rather than the commodities themselves, which further reduces the economic inefficiency of the securitization.

The preferred implementation permits a renewal of the smart contract, such that an actual transfer of the commodity is not required. Thus, the smart contract can be renewed, with underwriting performed in anticipation of renewal to verify that the secured assets have appropriate value, etc.

Users of the tokens may therefore anticipate variations in demand over time. However, while the security for the particular tokens may differ, the smart contract is designed and intended to provide sufficient margin between the security interest and likely range of values such that they are fungible, and various tokens for the same commodity would assume the same market value.

In some cases, the tokens will rise to a value above a right of substitution value or security interest value. In this case, there may be differences between tokens, but in that case, there would be an incentive for bidders to substitute the security as the commodity or a contract right for the commodity, and thus achieve fungibility in that manner. When the value of the token rises above the securitization, the result is somewhat similar to the pricing of non-securitized cryptocurrencies, i.e., the pricing is dependent on scarcity of the tokens and a demand established by usefulness or speculation. However, in contrast to current unsecuritized cryptocurrencies, the securitized tokens have a right, at smart contract termination, to liquidation at the exchange price for the commodity. Further, if the tokens are fully substituted (a likely occurrence if the token value exceeds the commodity exchange value), the smart contract may go into automatic renewal (with a right of token redemption), and thus expiration risk for token-holders abated. Therefore, the tokens have a minimum value expressed in a commodity exchange rate at or after smart contract expiration, and no maximum value.

The scarcity of the tokens is guaranteed by the limited nature of commodity production or creation facilities, and the securitization discount.

To the extent that commodities market risks are deemed unacceptable or undesired, the tokens may be backed by various types of portfolios that mitigate desired risks. Likewise, risks may be concentrated if desired. For example, instead of securitizing tokens based on a mineral mine or the like, the borrower may simply deposit futures contracts in exchange for tokens, which are then sold to raise cash. If the borrower owns a productive resource for the subject commodity, the future contract may be written based on production and fulfillment by the borrower. However, this is not required, and the borrower may engage in side transactions to fulfill its eventual obligations, so long as the security interest meets the smart contract and underwriting criteria.

Blockchain technology provides various known advantages. For example, entries made on the blockchain are permanent, immutable, and independently verifiable. Therefore, the use of blockchains is particularly valuable for verifying ownership of a token, recording transfers of tokens, and auditing transactions. The present technology does not require, and a preferred embodiment does not include, anonymity, and therefore parties to a transaction may be authenticated using biometrics, multifactor authentication, or various means. This avoids the need to rely on passwords or cryptographic credentials alone, as with some cryptocurrencies. Thus, a human service may be employed to verify participants in transactions. On the other hand, anonymous transaction capability is also possible.

The use of tokens provides advantages with respect to currencies, commodity trades, barter, contract obligations, and other payment methods. Currencies can require conversions, which incur costly and risk. Commodity trades may require delivery of significant quantities of material, storage (with required physical security), transactional costs, pricing fluctuations, etc. Barter has similar issues with commodities, but also incurs liquidity risk. (If an exchange is with respect to a liquid asset, it is either a security or a commodity, if it is neither, it would be considered a barter). Contract rights, such as futures, incur the risks of the underlying security or commodity, but also a greater pricing fluctuation risk to contract termination, and likely higher transactional costs.

The token according to the present technology has some characteristics of a derivative of a forward contract, with the advantage that, while the value of the token is secured by the value of the resource capable of delivering the commodity (with certain investment in delivering the commodity less than a differential between the value borrowed and the present value of the resource), and there is no actual requirement to develop the resource to deliver the commodity, maintaining efficient options for the borrower. The derivative, however, has characteristics of a virtual cryptocurrency, with asset backing, which can reduce volatility. While there is technically no limit on demand-based increases in the value of the tokens, and thus opportunity for speculation, according to a preferred embodiment, the borrower has a right of substitution, and therefore the risk of specific demand by a borrower seeking to recover the tokens and retire the debt, and therefore an opportunity for a hold-out, is limited. Further, the technology does not particularly seek to limit the amount of asset-backed tokens generated, and therefore scarcity of the class of tokens is not a driving principle for valuation. Therefore, the motivations for pricing volatility are limited, resulting in a cryptocurrency or cryptoasset token whose basis for valuation is the asset backing and usefulness in or in conjunction with commerce.

Typically, the tokens are fully fungible, and are issued in undistinguished form and as part of an indistinguishable pool. With respect to particular issuances, it is possible that a certain resource will turn out to be over-valued or fraudulent, and as a result, the tokens issued based on that resource lack asset backing. Since we presume that there are other tokens available in the same class that are properly asset backed, this would tend to negatively impact the valuation of those particular tokens. Of course, steps are taken during underwriting to avoid this possibility, but it is also possible to acquire insurance on the assets to assure the value of their securitization. When tokens are separately tracked, as compared to a normal token, a token under which the borrower has defaulted, and a surety is invoked may be worth more than other tokens, leading to higher "special" demand for riskier tokens. On the other hand, if the token value is above

redemption, the special tokens will be worth less. Therefore, a speculation opportunity is provided in this case.

If, upon termination of a smart contract, there is a default, and the security interest must be liquidated, it is possible that an investment is required to extract the resource to fulfill the contract. In that case, an administrator, trustee, receiver, or the like (as specified by the smart contract and underlying agreements) assumes control over the resource. Because of the margin between the debt amount and the value of the commodity that may be extracted from or produced by the resource, it is likely that the estate can be managed to fulfill its responsibility. Indeed, there may be cases where the token-holders are advantaged (at least theoretically) as compared to those holding tokens from non-defaulting debtors. Thus, a perverse incentive may occur where demand is high for the lowest "quality" tokens, due to an arbitrage opportunity at contract termination. For example, in a non-defaulting case, the debtor must acquire the previously issued tokens or deliver the substituted assets. Therefore, the market price for the tokens will be driven by the substitution paradigm, which is generally the same commodity as produced by the resource (though this is not a theoretical constraint). On the other hand, in the defaulting case, the debtor presumably lacks the liquid assets to re-acquire the tokens at market price, and lacks the commodity to fulfill the substitution, leading to a possibility that the rational market price is above the substitution value, especially if the terms of default provide an advantage to the token-holder, such as an above-market interest rate for delayed payment. Further, because these defaulting tokens become "special", and the smart contract under which they are issued will not be renewed, the expiring tokens may be subject to increased demand, and thus higher prices.

