

1/16

How to Build A Brain-To-Brain Internet:

Proof-of-Cognition Blockchains as a Foundation for Brain-to-Brain Security

April 30th, 2015evin Mauro

1. Motivation
2. Reverse-Engineering The Brain
3. Proof-of-Work Blockchains
4. Proof-of-Cognition Blockchains
5. Proof-of-Work Versus Proof-of-Cognition
6. Services in a Brain-to-Brain Internet
7. Neural States as a Method for Private Key Generation
8. Conclusion

Abstract

Security in a brain-to-brain network is crucial for maintaining individual autonomy. A digitized weave of human conscience will always be vulnerable to computational attacks unless its security relies

on biological barriers. The barrier must be easily sampled, difficult to model, and consciously harnessable. One such barrier is the human brain.

Within the realm of digital cryptography, blockchains offer a means of “hack-proof”, publically-verifiable information. Traditional blockchains, however, are currency-based, and rely on

mining hardware that could be re-routed or over-produced by artificial intelligence. When a number of

human consciences are at stake, a better solution would be a blockchain tied to a biological (“biocryptographic”) barrier. Neurally-based cryptography could one day bar artificial intelligence

from the network, while preventing abuse by dishonest human nodes. Publicized connection and termination to services prevent connections from exceeding an agreed upon duration and quality, and

alert users of potential network takeover. Future applications include efficient experimentation of brain-based medicine, decentralized elections and global governments, and an entirely new medium for

communicating information and experience.

This paper outlines a roadmap for constructing a “proof-of-cognition” blockchain. Various mechanisms for generating private keys from multi-unit neuron recordings of live rhesus macaques are

also evaluated.

1. Motivation The same way the digital internet transformed how we share information, a brain-to-brain

internet can transform how we share humanity.

Leading experts in the field of artificial intelligence estimate that by 2040, a technological singularity will cause an unpredictable “intelligence explosion”, after which the capabilities of machines will supersede that of humans (Armstrong [2012]; Carvalko [2012]; Eden [2013]). On the other end of the spectrum, leading neuroscientists contend that no machine will ever be able to compete with the non-linear, non-Turing prowess of the human brain (Nicoletis [2015]). Through biological barriers to computation, we can harness the power of our brains to isolate ourselves from unforeseen advances in machine-based artificial intelligence. Consequently, new methods of information transfer between humans may ignite an unpredictable intelligence explosion which rivals (or exceeds) that of machines.

For centuries, individuals have proclaimed their own existence with the phrase *cogito ergo sum* (“I think, therefore I am.”) (Descartes, 1685). Communicating thoughts between each other could, for the first time, prove the existence of individuals outside ourselves. Furthermore, an accused (or racially-persecuted) person could prove their innocence by sharing recorded neural patterns during the time of the arrest, a sports fan could experience the adrenaline and mechanical motions of their favorite athlete, a layperson could taste a restaurant's best meal from across the globe. A blind child could receive visual input from its mother, mental states (hunger, happiness, excitement) could be quantified and tracked, infrared-light detectors could expand our senses (Thomson et al. [2013]). Permanent external storage of thoughts and memories could greatly enhance information recall, and information could be translated and analyzed in ways not yet imagined. A brain-to-brain network would not be limited to humans. The neural intelligence and sensory input of other animals could also be harnessed (Pais-Vieira [2013]; Trimper et al. [2014]).

A secure mechanism to tie a human lifeform to a digital identity can push our governments onto the internet, enabling world passports, transparent elections, and a true, global democracy. In such an identity network, contracts and digital payments can be initiated by thought, files and assets can be forwarded elsewhere upon death, sensitive information can be shared only after a specific neural impulse. Note that DNA offers a mechanism for a biological identity, but not a digital one. DNA can be shed, and thereafter, copied. A better form of identity would be one that is unhackable, digital-friendly, and disposable. Such a form of identity could become the basis for bio-digital signatures, filling in the gap between the virtual and natural.

A proof-of-cognition blockchain as an underlying identity-network for a brain-to-brain internet would provide sufficient autonomy for each of its users. If cryptographic keys were generated and stored on an offline, physically-inaccessible, neurally-trained implant, hacking a

person's identity would be impossible. Decentralization of the network would guarantee that all users had equal power, and that a single ill-acting party could not cause sweeping changes across the network. In the event an ill-acting party did enter the network, the public nature of a blockchain would alert its users, ensuring honest nodes could exit or reject the dishonest node before harm were spread. So long as the majority of nodes remained honest, a proof-of-cognition blockchain can maintain the safety of an individual's conscious in a brain-to-brain network.

This paper proposes a pseudo-anonymous digital-biological network as a foundation for later brain-to-brain innovations. A rudimentary understanding of hashing, blockchains (Dai [1998]; Back [2002]; Nakamoto [2008]) and modern brain-machine interfaces (Lebedev and Nicolelis [2006]; Lebedev [2014]; Hildt [2015]) is recommended.