This is not to say that the system is design to include perverse incentives which drive objectively antisocial behavior and seeming irrational results. Rather, the terms of the smart contract may be designed to correct for aberrant conditions, and to ensure that token holders are assured of the asset backing under all of the various conditions, leading to a lower risk discount and reduced correlation of risks between the tokens themselves and external market-specific conditions. Thus, the tokens become available for cross-border transactions, use in unstable economies, during fiat currency contractions, and especially in transactions loosely linked to the commodity which secures the tokens. While the tokens may be initially issued in a private transaction, such as a financing of an acquisition of a mine, or expansion of a mine, they would become available for general transactions as an alternate to cryptocurrencies, and due to their characteristics designed to limit speculation-driven volatility, and artificial scarcity (i.e., scarcity due only to the fact that the tokens have limited liquidity and availability), the tokens become a good fit for use in consumer and business-to-business transactions.

The key advantage of the tokens over cash derives from their origination in an asset securitization transaction, which is designed to be more efficient, and yield lower discounts of net asset value to liquidity achieved, than traditional bank lending, and perhaps equity issuance. Thus, the generation of the tokens according to the present technology alleviates an artificial scarcity of fiat currencies, under traditional loan paradigms, thus unlocking vast amounts of wealth. Once unlocked, the tokens continue to represent an advantage over fiat currency in that they are readily exchangeable for fiat currency on an exchange or in a transaction, and are also backed by hard assets, a feature missing from fiat currencies.

Typically, the present technology is implemented in a manner fully compliant with all banking and securities laws, and the use or advantages of the technology are not predicated on US tax avoidance, US currency transfer restriction evasion, etc. However, the exchange network is independent of the financial services oligopolies, such as The Clearing House (New York), credit card networks (Mastercard, Visa, Discover, American Express, etc.), etc., and therefore is free to compete by providing lower cost, better service(s), advantages, etc.

According to another embodiment, the token system is tied to commodity-specific investment or speculation. For example, if a user seeks to invest in diamonds, gold, platinum, etc., he or she may acquire tokens having such characteristics. Typically, these would be tokens that have a high proportion of substituted assets in the form of the respective commodity, whose pricing is highly correlated with the commodity, and whose smart contract assures that significant deviation between the token price and the redemption price with respect to the underlying commodity are well aligned. Indeed, in such a scenario, short or staggered terms of the smart contract may be desired, so that speculation on commodity pricing over time can be effectively managed. Further, in some cases, the token owner may be provided with a right of demand, for example, to acquire an amount of the underlying commodity in exchange for the token. This would have the effect of converging the token price with the demand value as the contract nears expiration, and also permits use of the tokens as a market hedge, while maintaining liquidity.

Another issue involves international currency and financing issues. For example, a non-US commodity miner seeks to finance production of gold outside the US. Typically, a security interest in real estate or a business is recorded in the jurisdiction in which the business is location, and is denominated in the currency of that jurisdiction. This imposes difficulties where a lender does not wish to incur currency risk for the particular jurisdiction, but is quite willing to incur the business risk inherent in the loan, if denominated in commodity value. Therefore, the present technology provides a means to secure the loan, which in theory does not violate currency export restrictions of the jurisdiction, since at contract termination, the commodity itself may be delivered. The loan, in the form of the tokens, then represents a derivative of the value of the commodity, at least when issued, and near expiration, especially in a non-renewable smart contract. For a renewable smart contract, the value of the tokens is largely limited unilaterally, in that if the tokens exceed the market value of the underlying commodity, there will be a large incentive for the borrower to substitute commodity for tokens, thus tying their values together. However, if the commodity value is less than the token value, the token will remain floating in value.

Because of the discount between the value of the secured assets and the amount of the security, the tokens have a fair market value over equity in the resource. Therefore, the tokens have another basis for valuation, and another possible value correlation. In the event that the company that manages the resource becomes illiquid and its business prospects dim, the tokens assume properties similar to that of a secured lender. The redemption at contract expiration, in this case, appears as a zero coupon bond. The option of renewal of the contract, if provided, may provide the token-holder with an ability to cash out, but in any case typically requires a new underwriting process that assures that the value of the security interest is discounted well below the amount outstanding.

Therefore, the particular rules and features of the smart contract will define future risks and rewards of the participants, and can vary over a broad range of parameters. For example, instead of emulating a zero coupon bond, the tokens may yield dividends or pay interest. If this is a direct payment, this incurs tracking and tax reporting that might result in difficulties. However, as built into a smart contract, the result may be a programmed increase in value of the tokens over time with respect to redemption value, right of substitution amount, or other time-dependent features which are predicted to drive the value of the tokens up over time toward the expiration.

The present technology therefore features an asset which is subject to a legally-enforceable security interest, to secure a debt, which may be denominated in currency or a value of a commodity. In exchange for a sum, a series of tokens are generated and issued. The tokens are subject to a contract which provides for a future redemption, at a value secured by the security interest. The contract itself may be a smart contract, which includes automatically implemented rules and features, which in some cases may be independent of sovereigns, and has aspects which are enforced independent of the parties and their agents. The tokens may be traded on an exchange which relies on a blockchain. The future redemption may provide different options, such as currency, commodities, renewal upon terms, or otherwise.

In a preferred embodiment, the security interest is in a gold mine, with a debt of less than 20% of the proven reserves of the claim, after a due diligence investigation of the value of the claim and its productive capability, such that extraction of gold sufficient to redeem the tokens is commercially feasible without exceeding the value of the secured assets. The term may be 10 years, with a right of substitution at par value at any time over the term. At expiration, the borrower must reclaim all of the tokens, or substitute security in the form of a pre-specified amount of gold. However, upon redemption of any token, the borrower may reissue the debt as a new token, subject to a new smart contract. A term of the smart contract permits a token-holder with an outstanding token to automatically exchange tokens for replacement tokens, thus leaving pending transactions and markets uninterrupted.

The issued tokens are available for various transactions, similar to other known cryptocurrencies, and are traded on exchanges with respect to different currencies, tokens, or commodities, or between individuals on a secure digital ledger, which may be a blockchain.