2. Reverse-Engineering the Brain *The brain is poorly understood, but adequately complex.*

In the brain, sharp voltage spikes are propagated from neuron-to-neuron to communicate information (Hodgkin and Huxley [1952]). These voltage spikes are called action potentials. Action potentials can be elicited by a biological stimulus (such as eyes detecting light), or from the random fluctuation of ions crossing a neuron's biological membrane (Diba et al. [2004]; Kole et al. [2006]; Faisal et al. [2008]). The noisiness of action potentials has led movement-based brain-machine interfaces to make use of only a small fraction of neurons that are well-correlated with a particular movement (Chapin and Nicolelis [1999]; Laubach et al. [1999]). The remaining neurons are a source of entropy. How the brain makes use of such a high level of noise is poorly understood (Dorval and White [2005]; Faisal et al. [2008]). A complete understanding would likely require detailed modeling on a molecular scale.

Today's best recording implants can wirelessly record from hundreds to thousands of neurons (Schwarz et al. [2014]). For a brain-to-brain network to function optimally, stimulation and recording of neurons across multiple cortical layers must be engineered. Hundreds to millions of implantable, free-floating sensor nodes show promise for high-density, biocompatible brain recording (Seo et al. [2013]). Figure 1 depicts one recent pioneering innovation, aptly termed 'neural dust'.

Figure 1: Ultrasound waves are used to interrogate implanted neural dust. Changes in dust composition is correlated with changes in electrical potential. Neural dust offers an effective means of extracting a large amount of information from multiple cortical layers, but may have insufficient resolution for full-scale brain recording. Theoretically, neural dust particles of a different piezoelectric composition could cause neural stimulation following electromagnetic interrogation, allowing for parallel recording and stimulation of the brain. Taken from Seo et al. [2013].

Action potentials can also be elicited through foreign electrical stimulation ([Hodgkin & Huxley 1952]). Clinically, artificial neural stimulation has proven to be an effective treatment for depression, bipolar disorder, schizophrenia, (McNamara et al. [2001]) and more (STX-Med [2014]).

In a binary model of a neuron, a neuron in the midst of an action potential is considered a '1', while a neuron at rest is equivalent to a '0'. With greater than tens of billions of neurons in the brain, there are at least 210,000,000,000 possible states of the brain at any given moment. Given the brain is not truly binary, but non-Turing, modeling the human brain on a metabolic, molecular, and electrical scale remains a challenging computational problem (Yoosef et al. [2014]). Based on Intel's BlueGreen experiments (Yoosef et al. [2014]), Moore's Law indicates that it will take approximately 60 years until a computer may be capable of fully modeling a simplified, generic rat brain. Considering the significant variations from person-to-person in the upper cortex (Kelly et al. [2012]), it may never be possible to successfully model the brain of a living being.

In 2013, the world's first brain-to-brain interface was constructed between two rats in the laboratories of Miguel Nicolelis (Pais-Vieira [2013]; Nicolelis [2015]). Sensory information was

translated to motor action between one rat located in Durham, North Carolina, to another rat in Natal, Brazil. Critics of this experiment claim that information transfer rates between the two rats were not close to the information transfer rates of computers, or even modern brain-machine interfaces. This should not discourage research into brain-to-brain interfaces, but rather, signify a dangerous lack of inquiry into brain-to-brain interfaces relative to computer circuitry.

In the same year that the first brain-to-brain interface was constructed, U.S. President Barack Obama founded the BRAIN Initiative, a funding effort intended to boost the U.S. to the forefront of brain innovation (NIH [2014]). Unfortunately, our lack of understanding in neuroscience is correlated with a stark lack of funding compared to computational research. Even with the BRAIN Initiative in mind, when considering investment from both governments and private institutions, yearly U.S. funding into brain research is less than 10% of the hundreds of billions poured into computer science research (SFN [2011], Kennedy [2012], NIH [2014]). In order to reverse-engineer the brain at a rate similar to advances in computer science, these numbers must be flipped and exceeded. Cooperation is required from industry leaders, governments, and philanthropists to fund neurobiological and brain-machine interface research, particularly because the additional regulations and experimental time necessary for biological research will always exceed that of hardware and software research. So long as we continue to innovate computer circuitry and neglect biological integration, the computational abilities of machines will ultimately surpass humanity's collective intelligence.

3. Proof-of-Work Blockchains On a blockchain, data is made public, verifiable, and is resistant to change.

Publicized, mathematically-verifiable data can transform digital communication the way Darwin's theory of evolution transformed biology. For the first time in history, digital data has an essence of physicality. On a blockchain, data is “collectable”, uncopyable, and unforgable. By announcing information publically and locking it in place cryptographically, blockchains permit an entirely new means of validating digital information.

A blockchain connects blocks of data into a time-stamped chain. A blockchain is engineered such that anyone can efficiently validate what is in the chain, using a hashing algorithm (Nakamoto [2008]). Opposingly, a centralized structure, such as a server, holds data which can be manipulated by a single, ill-acting operator. If an ill-acting party attempted to manipulate data in a blockchain, the party's new blockchain would be mathematically rejected by other participants in the network. This enables a blockchain to be a source of information where responsibility for data-integrity relies on every participant. Thus, a blockchain is a “decentralized” data structure.

To add a block to a blockchain, a participant in the network (called a 'miner') must 'mine' three parts of a block. The first component of a new block is a hash of the previous block. This is

easy to find, but arguably the most important. The hash of the previous block ensures that a new block is not tied to a forged blockchain (that is, forged information). The second part of a block is a data type. Bitcoin, a blockchain currency system, publishes currency transactions between “bitcoin accounts” (ECDSA-generated public keys) (Nakamoto [2008]). By listening for transactions over the internet, it is somewhat easy for a miner to attain a valid list of recent bitcoin transactions. The third, most sought- after component, is a nonce, or a random value. This value is altered by the miner. By altering the nonce, a miner changes the hash value of the block. A block will only be validated by other miners if its hash begins with a certain number of leading zeros. Upon finding a block, a miner is rewarded by the network, typically through newly-generated digital currency.

Figure 2: Blocks integrate a hash of the previous block, a nonce, and groups of transactions. Transactions consist of inputs from previous transactions, and outputs to new public keys. They are digitally-signed to transfer ownership. Taken from Nakamoto [2008].

A very fast computer can make more guesses in less time, and is more likely to find a block. When many powerful computers are competing to find a block, the network adjusts itself to increase the required number of leading zeros for a block's hash. The more leading zeros required, the more difficult it is to find a block. This increased difficulty ensures that the average amount of time to find a block is predictable, allowing for a reasonable amount of time to receive transactions (typically 2 or 10 minutes). However, this also necessitates increasingly powerful computers for one to find a block and receive a reward. This 'computational arms race' leads to an absurd amount of wasted energy to be spent guessing nonces. Fortunately, this arms race also ensures data integrity. If someone wished to alter (forge) a transaction in a blockchain's history, they would have to find many hard-to-find nonces in a short amount of time. As of 2013, the computational barrier in bitcoin's case is greater than 128- times the power of the world's top 500 supercomputers combined (Cohen [2013]). This qualifies a "proof-of-work" blockchain.

Once a new block is found, it (including the nonce) is broadcast to all other miners in the network. After including the nonce in their own block, a miner can verify the included transactions correspond to those received by the network during the time the block was mined. They do so by hashing the new block and determining if the output has the required number of leading zeros. (Hashing is a very fast process as opposed to guessing nonces.) Thus, data is verifiable by anyone wishing to listen for transactions and hash a proposed block.

Additional security follows with the open-source nature of blockchains. Anyone can see the code underlying the proof-of-work protocol. If there were a flaw, the protocol could be modified and re-distributed. The longer the code has been made public, the less likely there is a fatal flaw yet undiscovered.

Transactions in a blockchain simply credit and debit public keys with a broadcasted value. Using digital signatures, coins can be traced back to the newly-generated block from which they originated. A full description of proof-of-work blockchains, including the detailed structures of transactions, is covered in Nakamoto [2008].

4. Proof-of-Cognition Blockchains Proof-of-cognition relies on human recognition to prevent nefarious action. Valid chains are the ones most-verified by nodes that have proven their humanity.

A proof-of-cognition blockchain links a published trail of transactions to a human-cryptographic identity. In the following descriptions, a typical node is assumed to be human. An ill-acting node is assumed to be a computer, or a human who wishes to bring a computer onto the network. An identity on the network can be disposed of at will, though a new

identity will then have to work to regain the

trust of other nodes on the network if it wishes to mine. Care is taken to preserve a node's anonymity. Proof-of-cognition only proves that a node is human– not necessarily which human they are. Two mechanisms for validating humanity are random and local verification. Random verification (4.3) would connect nodes across the planet, ensuring one group could not isolate another group from a particular blockchain. Local verification (4.4) would rely on neural models that involve recognizing encounters with another human, or a well-trusted friend. These two methods can work together to ensure optimal security for a proof-of-cognition blockchain. Novel methods of verification should be considered in addition to the two methods outlined here.

A proof-of-cognition protocol must rely on tasks that only humans can perform (4.3, 4.4). It should be immune to unforeseen advances in computing. An optimal solution would be one that relies heavily on biological verification of another's humanity, with neurocryptographic barriers against nefariously-acting humans and computers. The following sections outline a proof-of-cognition protocol for achieving these goals.

4.1 Mutual Transactions Quantify Trust

A node will only transact with another node if it deems the latter node to be human. Simultaneous, mutual transactions between two nodes signify a neural 'handshake', indicating that each node acknowledges the other's humanity. This acknowledgement relies on both conscious and unconscious verification (4.3, 4.4). Mutual transactions included in the same block are considered 'simultaneous' (Figure 3).

Figure 3: Only transactions in which each node reciprocates verification are included in blocks. This prevents network spam and limits blockchain growth. Mutual transactions between poorly-verified nodes are not recorded in the blockchain.