The right of redemption may permit substitution of other tokens, having distinct security, which meet system-wide criteria, and therefore the tokens become fungible. However, each debtor remains responsible at contract termination to fulfill its obligations, or the security interest may be foreclosed. When the process is aggregated for a number of borrowers, the management of redemption, substitution, and foreclosures becomes a normal business activity, and therefore can be managed accordingly. Further, the aggregation leads to greater liquidity, reduced search cost and individual risk speculation, and more orderly markets. Further, from the perspective of a token-holder, default risks are also aggregated, and therefore have lower volatility.

In some cases, the value of resource subject to the security interest is far in excess of the value of the debt, and respective tokens. One option is, rather than issuing all tokens having the same value, a set of tranches of tokens, which represent priority of liquidation in event of default. This has the effect of creating tokens with greater and lesser degrees of correlation with the underlying business involved

in the secured resource. The tranches may also have other differences. Thus, it is possible to create derivatives which largely isolate different risks, and thus satisfy different investment objectives.

In some cases, the security is not in real estate, mines or leases per se, but may be with respect to in-process inventory. For example, mined, unrefined ore may be subject to the monetization transaction, with the redemption based on a fraction of the metal or mineral contained in the ore. Therefore, the technology may be used for short term business financing. As the ore is processed, if the tokens are not redeemed, the miner may replace the ore to maintain the pool or secured assets.

As discussed above, the smart contract may be insured. This permits an insurer, such as an insurance company, to assume various risks independent of the token values, and may therefore arbitrage the market value of the tokens with respect to the implied insurance risk. This therefore incentivizes the insurer to act as a market-maker with respect to the tokens under various circumstances, assuring liquidity and orderly markets for the tokens. Even if the smart contract does not require a captive sinking fund, an insurer may require this as a term of the contract. The use of insurance is one way to make different tokens fungible; if tokens having various underlying distinctions are insured to have the same risks (e.g., a common guaranty by the same insurance company), then the markets may treat these as being of the same type. This can be reinforced if the right of substitution includes alternate tokens (though this will generally sink the market value to the lowest valued token, which may be undesirable from an efficiency standpoint).

The blockchain employed is preferably the Symbiont system, which is a permission-based blockchain. See, symbiont.io/technology/.

By applying the power of a distributed ledger and smart contracts, the present invention can offer commodity asset owners a method to attain liquidity from their pre-commodity assets by digitizing those assets, and providing a new liquid asset representing a liquid derivative of the pre-commodity assets. These new liquid assets are derived from a fractional representation of the commodity assets. Commodity asset owners who desire liquidity can use this fractional representation (the new liquid asset) for alternative financing. This method will considerably speed up the process of gaining liquidity while allowing the asset owner to avoid ongoing financing charges.

The digitization of commodity assets allows for the entry of previously excluded asset classes into the existing securitization marketplace. The problem was previously addressed through a traditional and less efficient securitization market. In the current marketplace, a pre-commodity asset owner can go to a lender and securitize the pre-commodity assets thereby gaining liquidity. The problem with the current model is that it is painfully long and expensive due to the deep discounting of the value of the pre-commodity assets as well as ongoing interest charges due to the business model to the lender.

The previously available solution was not fully satisfactory because it required ongoing finance charges as well as steeper discounting of the underlying assets. The known solutions require that there be a one to one relationship between the commodity asset holder and the lender. Because of this, the commodity asset holder is typically restricted in how to use the borrowed funds. The funds received were simply cash in the bank and not further available for

securitization or other advanced financing techniques including the ability to sell to and/or repurchase from third parties.

These, and other objects and advantages of the present invention are achieved in accordance with the present invention by the method according to the invention. In one preferred embodiment, the method is for creating precious metal backed token assets, secured by unrefined precious metal reserves, utilizing blockchain and/or distributed ledger technology. It is of course understood that other unrefined or pre-commodity reserves can be digitized in accordance with the present invention. Typically, if the assets achieve commodity status, there are existing markets that can be employed; however, such assets are not excluded by the present technology.

Similarly, non-commodity assets may also be securitized, with a slightly different security predicate. For example, an income-producing real estate investment trust can be monetized according to the present technology. In terms of substitution or replacement, the options include, for example, real estate interest substitution, heterogeneous asset or commodity type substitution, cash substitution, or other security.

In accordance with one preferred embodiment of the present invention, the digitization of illiquid pre-commodity assets creates a liquid asset, utilizing smart contracts and distributed ledger technology. The commodity asset holder is required to pledge the assets into a collateral pool for digitization and representation on the distributed ledger. Once recorded on the distributed ledger as inventory, the collateral will be used to digitize a fractional representation of the commodity assets.

Digitization occurs on the ledger via a smart contract. The smart contract contains all the necessary parameters needed to digitize including (but not limited to) the commodity asset description, the owner, the quantity, the location and the appropriate risk adjusted discount for the commodity asset and the title for the duration of the smart contract. The terms of digitization embedded in the smart contract allow for the allocation of the fractional representation of the commodity asset to the commodity asset owner. Also, embedded in the terms of the smart contract, is a maturity date which triggers the release of any pledged collateral back to the commodity asset owner while simultaneously requiring the return of the original fractional representation to the collateral pool. The commodity asset owner has control of the now liquid fractional representation created from the collateral. Of course, other encumbrances or side-deals may take place, represented in smart contracts or otherwise.

It is therefore an object to provide a token, representing an interest in a smart contract, the smart contract representing an agreement, secured by a security interest in the real property or a right in real property, to return the token within a predetermined period.

It is a further object to provide a method of defining a token, comprising: defining a smart contract, representing an agreement, secured by a security interest real property or a right in real property, to return the token within a predetermined period; pledging the real property or a right in real property to secure the security interest in the smart contract; and issuing the token. The method may further comprise returning the token, and releasing the real property or a right in real property from the security interest. The smart contract may be implemented in conjunction with a distributed ledger. The method may further comprise exercising a contingent property interest, e.g., the security interest, in the real

property or a right in real property, after the predetermined period if the token is not returned.

The token may represent a fractional interest in the real property or a right in real property.

The smart contract may be implemented in conjunction with a distributed ledger.

The predetermined period may be tolled if a substitute asset is tendered.

The real property or a right in real property may comprise a mine having proven available reserves of the substitute asset. The substitute asset may be gold. The proven available reserves may be a predetermined multiple of the substitute asset.