An inaccessible device, designed and surgically implanted in a transparent, open-source manner, would ensure a person's cryptographic miner could not be hacked through hardware or software means. Unsigned transactions could be locally communicated to the offline implant. They would then be digitally signed and returned without risk of intervention. Keys would be used for one transaction, then disposed of. Each transaction would have one input and two outputs– one output to a user's newly generated public key, a second output to a verified human, or secondary service (6.0). The trail of transactions between one's previous and newly-generated public keys would qualify as a person's identity. Identities could be efficiently tabulated (and audited) in conjunction with the blockchain, rather than having to trace an identity each time a user wished to evaluate another node's level of verification.

A proof-of-cognition blockchain need not trade currency. Instead, nodes would trade a valueless data structure called 'trust'. Whether a node's humanity was accepted or rejected by another node

would necessitate posting on the blockchain. Transaction inputs and outputs would have no inherent value, other than to signify the acceptance or rejection of a node. The sum of “trust” in a wallet has no meaning. Rather, mutual transactions would quantify trust.

4.2 Verification Webs Enable Blockchain Innovations

A node's trust level can be evaluated by tracing the history of blockchain transactions to determine who the node has mutually traded trust with. If a node has a small degree of separation from another node (Newman et al. [2006]), the two are assumed to trust one another. An over-arching web of verification can then be constructed among nodes that are shown to be human (Figure 4a). Even if dishonest (artificial) nodes begin verifying each other through non-biological means, other honest (human) nodes would reject their proposed transactions, barring them from the blockchain(s) that relied on the web of human trust (Figure 4b). Due to unconscious verification (4.4), honesty is inherently tied to a person's biology.

Figure 4: Verification webs develop with increasing mutual transactions between nodes. A solid line indicates a mutual transaction. The perspective of trust is given by the node with the black star, though as more nodes become verified, any node near the center of trust is similarly-defined. a) depicts the growth of the web as more nodes pass mutual connections. b) demonstrates that nodes disconnected from the star-node are not accounted for in the verification web. From the second to third panels, the center of trust moves leftwards due to the increasing trust among those nodes. The node with the black 'X' is deemed questionable due to its singular, distant connection from the center of trust. It is deemed fraudulent (and ignored) after trusted nodes have tested it and communicated their rejection to the blockchain (dotted lines). Note that because verification is partly unconscious, many intuitive methods of abuse are impossible (4.4). Contrary to the 2D representation above, a web of verification would be n-dimensional. This concept is quite similar to “six degrees of separation” (Newman et al. [2006]).

With an overarching web of verification, maintaining a blockchain need not be a waste of energy. By digitally-signing found blocks with their identity, miners could reject blocks if they were proposed by nodes that 1) had a low-level of trust 2) had recently found an improbable number of blocks based on the assumed hashing power. A node could be assumed to have a limited hashing power,

perhaps even a small enough hashing power to be biologically-powered (Scholz and Schröder [2003]). Currency rewards are not necessary for honest mining, due to the negligible cost required for hashing, and the valuelessness of the data type being traded. Human-to-human transactions on a proof-of-cognition blockchain would not have fees because transactions between unverified nodes can be ignored by the central verification web, preventing network spam (Figure 3). However, human-to-service transactions would necessitate digital currency fees to both prevent network spam and ensure non-miners can dispose of identities when they wished to become anonymous (see 6.0).

How resilient is proof-of-cognition to double-spend attacks? In proof-of-work, double-spend attacks re-write a miner's transaction to reward a miner with digital currency. In proof-of-cognition, double-spends could be used to bring an ill-acting node onto the network by changing a verification transaction from one node to another. However, by modifying one's own transaction to link to a distal, ill-acting node, a particular human would only decrease their own verification level. Since mining difficulty is partly based on the verification level of the node mining the block, it would become increasingly difficult for a miner to bring ill-acting nodes onto the network the more blocks they modified.

A computer posing as a human must attain verification from a substantial number of humans in order to gain entrance into the web of trust. A node's trust could be liquidated if it attempted to bring new nodes onto the network too quickly, or if it repeatedly tried to enter the same node onto the human web.

Human nodes gravitate toward one another through continued mutual transactions. Similar to a proof-of-work protocol, a proof-of-cognition protocol only functions if a majority of nodes remain honest.

4.3 Neurocryptographic Barrier: Neural Hashing

In addition to the barriers offered by normal cryptography, neural hashing may offer a biological means of cryptographic protection. If the brain is modeled as a one-way black box function, where the input is a random stimulation, and the output is the recorded firing of neurons, the brain becomes a biological hashing function incapable of being modeled by machines. With full-fledged recording and stimulation of the brain, an initial, dense, time-varying stimulation could produce an output that propagated across many neurons. A second, diffuse, low-intensity stimulation could then be initiated which mimics the first stimulation's neural output (Figure 5). This second stimulation would stimulate individual neurons, rather than propagating across many. If the two stimulations 'feel' the same to a human, a human could consciously verify paired stimuli proposed by another human. It would be difficult or impossible for a human to forge such computations, especially if it were also made to resemble a human being (4.4).

Figure 5: a) Human A sends a neural stimulus, along with their respective output to said stimulus. The output is used to create a low-intensity stimulus that mimics the distal propagations of the primary, dense stimulation. Human B then stimulates themselves and determines if the two stimuli are sufficiently similar. b) Human B then sends its own stimulus and output to A. If A finds its own output to this stimulus to be sufficiently similar to B, two transactions are then broadcast. c) Trust is built between A and B, as publicized on the blockchain through a mutual transaction.

Alternatively, computational verification (an algorithm that compares a stimulus to an output), would probably require computational power on the order of modeling (and thus forging) neural outputs. Further research is necessary to determine what sequences of stimulation would provide outputs with the greatest entropy, as well as whether biological (conscious) verification in such a method is even possible.

This method offers a mean for random verification, in the sense that it can be performed between two nodes that have never met in person. This can prevent a group of nefarious individuals from bringing an ill-acting node onto the network, or from excluding honest-acting nodes from a particular blockchain. In the next section, local verification shows promise for realistically authorizing a familiar human onto the network (4.4).

4.4 Biological Verification: Human Recognition

While computers are rapidly improving their ability to recognize humans, the opposite is unfeasible— humans have no problem recognizing that a computer is not a human being. To design a computer that was indistinguishable from a human being, one would need a plethora of biological pheromones, realistic skin and eyes, complex motor movements, proper body heat, and the ability to package all that and more into a fully-functioning model organism. Neural models for human recognition are essential for a proof-of-cognition protocol. In this context, two mechanisms worth considering are conscious and unconscious verifications.

A second conscious verification could work externally. For example, Human A encounters Human B. Each human uses an external device (a simple NFC tap of a cell phone) to mutually share verification. Unfortunately, this level of conscious action would not only be tedious, but would have significant potential for abuse. An external device could be hacked, or a group of nefarious individuals (such as a political party) could introduce a computer onto the network for personal gain. Thus, a purely conscious protocol is invalid.

Unconscious verification would rely on recording a person's neurons as they encountered another human being. A nearby human's current public key on the blockchain could be communicated through low-energy bluetooth, unconsciously, as they encountered another human in day-to-day life. A simple touch (even a physical handshake) could be recognized by an adaptable actor-critic learning algorithm (Gürel and Mehring [2012]; Prins et al. [2014]; Roset et al. [2014]) as a meeting with another

human being. Alternatively, models that rely on recognizing friends rather than superficial encounters can further ensure integrity of the network. A second conscious verification may come with disposing accidental unconscious verifications.

A combination of the random verification in 4.3 and the local verification in 4.4 appears to coalesce the best of both worlds. A node's trust level could be calculated based on a combination of both random and local verifications. Conscious, computationally-resistant verification prevents a computer from entering the network. Unconscious, equally-resilient verification prevents human nodes from acting in a self-serving manner. The public keys of each human that a person encounters during a day could be queued internally, selected at random, then verified after sharing a stimulus and output at a later time. The queued public keys would be dissociated from the original identity, preventing lock-out of particular humans, and ensuring a specific human's encounters could not be tracked.

Models for accurate human recognition based on a number of neural indicators are necessary for a proof-of-cognition protocol. As the section 6.0 illustrates, the proof-of-cognition protocol outlined above is necessary for a fully-secure brain-to-brain internet.

5. Proof-of-Work Versus Proof-of-Cognition Each protocol serves two distinct purposes.

Can a proof-of-work (POW) protocol substitute or supplement a proof-of-cognition (POC) protocol? It is possible, but not ideal. Humans in a POC protocol have equal mining power, instead using human biology to secure human conscience. If mining power were unequal (as is the case with POW), human consciences could be manipulated— a debatably far worse outcome than a simple double spend in a currency system. By relying on machines rather than biology, the network can be overpowered by artificial intelligence producing their own mining hardware, or re-routing existing mining power to reap digital currency rewards (BGP hijacking). Furthermore, miners in a POW protocol are motivated by currency rewards for honest mining. Human-to-human transactions would need to be made feeless if humans were to continuously verify each other. Determining which verification webs were human-based would be difficult or impossible from a POW miner's perspective.

Can POC replace POW? The simplest reason it cannot is that miners in a POC protocol would be able to inject bad blocks into the network, sending themselves currency when they had previously had sent it elsewhere. Since POC trades a valueless data structure, this risk is nonexistent.

Proof-of-work is ideal for currency, while proof-of-cognition is ideal for identity.

6. Services in a Brain-to-Brain Internet A proof-of-cognition protocol prevents brain-hacking.

Imagine a service called YouLive. Rather than sharing user-uploaded videos (as YouTube

does), YouLive would share user-uploaded experiences. In such a system of shared experience, care must be taken to ensure that a user's downloaded experience does not exceed the length of the recording, and that the downloaded experience is of the quality expected by the receiver. A user's brain must be fundamentally protected from becoming 'trapped' in another (potentially ill-acting) user's experience.

On a POC blockchain, safety is ensured by publically-announcing connections between humans and services. The value or script of the transaction would be signify the length of time a connection was assumed to be valid. A human would confirm a connection time after receiving a service's proposal. Once a connection was terminated, a human would then send a second transaction to signify

they had been disconnected. The publicized initiation and termination of a connection is crucial to the safety of the network. If a connection were not terminated after the specified time, users on the network would notice on the blockchain, and could remove the ill-acting service from the network, or simply leave the network themselves. Once again, the decentralized nature of blockchains offer a safer means for these broadcasts than a centralized server, the latter of which can be hacked. Thus, a blockchain is absolutely necessary as a foundation for brain-to-brain security.