The token may be generated as a transaction of a cryptographically-authenticated, distributed ledger comprising a database held and updated independently by each of a plurality of distributed elements, forming a consensus determination of transaction validity.

The agreement, secured by a security interest in the real property or the right in real property, may be terminated if the token is returned within the predetermined period, else a contingent property interest, e.g., the security interest, in the real property or a right in real property, may be exercised.

It is a further object to provide a method for creating a liquid token representation from an illiquid asset comprising: receiving a pledge of an illiquid asset; and digitizing the illiquid asset into fractional representations using a smart contract on a distributed ledger network, the fractional representations being secured by the pledge of the illiquid asset as collateral. The method may further comprise trading a fractional representation on an exchange, recorded in a distributed ledger network. The smart contract may comprise an illiquid asset description, an illiquid asset owner, a quantity, and at least one redemption rule. The at least one redemption rule may comprise a maturity date which triggers a release of the pledged illiquid asset as collateral back to the illiquid asset owner in exchange for return of all of the original fractional representations. The illiquid asset may comprise ore of a precious metal. The precious metal may be gold.

It is a still further object to provide a distributed ledger comprising terms of at least one smart contract representing an agreement which imposes a security interest in real property rights, comprising a term which authorizes creation of a token subject to the security interest, and subsequently deactivates the token and releases the security interest upon fulfillment of the smart contract terms.

It is another object to provide a computational node of a distributed communication network, configured to execute a portion of a distributed ledger comprising at least one smart contract representing an agreement which imposes a security interest in real property rights, comprising a contract term which authorizes creation of a token subject to the security interest, and subsequently deactivates the token and releases the security interest upon fulfillment of the smart contract terms.

The smart contract may comprise at least one term which imposes a predetermined period during which the terms must be fulfilled. The smart contract may comprise at least one term which provides a right to substitution to toll a foreclosure of the security interest. The distributed ledger may be provided in combination with computational nodes of a distributed communication network, configured to authenticate transactions involving the token, and automatically execute the terms of the smart contract, without centralized control.

These and other features of the present invention will become apparent from the drawings and the following detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of the distributed ledger (blockchain) Network used by a method according to the present invention;

FIG. 2 is a flowchart of a first embodiment of the method according to the present invention;

FIG. 3 shows a flow diagram of underlying asset verification and valuation;

FIG. 4 is a flowchart of a second embodiment of the method according to the present invention;

FIG. 5 is a flowchart of a proposed information flow according to the second embodiment of the invention; and

FIG. 6 is a state diagram of a sovereign-backed securitization model according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Example 1

FIG. 1 shows a block diagram of the network utilized in accordance with the method of the present invention. At the heart of the system is a distributed ledger (e.g., transaction chain or blockchain) network **10** preferably implemented on the Internet and providing a distributed ledger, which is immutable. The network may use public key/private key cryptography to insure identification integrity and other algorithms to insure trust before a block of at least one transaction is added to the distributed ledger. The network can be implemented on any platforms that permit the running of smart contracts, such as the Hyperledger blockchain, Symbiont.io or the Ethereum blockchain.

Connected to the network **10** is a commodity asset owner **20** who is interested in pledging illiquid assets, such as unrefined gold that is still in the ground, to **3** create a liquid asset. Also connected to the network **10** is a digitizer party **30** who agrees to take the pledge of the illiquid assets subject to terms enforced in smart contracts running on the network **10**, and digitize the asset into fractional representation that can be sold to account holders **40**. For example, if the commodity asset was unrefined gold for which the owner can demonstrate that there is a proven gold reserve, the digitizer will provide 1 fractional representation (e.g., an Orbit.au) token for a defined amount of Reserves of gold. The assets will be pledged for 10 years, after which the asset owner must replace the entire reserve that was fractionalized.

The token can be sold to account holders on the network and each transaction is recorded immutably in a block on the distributed ledger to establish unquestionable ownership rights. The smart contracts, which are computer programs designed to operate on the distributed ledger network and carry out the terms of the method, automate the process and eliminate the need for human intervention in many steps.

FIG. 2 is a flowchart of one embodiment of the method according to the invention. According to the method of FIG. 2:

A reserve is created by the asset holder. A first Smart Contract creates the inventory of the total amount of reserves being placed in the pool by the asset holder for digitization.

The diamond indicates that an executive must sign off on the reserve once it is created before it can go to the ‘Signed Off’ state. If there is no sign off, the system waits for the proper approval.

If the reserve has been approved (signed off), it is now ready to be digitized by the digitizer party using a smart contract process.

In order for the digitization to occur, the digitizer runs a second smart contract which is called an ‘ADSA’ (Asset Digitization Service Agreement) which is shown in a flow diagram in the second column. This contract knows the haircut (discount) and also is where the digitizer party puts in the maturity date, the digitization date and when digitized, creates the tokens. This is also where the digitizer associates the ADSA to a reserve and the number of tokens created for a particular reserve can be seen.

The created ADSA now waits for executive sign off similar to the reserve to create an object.

Once the digitization date is reached the reserve is bound, the ADSA is marked digitized and the tokens are created. FIG. 2 illustrates that until the digitization date comes and digitization occurs, the asset holder can still back out. However, if the digitization occurs, there is no backing out.

The ADSA now distributes the tokens created to the asset holder, and also deducts relevant on-ledger fees which are paid to both the network operator and the digitizer party.

The 10-year period now begins and after the maturity date, everything underneath happens at termination, i.e., the debt is repaid, the tokens are destroyed (e.g., redeemed and/or retired), the ADSA is terminated and the Reserve is terminated as long as there are no associated ADSA’s.

Specifically, on termination the ADSA will look at the asset holder’s account and sweep the tokens (exactly the number that was digitized) back for inactivation so that the lien can be lifted off the reserve. The inactivation, redemption, or retirement of a token is a transaction on the distributed ledger that labels the tokens with an updated status, to alert future buyers of that status, which would normally render them valueless, and thus block future transactions. (Note that, according to existing non-asset backed cryptocurrencies, the lack of asset backing does not preclude use, so technically, the transition from asset-backed to non-asset backed does not require that no party attribute value to the redeemed tokens.)