In the bitcoin protocol, a small number of 'emergency broadcast keys' are given to core developers to alert users in the case of network failure. These keys could be used to automatically disconnect users in the event a number of downloaded thoughts or service connections did not receive their termination transactions.

In the case of human-service handshakes, a minuscule digital currency fee would be incurred to allot for the cost of blockchain data storage, as well as to help support the centralized structure hosting the experiences. If the service were decentralized (which it should be), a portion of the digital currency fee could be sent to the user who uploaded the experience, or whatever other feature the service offered. This fee would be necessary as an incurred cost for network spam, and relatedly, would ensure human nodes with a poor level of trust could still have their human-service transactions posted on the blockchain.

The protocol for a brain-to-brain internet outlined here is one that is discontinuous. Downloaded thoughts and service connections would only be utilized for an agreed-upon length of time. Code signing (Kiehlreiber and Brouwer [2006]) would ensure that downloaded thoughts were unaltered, similar to how file downloads on the internet are secured. Continuous, live stimulation of neurons will probably never be a safe mechanism for brain-to-brain communication, though refreshable service connections may allow for continued data sharing.

7. Neural States as a Method for Private Key Generation Can we digitize our thoughts to keep our brains secure?

Storing passwords on a neural implant offers significant advantages to other means of storage. A password on a contact lens can be removed without harm to the user. To remove a password from a brain, the attacker must commit first-degree murder, and even then, the neurally-trained implant would offer a high-level of protection from someone trying to steal one's password. Thus, it is unlikely one would even attempt murder for a password, nor would torture be effective since passwords would never be sent outside the brain. Thus, on a human brain, passwords have never been safer.

Private keys are essentially very random passwords used in cryptographic algorithms. They can be generated through a recording sample taken from a brain. Brain-generated

passwords are ideal because an individual's connectome may show sufficient variation to prevent a hacker from all-out network attack (Kelly et al. [2012]). The far greater number of states in the human brain (210,000,000,000) versus typical cryptographic algorithms (277) are encouraging for this endeavor. If a (quantum) computer were ever powerful enough to guess a private key generated by an implant-stored algorithm, the network would collapse.

To evaluate the potential of neurons for private key generation, recordings from 384 to 640 neurons were taken from rhesus macaque primates during a center-out task (Li et al. [2009]). Data was sorted into 10ms and 100ms bins. A neuron that had fired within the bin window was given a '1', while those that were quiet were given a '0'. For each time window, the neurons were randomly sorted into groups of 32, constructing a 32-bit number based on firing state. Over 14 million 32-bit numbers were generated in this manner. After excluding 0 (the 32-bit number), each of the 14 million numbers were

unique. However, the greater prevalence of the '0' bit-state (quiet neurons) resulted in a non-uniform distribution (Figure 6), invalidating this rudimentary method as a means for private key generation.

Figure 6: The top panel depicts a histogram for the distribution of the total dataset of 32-bit numbers, sorted into 300 bins. The middle and bottom panels depict histograms for two particular days of recording, sorted into 300 bins. The non-uniform distribution in each case indicates a vulnerability in private key generation.

While only 32-bit numbers are utilized here, fully uniform 32-bit numbers can be concatenated to create larger numbers, and thereafter, significant computational barriers for private key guessing. It is important to use a dataset as uniform as possible, else attacks can be utilized which make use of the most commonly-generated numbers. Simplifying to the 32-bit-state allows for increased sample size and more efficient data processing.

A second approach utilized a combination of biological and computational means. Using MATLAB's random number generator, a neuron that had been given a '0' bit-state had a 25% or 35% chance of being altered into a '1'. Figure 7 depicts a new level of uniformity of the dataset. While the data is increasingly uniform, the partial reliance on a random number generator is somewhat concerning. Fortunately, this combination is much more secure than simply using a random number generator on its own. With the neurons partially-generated by a persons brain, the advantages to individuality are maintained without full reliance on a random number generator.

Figure 7: a) Full distribution for neurons that had a 25% chance of changing their bit-state from a '0' to a '1'. c) Full distribution for neurons that had a 35% chance of changing their bit-state from a '0' to a '1'. It is probable that the 'sweet spot' lies somewhere between 25% and 35%.

A third approach was again purely biological. This time, the number of times a neuron had fired in a particular bin was used to influence the bit-states in the neuron's past bins. In Figure 8a, if a neuron had fired greater than twice in a particular bin, the previous three bins for that neuron had their bit-state set to '1'. In Figure 8b, if a neuron fired x times in a particular bin, x previous bins were set to '1'. There again appears to be increased uniformity when just comparing these two mechanisms, but not enough for private-key generation. It is possible that a variant of this mechanism could be used, however, the storage and manipulation of past neural recordings offer a potential security vulnerability, as a hacker may be able to edit the stored bit-states to a private key of their choosing. Real-time private-key generation is preferred for optimal security.