If for some reason the number of tokens in the asset holder’s account are less than the original amount created the process goes into a default scenario. In event of a default, the legal process of foreclosure on the secured assets proceeds, and this provides security for the token-holders. Since the amount of feasible asset recovery exceeds the redemption value of the tokens, it is most probable that all token holders will be made whole, and indeed, the default process may make outstanding tokens more valuable than those that are redeemed in the normal course.

As a result of the method of the present invention, fungible liquidity is obtainable from commodity assets in various states of extraction or non-extraction and refinement. The method can create liquidity from pre-refined, combined and disparate commodity assets for each of those disparate commodity assets.

Example 2

System Architecture

It is a challenge for typical investors to get exposure to unrefined assets while, at the same time, owners of such assets often struggle to access liquidity. Asset digitization

can provide investors with exposure to illiquid assets in a form that can be easily registered, traded and transferred. It also provides owners of illiquid physical assets with an opportunity to access new sources of liquidity. A distributed ledger technology platform is ideal for asset digitization because it provides an immutable record of the origination and provenance of digitized assets as well as a tamper-proof repository for all documentation supporting a given origination.

Tokens created according to the present paradigm are fundamentally different from most blockchain or distributed ledger offerings. Two core differences are that they are a cryptoasset, and not a cryptocurrency loosely backed by a hard asset, and directly represent the hard asset which is available as security for the set of transactions. The tokens are not decoupled from central management or rules of law. The platform is de-centralized from a resiliency and technological perspective but it is centrally managed by a service company. Unlike cryptocurrencies, such as Bitcoin or Ethereum, which are not centrally governed, it does not make sense to decouple a hard asset, which is subject to rules and regulations nor would you want to. Therefore, the divergence on this issue is both fundamental and philosophical. Since the linkage to a hard asset affords contingent property rights, the ability to operate with and within law is important. The same rules and regulations that protect the hard asset, i.e., preventing someone from usurping ownership of mineral rights, will protect any investment in the tokens, or the underlying smart contract. The tokens are subject to rule of law, and can and will be transferred in accordance with court rulings.

The hard asset backing the tokens may be, for example, proven reserves of unrefined gold. Primarily this will be in the form of in situ gold deposits but could include gold assets in interim stages of the refining process including head ore, concentrate, miner bars, and doré. The tokens represent a single mineral or commodity, and tokens tied to different security types will generally not be fungible among asset types. However, it is possible to define diversified pools, which consistently represent a plurality of asset types in predetermined ratios. Each token is fungible so a token from one source is equivalent and interchangeable to a corresponding token from another source without the need to trace its provenance. However, the provenance of each and every token is traceable within the pool and documented. The history of every transaction is stored in an immutable and tamper-proof ledger along with all supporting documentation.

Previously there was no platform for lending against unrefined proven reserves of gold. With this model, a secure auditable platform enables these transactions.

To validate the mineral deposits of any reserve or claim, one can use the “Canadian Institute of Mining” (CIM), National Instrument 43-101 (NI 43-101) guidelines for reporting. The NI 43-101, although Canadian based, is commonly used throughout North America, and even internationally. Although it is common to North America, the CIM/NI 43-101 reporting guidelines adhere to the “Committee for Mineral Reserves International Reporting Standards” (CRIRSCO). In the future, other CRIRSCO member guidelines may be used, including but not limited to the Australian “Joint Ore Reserves Committee” (JORC) reports. Each NI 43-101 report is conducted by an independent “qualified person(s)”. At a minimum, the qualified person must be an experienced and accredited engineer or geoscientist with experience relevant to the subject matter of the

mineral project (NI 43-101, 2011, p. 4). This report may be commissioned by the claim or reserve holder, but it is not carried out by the holder.

In order to understand the NI 43-101 report, the following definitions are required:

A Mineral Resource is a concentration or occurrence of solid material of economic interest in or on the Earth's crust in such form, grade or quality and quantity that there are reasonable prospects for eventual economic extraction.

The location, quantity, grade or quality, continuity and other geological characteristics of a Mineral Resource are known, estimated or interpreted from specific geological evidence and knowledge, including sampling.

Mineral Resources are sub-divided, in order of increasing geological confidence, into Inferred, Indicated and Measured categories. An Inferred Mineral Resource has a lower level of confidence than that applied to an Indicated Mineral Resource. An Indicated Mineral Resource has a higher level of confidence than an Inferred Mineral Resource but has a lower level of confidence than a Measured Mineral Resource.

The full report goes into greater detail on each Mineral Resource classification, however we are mainly concerned with minerals with the highest degree of confidence and those which can be converted into proven reserves. The following is a definition of a measured mineral resource:

A Measured Mineral Resource is that part of a Mineral Resource for which quantity, grade or quality, densities, shape, and physical characteristics are estimated with confidence sufficient to allow the application of Modifying Factors to support detailed mine planning and final evaluation of the economic viability of the deposit. Geological evidence is derived from detailed and reliable exploration, sampling and testing and is sufficient to confirm geological and grade or quality continuity between points of observation. A Measured Mineral Resource has a higher level of confidence than that applying to either an Indicated Mineral Resource or an Inferred Mineral Resource. It may be converted to a Proven Mineral Reserve or to a Probable Mineral Reserve. (CIM Definitions, 2014)

Measured Mineral Resources offer the highest level of confidence such that there is sufficient sampling and testing to confirm grade or quality between points of observation. However, to be accepted by the present system, the analysis must also ensure the economic feasibility of the deposit. The presence of gold is insufficient if it is in such low concentration that extraction is not economical, or if there are legal or environmental restrictions. The NI 43-101 also takes this into account as well as defined by "Mineral Reserves":

A Mineral Reserve is the economically mineable part of a Measured and/or Indicated Mineral Resource. It includes diluting materials and allowances for losses, which may occur when the material is mined or extracted and is defined by studies at Pre-Feasibility or Feasibility level as appropriate that include application of Modifying Factors. Such studies demonstrate that, at the time of reporting, extraction could reasonably be justified. The reference point at which Mineral Reserves are defined, usually the point where the ore is delivered to the processing plant, must be stated. It is important that, in all situations where the reference point is different, such as for a saleable product, a clarifying statement is included to ensure that the reader is fully informed as to what is being reported. The public disclosure of a Mineral Reserve must be demonstrated by a Pre-Feasibility Study or Feasibility Study. (CIM Definitions, 2014)

It is important to understand that the Modifying Factors are not limited in scope to just the economics of the extraction process:

Modifying Factors are considerations used to convert Mineral Resources to Mineral Reserves. These include, but are not restricted to, mining, processing, metallurgical, infrastructure, economic, marketing, legal, environmental, social and governmental factors. (CIM Definitions, 2014)

Once Modifying Factors are applied to the technical reports on the Mineral Resource Estimates a new category of Mineral Reserves are generated including Proven Mineral Reserves:

A Proven Mineral Reserve is the economically mineable part of a Measured Mineral Resource. A Proven Mineral Reserve implies a high degree of confidence in the Modifying Factors. Application of the Proven Mineral Reserve category implies that the Qualified Person has the highest degree of confidence in the estimate with the consequent expectation in the minds of the readers of the report. The term should be restricted to that part of the deposit where production planning is taking place and for which any variation in the estimate would not significantly affect the potential economic viability of the deposit. Proven Mineral Reserve estimates must be demonstrated to be economic, at the time of reporting, by at least a Pre-Feasibility Study. Within the CIM Definition standards the term Proven Mineral Reserve is an equivalent term to a Proven Mineral Reserve. (CIM Definitions, 2014)

FIG. 3 illustrates the relationship between the confidence of testing and samples versus Modifying factors.

Once a claim with an NI 43-101 has been reviewed and the Proven Reserves of gold have been validated, the value of the collateral, for all intents and purposes, is pegged at 1:5. For every 5 troy ounces of Proven Reserves of gold, exactly one token is issued. Ultimately, there is no definitive way to determine the value of the entire pledged claim, however, even by CRIRSCO reporting guidelines there is at least 5 times as much feasibly extractable gold. There are also Indicated Resources, there are also Inferred Resources and it is likely there are accompanying metals such as copper, molybdenum, silver and others that may or may not also have extractable value.

It is important to note that the pool is not purchasing the claims or the owner of the derived assets, it is only a pool for managing the title. The title to a claim is fully pledged to pool for the life of the loan and is to be returned to the claim holder once the loan has been repaid at Maturity. The asset represented in the form of a token are held by the account holder on initial digitization. The account holder may then use the platform to exchange tokens to another account holder. That account holder may then transfer tokens to another account holder, who may or may not be the original claim holder. The nature of the agreement, and the exchange of any non-token assets including but not limited to currency, stock or hard assets is between the two parties.

The services provided to run the pool include: the review and processing of Asset Digitization Service Agreement applications for perspective claim holders; the digitization of claims into smart certificates known as the token; the smart contracts to manage the lifecycle of the Asset Digitization Service Agreement; providing a secure platform to facilitate the trading of token; to provide an immutable and auditable history of transactions and documents including but not limited to claim titles. The services do not typically hold the derived value of assets; set the value of token; negotiate the terms of any loans or transactions; facilitate the transaction of assets other than token; or arbitrate agreements.

The business logic for the pool is codified within immutable Smart Contracts. The Smart Contracts ensure the agreed upon rules are correctly adhered to for the lifecycle of the ADSA. There are (2) primary Smart Contract types: ADSA: which represent the Asset Digitization Service Agreement; and Reserve: which represent the titled asset, typically a claim. The Smart Contracts allow for optional extensions. One such example is that presently only gold can be digitized. However, a claim is a defined section of land with Mineral Resources and is not limited to a particular mineral type. Therefore, other asset types may also be digitized, and those assets will be bound to the same Reserve object. If the digitization start and end periods of different asset types within a Reserve are not aligned, freeing up the title of a claim prematurely would cause ownership issues. Therefore, the Reserve is a separate object and is bound until the last ADSA is terminated.

The Reserve object or Smart Contract represents the titled claim and has the following attributes:

“approved_timestamp”; “approver_id”; “country”; “created_timestamp”; “creator_id”; “documents”; “geolocation”; “internal_reserve_id”; “last_modified_timestamp”; “last_modifier_id”; “owner_id”; “proven_reserves”; “ready”; “reserve_description”; “reserve_id”; “reserve_type”; “signed_date”; “state”; “status”; “terminated_timestamp”; “terminator_id”.

The ADSA object or Smart Contract represents the fungible asset and has the following attributes:

“approved_timestamp”; “approver_id”; “country”; “created_timestamp”; “creator_id”; “documents”; “internal_reserve_id”; “last_modified_timestamp”; “last_modifier_id”; “owner_id”; “proven_reserves”; “asset_type”; “tokens”; “quantity”; “ready”; “reserve_description”; “reserve_id”; “reserve_type”; “signed_date”; “state”; “status”; “terminated_timestamp”; “terminator_id”.

The Substitution objects or Smart Contracts represent the Reserve and ADSA equivalent except for Refined Gold. The only difference is the reserve_type for Reserve is substitution and the resulting ADSA quantity is digitized at a 1:1 ratio instead of the standard 1:5.

The flow chart in FIG. 4 traces a simplified Reserve/ADSA lifecycle through the various states. The reserve status may be: Pending, Signed Off, Bound, Terminated ADSA status: Pending, Signed Off, Digitized, Terminated. According to FIG. 4:

1. Asset/Collateral Holder initiates a claim for review; Reserve (Pending).
2. The Reserve title is clear; Reserve (Signed Off).
3. The ADSA is submitted for review; ADSA (Pending).
4. The ADSA passes review; ADSA (Signed Off).
5. The Digitization date is set for the future; Reserve (Signed Off)/ADSA (Signed Off).
6. The Digitization date arrives and token are digitized; Reserve (Bound)/ADSA (Digitized).
 - a. The claim title is fully pledged to the pool.
 - b. The token are deposited into the Asset/Collateral Holder’s account.
 - c. The Asset/Collateral Holder may transfer token to other accounts.
7. The Maturity date arrives and the original amount of digitized token are removed from the account; ADSA (Terminated)/Reserve (Terminated)
 - a. The claim title is returned.
 - b. The token are retired from the system.

It is possible to renew an ADSA past the original Maturity date, as defined by the Smart Contract. In order to terminate an ADSA, the Asset/Collateral Holder must acquire the

original amount of tokens from the market to be retired. All of the tokens are accounted for on the system, but it is possible there may not be any for sale. To account for this anomaly an option exists for, and only for, ADSA owners who are approaching the maturity date, to pledge Refined Gold from a vault for the outstanding amount in a new Reserve. The original Reserve is terminated and the title to the claim is returned, however, there is a new obligation to recover sufficient tokens to release the pledged Refined Gold.