Figure 8: a) Full distribution for active neurons that, after firing at least three times, had their previous three bins changed from a '0' to a '1'. b) Full distribution for neurons that had x previous bins set to '1', where x is the number of spikes in a particular bin. In this case, bins were not "changed" from 0 to 1 – if a neuron already had a '1' in a previous bin, the bin was included in the x count. Comparison between a) and b) shows increased uniformity.

Figure 9: a) Full distribution for outputs in which the results of two random neurons were merged. If either neuron had at least one spike, a '1' was recorded, otherwise a '0' was recorded.. b) Full distribution for a similar case in which three neurons were used. There appears to be slightly more uniformity in the 3-neuron case, indicating that grouping neurons could produce increased uniformity.

A final approach merged neurons. If either of two neurons had a spike, a '1' bit-state was recorded. A similar method was used for three neurons. Outside the first peak in Figure 9, there appears to be increased uniformity elsewhere in the dataset. This is promising because merging neurons offers no potential security vulnerabilities. When recording from 1,000 neurons, 10 neurons may be merged to offer a possibility of 2100 bit-states. This still exceeds the 277 offered by most private-key generation algorithms. As the number of recorded neurons increases, one can expect merged neurons to play a greater role in private-key generation.

8. Conclusion

This paper explored a simple binary model in a small number of neurons. The simple, real-time characteristics of this method is enticing because it allows more time for spike sorting and does not “hide” potential vulnerabilities behind a complex algorithm. However, utilizing a stored random number generator may be a security vulnerability (Figure 7), and so is minor data storage (Figure 8). Further research into merging neurons (Figure 9) should be conducted to determine whether a fully- uniform dataset can be constructed. It is anticipated that as the number of recording neurons increases, so will the ability to generate private-key generation algorithms from neurons. Note that specific motor movements did not consistently elicit particular private keys. This can be explained by the immense amount of noise rampant within the brain.

Encapsulated within each of us is a relatively untapped, multi-billion dollar supercomputer. Our brains are capable of performing particular tasks with a power-efficiency and computing-strength greater than any arrangement of hardware. They are also incredibly talented at recognizing what makes us human, offering potential for a proof-of-cognition protocol to tie digital identities to biological ones. Additionally, an identity-based blockchain offers a safer means of data communication than a centralized server, which can be hacked. Unfortunately, our understanding of the brain lags significantly behind our understanding of machines. Further research into neural engineering may prove that a brain-to-brain network will exponentially increase humanity's collective intelligence.

Citations

- Armstrong, S. (2012), 'How We're Predicting AI'2012 Singularity Conference'.
- Azuma, M.; Ito, Y.; Takano, K.; Tachibana, T.; Koyama, S.; Takada, Y. & Wakisaka, T. (2009), 'Glucose fuel battery powered by yeast—analysis of the battery for high performance ', *Journal of Bioscience and Bioengineering* 108, Supplement 1(0), S129 - S130.
- Back, A. (2002), "'Hashcash - a denial of service counter-measure'".
- Barnett, Mark W; Larkman, P. M. (2007), 'The action potential', *Practical Neurology*.
- Brown, R. G. (2015), 'dieharder'.
- Carvalko, J. (2012), The Techno-human Shell-A Jump in the Evolutionary Gap, Sunbury Press.*
- Chapin, J. K. & Nicolelis, M. A. (1999), 'Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations ', *Journal of Neuroscience Methods* 94(1), 121 - 140.
- Cohen, R. (2013), 'Global Bitcoin Computing Power Now 256 Times Faster Than Top 500 Supercomputers, Combined!', *Forbes.com*.
- Dai, W. (1998), 'b-money'.
- Descartes, René. (1685) *Principia Philosophiae*. (Amstelodami): n.p., Print.
- Diba, K.; Lester, H. A. & Koch, C. (2004), 'Intrinsic Noise in Cultured Hippocampal Neurons: Experiment and Modeling', *The Journal of Neuroscience* 24(43), 9723-9733.
- Dorval, A. D. & White, J. A. (2005), 'Channel Noise is Essential for Perithreshold Oscillations in Entorhinal Stellate Neurons', *The Journal of Neuroscience* 25(43), 10025-10028.
- Eden, Amnon; Moor, J. S. J. Steinhart, E., ed. (2013), Singularity Hypotheses: A Scientific and Philosophical Assessment, Springer.*
- Faisal, A. A.; Selen, L. P. & Wolpert, D. M. (2008), 'Noise in the nervous system', *Nature Reviews Neuroscience* 9, 292+--.
- Gürel, T. & Mehring, C. (2012), 'Unsupervised Adaptation of Brain-Machine Interface Decoders', *Frontiers in Neuroscience* 6, 164--.
- Hildt, E. (2015), 'What will this do to me and my brain? Ethical issues in brain-to- brain interfacing', *Frontiers in Systems Neuroscience* 9, 17--.
- Hodgkin, A. L. & Huxley, A. F. (1952), 'A quantitative description of membrane current and its application to conduction and excitation in nerve', *The Journal of Physiology* 117(4), 500--544.
- Kelly, C.; Biswal, B. B.; Craddock, R. C.; Castellanos, F. X. & Milham, M. P. (2012), 'Characterizing variation in the functional connectome: promise and pitfalls ', *Trends in Cognitive Sciences* 16(3), 181 - 188.
- Kennedy, J. V. (2012), 'The Sources and Uses of U.S. Science Funding', *The New Atlantis*.
- Kiehlreiber, P. & Brouwer, M. (2006), 'Method and apparatus for incremental code signing', *Google Patents, US Patent 7,103,779*.
- Kole, M. H. P.; Hallermann, S. & Stuart, G. J. (2006), 'Single Ih Channels in Pyramidal Neuron Dendrites: Properties, Distribution, and Impact on Action Potential Output', *The Journal of Neuroscience* 26(6), 1677-1687.
- Laubach, M.; Shuler, M. & Nicolelis, M. A. (1999), 'Independent component analyses for quantifying neuronal ensemble interactions ', *Journal of Neuroscience Methods* 94(1), 141 - 154.
- Lebedev, M. (2014), 'Brain-machine interfaces: an overview', *Translational Neuroscience* 5(1), 99-110.
- Lebedev, M. A. & Nicolelis, M. A. (2006), 'Brain-machine interfaces: past, present and future ', *Trends in Neurosciences* 29(9),

- Li, Z.; Hanson, T. L.; Lebedev, M. A.; Henriquez, C. S. & Nicolelis, M. A. L. (2009), 'Unscented Kalman Filter for Brain-Machine Interfaces', *PLoS One* 4(7).
- McNamara, B.; Ray, J.; Arthurs, O. & Boniface, S. (2001), 'Transcranial magnetic stimulation for depression and other psychiatric disorders', *Psychological medicine* 31(07), 1141--1146.
- Nakamoto, S. (2008), 'Bitcoin: A Peer-to-Peer Electronic Cash System'.
- Newman, M.; Barabasi, A.-L. & Watts, D. J. (2006), *The structure and dynamics of networks*, Princeton University Press.
- Nicolelis, M. A. (2015), 'Miguel Nicolelis: Brain-to-brain communication has arrived. How we did it', TED.
- NIH (2014), 'BRAIN Initiative', U.S. Government.
- Pais-Vieira, Miguel; Lebedev, M. K. C. W. J. N. M. A. L. (2013), 'A Brain-to- Brain Interface for Real-Time Sharing of Sensorimotor Information', *Sci. Rep.*.
- Prins, N. W.; Sanchez, J. C. & Prasad, A. (2014), 'A confidence metric for using neurobiological feedback in actor-critic reinforcement learning based brain- machine interfaces', *Frontiers in Neuroscience* 8(111).
- Rao, R. P. N.; Stocco, A.; Bryan, M.; Sarma, D.; Youngquist, T. M.; Wu, J. & Prat, C. S. (2014), 'A Direct Brain-to-Brain Interface in Humans', *PLoS ONE* 9(11), e111332--.
- Rosenfeld, M. (2012), 'Overview of colored coins', Technical report.
- Roset, S. A.; Gant, K.; Prasad, A. & Sanchez, J. C. (2014), 'An Adaptive Brain Actuated System for Augmenting Rehabilitation', *Frontiers in Neuroscience* 8(415).
- Scholz, F. & Schröder, U. (2003), 'Bacterial batteries', *Nature biotechnology* 21(10), 1151-2.
- Schwarz, D. A.; Lebedev, M. A.; Hanson, T. L.; Dimitrov, D. F.; Lehew, G.; Meloy, J.; Rajangam, S.; Subramanian, V.; Ifft, P. J.; Li, Z. & others (2014), 'Chronic, wireless recordings of large-scale brain activity in freely moving rhesus monkeys', *Nature methods* 11(6), 670--676.
- Seo, D.; Carmena, J. M.; Rabaey, J. M.; Alon, E. & Maharbiz, M. M. (2013), 'Neural Dust: An Ultrasonic, Low Power Solution for Chronic Brain-Machine Interfaces', *ArXiv e-prints*.
- SFN (2011), 'U.S. Neuroscience Funding Process', Society for Neuroscience.
- STX-Med (2014), 'FDA allows marketing of first medical device to prevent migraine headaches', FDA.
- Thomson, E. E.; Carra, R. & Nicolelis, M. A. (2013), 'Perceiving invisible light through a somatosensory cortical prosthesis', *Nature Communications* 4, 1482--.
- Trimper, J. B.; Wolpe, P. R. & Rommelfanger, K. S. (2014), 'When “I” becomes “We”': ethical implications of emerging brain-to-brain interfacing technologies', *Frontiers in neuroengineering* 7.
- Yoosef, A.; Parasuram, H.; Medini, C.; Solinas, S.; D'Angelo, E.; Nair, B. & Diwakar, S. (2014), Parallelization of a Computational Model of a Biophysical Neuronal Circuitry of Rat Cerebellum, in 'Proceedings of the 2014 International Conference on Interdisciplinary Advances in Applied Computing', ACM, New York, NY, USA, pp. 48:1--48:6.