In the event there are insufficient tokens to cleanly terminate an ADSA on the Maturity date, the ADSA will default. Because the circumstances of each default are different and involve externalities, the Smart Contract principally flags the issue for remediation. However, to maintain the integrity of the platform, the only way to terminate an ADSA is to fully return the Digitized token.

There is no concept of a complete cycle. The smart contracts allow for any number of valid combinations. An Asset/Collateral Holder could Digitize an ADSA, renew, renew again, Digitize another ADSA, use the balance with the new claim to Terminate the original ADSA, default on the new ADSA, renew with penalties, pledge a substitution, Terminate the new ADSA, then repay the substitution at a later date, for example.

In its simplest form of the system uses distributed ledger technology (e.g., provided by Symbiont) to create a new digital asset class. While the term “blockchain” is the more widely recognized it is more technically correct to describe the platform of choice as a distributed ledger platform because transactions are appended one at a time rather than in “blocks”. Indeed, in some cases, the technology may be implemented in blocks. Despite this minor distinction, the ledger retains the properties of traditional blockchains including replication, resiliency, immutability and enforced consistency. However, when implemented as a private ledger, many of challenges of the around privacy and performance of public distributed ledgers are inherently addressed.

The distributed ledger network may include trusted member nodes so the ledger is never publicly exposed. Alternate technologies employ cryptography that permit untrusted member nodes, which process the transactions in a verifiable and authenticated manner without access to the underlying data. Each member’s data is encrypted and decrypted only by authorized members on the network. As a member, the pool leverages Smart Contracts which strictly enforce predetermined business rules. All activity is recorded on a tamper proof, append-only ledger along with times stamps and digital signatures. The pool preferably operates on a permissioned network negating the need for mining to enforce consensus. The distributed ledger preferably uses an implementation of a *Byzantine* Fault Tolerant algorithm (BFT-SMaRt, n.d.) that enforces consensus across the network. This approach provides resilience and performance orders of magnitude greater than mining, e.g., Bitcoin. All ledger data is encrypted and accessible only by authorized parties. When the pool queries the Smart Contracts, the encrypted data is read from the ledger, which only the pool member is able to decrypt. (In an alternate implementation, public verification is supported).

The news is full of stories where wallets have been compromised, cryptocurrency is stolen and there are few mechanisms to undo the damage. Private or permissioned based platforms mean everyone on the network or who has an account is a known entity who has passed a “Know Your Customer” (KYC) and AML. Compromising the system to benefit a particular account holder on a tamper-proof immu-

table ledger would be highly risky. In addition, since a service provider can administer the system, any transactions can be undone with complimentary transactions. This cannot be done in decentralized blockchains.

In addition to the data itself, the Smart Contract enforces business logic, which is also stored on the immutable ledger. Should a security hole, error or bug be identified in a Smart Contract, the distributed ledger platform may provide a straightforward mechanism for correcting it. Since all smart contract code is recorded as data on the immutable ledger, all parties have a record of both the error and the fix, and may employ legal recourse as necessary.

Along with financial data there it is possible to store various legal documentation data on the ledger as well. Examples include PDF documents, signed and scanned legal documents and stamped geological reports among others. Data, documents and business logic are all encapsulated on an immutable ledger for a completely secure and auditable solution.

The system is resilient and tolerant of failures. It can scale. Most importantly, it is a cryptographically-enforced, append only, immutable chain of all the history since inception. It is an ideal system for accountability and auditability. A permissioned distributed ledger adds no more attack vectors than the traditional stack while enabling full auditability in the event if they did happen. The present system may provide customized and modular APIs to securely interface with the platform.

A proposed information flow diagram is provided in FIG. 5.

REFERENCES

- BFT-SMaRt. (n.d.). State Machine Replication for the Masses with BFT-SMART. www.di.fc.ul.pt/~bessani/publications/dsn14-bftsmart.pdf
- BusinessWire. (2017, Mar. 15). Orebits & Symbiont Deploy Distributed Ledger Technology to Digitize Gold Ownership. www.businesswire.com/news/home/20170315005332/en/Orebits-Symbiont-Deploy-Distributed-Ledger-Technology-Digitize
- CIM. (n.d.). CIM. Canadian Institute of Mining, Metallurgy and Petroleum: cim.org/CIM-Definitions. (2014, May 10). CIM Definition Standards for Mineral Resources and Mineral Reserves. www.cim.org/~media/Files/PDF/Subsites/CIM_DEFINITION_STANDARDS_20142
- NI 43-101. (2011, Jun. 24). NI 43-101 Standards. Canadian Institute of Mining, Metallurgy and Petroleum: web.cim.org/standards/documents/Block484_Doc111.pdf

Example 3

Sovereign Financing

Financing a sovereign's deficit has become prominence for economic growth and stability, with the financial challenge, a sovereign is always in-need for a cash infusion. Recently, sovereignty auspices are not enough to get finance at a prime rate. When a government is looking to add liquidity to the economy on a non-inflationary basis, it needs to look to foreign investment and trade to accomplish this.

Many governments have proven mineral reserves but due to many reasons, these assets are not being utilized in any fashion. There is no mechanism to get the proven value of mineral reserves onto the central bank ledger to issue currency against these mineral reserves (e.g., gold) on a non-inflationary basis. The present technology permits a sovereign to finance and issue a debt obligation against a

new asset class, such as a proven mineral reserve, which can be allocated in a way to provide leverage and a line of credit, without requiring extraction of the mineral, thus permitting preservation of the resource, with extraction only required in event of default.

Utilizing the Sovereign Government Proven Gold Mineral Reserve.

The sovereign government, in this case, assigns the proven reserves to an international mining corporation as part of a public-private partnership. This is a common practice where mining rights and the reserve ownership is given out under different systems depending on the company and record keeping of the reserves through the government department assigned these duties but generically they are referred to as mining claims.

To mitigate the political risk and logistical risk of adding any new territory, the new government also provides a government guarantee as for the value of the proven reserves and to further guarantee they will allow the mine to operate and export the product without hindrance. The goal is to bring the risk in-line with the existing ledger assets so all respective securities stay fungible.

The international mining corporation then takes this claim ownership and the sovereign guarantee information to a monetization entity, which "digitizes" the in-ground value through a smart contract.

The international mining corporation then has the fungible digital assets they can be put into a trust and securitized within international financial markets with an audited value.

The international mining corporation can now pledge the digital assets in trust to a commercial bank (e.g., in the originating country). Utilizing normal bank protocols, the new asset can be pledged to the commercial bank, and the commercial bank can apply to the central bank for approval of the new crypto asset as well a pricing. The central bank can now create funds on a non-inflationary basis in the local economy including the finance activity to get the new mine into production. This scheme is represented in FIG. 6.

Although the disclosure is described above in terms of various example embodiments and implementations, it should be understood that the various features, aspects and functionality described in one or more of the individual embodiments are not limited in their applicability to the particular embodiment with which they are described, but instead can be applied, alone or in various combinations, to one or more of the other embodiments of the disclosure, whether or not such embodiments are described and whether or not such features are presented as being a part of a described embodiment. Thus, the breadth and scope of the present disclosure should not be limited by any of the above-described example embodiments, and it will be understood by those skilled in the art that various changes and modifications to the previous descriptions may be made within the scope of the claims.

What is claimed is:

1. A token system, employing a token representing an interest in a smart contract, the smart contract representing an agreement, secured by a security interest in real property or a right in real property, to return the token within a defined period, comprising:

an automated distributed virtual state machine, hosted by a plurality of cryptographic hardware processors, employing a distributed consensus model, and being configured to execute immutable transactions recorded on a blockchain, the blockchain having blocks which are created in accordance with a competitive consensus process;

55

automatically executable smart contract code which controls the automated distributed state machine, after the defined period of time, to permit exercise the security interest with respect to the real property or the right in the real property if the token is not returned; and

a communication port configured to interface with an automated distributed communication network, to communicate distributed consensus messages through the automated distributed communication network, to receive the returned token, and to communicate an immutable message for exercise of the security interest through the automated distributed communication network.

2. The token according to claim 1, wherein the automated distributed state machine comprises an Ethereum virtual code machine.

3. The token according to claim 1, wherein the executable smart contract code further permits tolling of the defined period of time for exercise of the security interest dependent on whether a substitute asset is tendered.

4. The token according to claim 3, wherein the real property or a right in real property comprises a physical mine having proven available reserves of the substitute asset.

5. The token according to claim 1, wherein the automated distributed virtual state machine charges a transaction fee for execution of the executable smart contract code.

6. The token according to claim 4, wherein the proven available reserves are a predetermined multiple of the substitute asset.

7. The token according to claim 1, wherein the token represents a fractional interest in the real property or a right in real property.

8. The token according to claim 1, wherein the token is generated as a transaction of the automated distributed virtual state machine comprising a cryptographically-authenticated, distributed ledger having a database held and updated independently by each of a plurality of distributed processing elements, forming a consensus determination of transaction validity.

9. The token according to claim 1, wherein the blockchain comprises a transaction list and a state of the smart contract.

10. A method of transacting with a token, comprising:

defining an executable smart contract, representing an agreement, secured by a security interest in real property or a right in real property, to return the token within a defined period of time, the executable smart contract being executed on an automated distributed virtual state machine, hosted by a plurality of cryptographic hardware processors, employing a competitive distributed consensus model which produces immutable blocks, and being configured to execute immutable transactions recorded on a blockchain comprising a series of the immutable blocks;

pledging the real property or a right in real property as the security interest to secure the smart contract;

issuing the token and recording issuance of the token by the on the blockchain;

controlling the automated distributed state machine in accordance with the executable smart contract to extinguish the token and the security interest when the token is tendered for return; and

controlling the automated distributed state machine in accordance with the executable smart contract to, after the defined period of time, communicate a message

56

authorizing exercise of the security interest with respect to the real property or the right in the real property if the token is not returned.

11. The method according to claim 10, further comprising returning the token, and releasing the real property or a right in real property from the security interest.

12. The method according to claim 10, wherein the smart contract is implemented in conjunction with a distributed ledger, wherein the distributed virtual state machine charges a transaction fee for execution of the executable smart contract code.

13. The method according to claim 10, wherein:

the token comprises a fractional interest in the real property or a right in real property,

the define period of time is tolled if a substitute asset is tendered,

the real property or a right in real property comprises a mine having proven available reserves of the substitute asset, and

wherein the proven available reserves are a predetermined multiple of the substitute asset.

14. The method according to claim 10, wherein the distributed virtual state machine is Turing complete, and the executable smart contract has a predefined maximum number of executable instructions.

15. The method according to claim 10, wherein the token is generated as a transaction recorded on the blockchain, and the blockchain comprises a cryptographically-authenticated, distributed ledger held and updated independently by each of a plurality of distributed computing elements, the method further comprising forming a consensus determination of transaction validity.

16. The method according to claim 10, further comprising exercising the security interest in the real property or a right in real property after the defined period of time if the token is not returned.

17. A method for creating a liquid token representation from an illiquid asset comprising:

receiving a pledge of an illiquid asset; and

digitizing the illiquid asset into fractional representations;

issuing a token representing the fractional representations, subject to an executable smart contract executed on a distributed virtual state machine, hosted by a plurality of cryptographic hardware processors, employing a distributed consensus model, and being configured to execute immutable transactions recorded on a blockchain, the fractional representations being secured by the pledge of the illiquid asset as collateral;

controlling the immutable distributed state machine in accordance with the executable smart contract to extinguish the token and the security interest upon return of the liquid token; and

controlling the immutable distributed state machine in accordance with the executable smart contract to, after the defined period of time, communicate a message authorizing exercise of the pledge of the illiquid asset as security for the outstanding liquid token.

18. The method according to claim 17, further comprising trading a fractional representation on an exchange, recorded in a distributed ledger network.

19. The method according to claim 17, wherein the smart contract comprises an illiquid asset description, an identification of an illiquid asset owner, and at least one redemption rule.

20. The method according to claim 19, wherein the at least one redemption rule comprises a maturity date establishing the defined period of time and a tolling rule for the delaying

the maturity date, wherein a satisfaction of the at least one redemption rule triggers a release of the pledged illiquid asset back to the illiquid asset owner in exchange for return of all of the original fractional representations or substitute collateral.

5

* * * * *